Global Knowledge ®

Expert Reference Series of White Papers

SDN and Cloud Computing

SDN and Cloud Computing

John Hales, Global Knowledge Instructor,
A+, Network+, CTT+, MCSE, MCDBA, MOUS, VCP, VCAP, VCI, EMCSA

Introduction

In this white paper, we'll review Software-Defined Networking (SDN) and briefly touch on its close cousin Network Functions Virtualization (NFV). After highlighting a few relevant cloud computing concepts and terms, we'll look at the three main deployment scenarios: SDN without deploying cloud computing, cloud computing without deploying SDN, and deploying cloud computing in conjunction with SDN. For each of these three scenarios, we'll look at use cases, when the approach makes sense, and any applicable limitations.

Finally, we'll conclude with the reason SDN and cloud computing are often mentioned together and determine which approach is best for solving various business challenges.

Brief Review of SDN and Cloud Computing Terms

This white paper is not an in-depth review of SDN, NFV, or cloud computing, but rather explores the relationship between them. Therefore, we will not explain them in detail but instead offer a brief review of terminology and concepts to help you understand the entire white paper.

SDN

SDN changes how networking is fundamentally done. Instead of having network intelligence distributed across every device, SDN aims to centralize command and control in a master device (or a few of them for redundancy) and to split networking into three planes, namely:

- Management: This is the interface you use to orchestrate the entire network, specifying the way you wish the network to be run at a high level.
- Control: This is where all the individual devices are directed from using the inputs from the management layer, translating management directives into the actual commands the data layer use to move traffic around.
- Data: This is where the data is actually moved from one device to another.

A major advantage of SDN is that the actual data layer devices can be much simpler and thus less expensive as they don't have to decide what to do with each packet they receive. From a human perspective, each device does not need to be individually programmed. SDN's purpose can thus be summarized as centralized command and control.

NFV

NFV takes the physical networking devices commonly used today (switches, routers, load balancers, firewalls, antivirus, etc.) and virtualizes them in much the same manner as servers. NFV is used to scale out across devices less expensively (scaling by simply adding compute power) and to automatically deploy devices as needed. Thus each project does not require separate equipment or reprogramming of existing equipment. Relevant devices can be centrally deployed via your hypervisor management platform and configured with rules and policies. NFV is almost exclusively used in conjunction with virtualization of servers.

NFV's goal can thus be summarized as automated provisioning of devices.

SDN and NFV Together

While SDN can be used without NFV and vice versa, the real power, especially as it relates to cloud computing, comes when they are used together. When combined you get automated provisioning along with centralized command and control. In the context of this white paper we will combine both under the SDN banner to simplify the discussion.

Cloud Computing

Cloud computing is aimed at self-service provisioning across tenants. A tenant may be a project, department, division, or even a different company. As such, security becomes very important. There are multiple models associated with cloud computing; major categories include:

- Infrastructure as a Service (IaaS): Making VMs available to customers with the physical hardware (servers, storage, and networking) managed by the service provider. A variation of IaaS allows physical servers to be used in place of VMs (called MaaS or Metal as a Service by some, though not officially part of the NIST definition of cloud computing).

- Platform as a Service (PaaS): The development platform for programmers is provided as a service while all the details about the physical and virtual equipment is abstracted from the developer and managed by the service provider.

- Software as a Service (SaaS): An application (such as email or contact management) is made available to customers while all the details of the underlying platform are abstracted from the customer.

There are other models and services that can be deployed in conjunction with cloud computing, but are derivatives of those already listed. This white paper will mostly be addressing the first model (IaaS).

The other aspect of cloud computing is where the resources are located, in other words who is the service provider. The cloud computing infrastructure can be handled within a company by their IT department (private cloud), outsourced to a third party on the Internet (public cloud), or some combination of the two where some resources are at the customer's datacenter and some are hosted on the Internet (hybrid cloud). SDN and NFV can be used with any deployment model, but in the public cloud model really is only relevant to the service provider, not most companies, while in the private and hybrid models, each customer must consider how the company's resources will be allocated, managed, and so forth. That is the primary context of this white paper as it has the broadest application.

SDN without Cloud Computing

SDN can be deployed without cloud computing or virtualization, though this is probably unlikely today. It is most likely deployed at large organizations that have many networking devices and need centralized command and control as well as network omniscience to handle changes in traffic in real time for optimum performance. In this design, provisioning and configuration of servers is under the control of the IT department. While this design automates and controls networking, server provisioning may still be a bottleneck as a user still needs to request a resource from IT and deploy it.

Use cases of this design include companies such as Google (search, not Google Compute Engine—GCE—which is a cloud platform) or Facebook that have their own provisioning process and need a way to scale as new devices are added simply and easily. We as end users do not tell Google to provision new servers when performance is slow. While there are use cases for this design, they are limited.

Cloud Computing without SDN

Cloud computing can be deployed without SDN, and for the last few years this has been a common deployment model. The problem is that while users can quickly request a new VM or even physical server be deployed, finalizing the configuration requires either a pool of networking configurations (such as separate VLANs) that already exist or the process is held up waiting for networking to be provisioned. A third option is that in a private cloud situation security may be less of a concern and thus complete isolation from other servers may not be required.

The most common use case is in private cloud where security is not a key concern or where customers can deploy a limited set of options with security preconfigured. In these cases, SDN is not required (at least from a cloud computing perspective). This is probably the most common scenario today, but there are compelling reasons to move to the next combination.

SDN and Cloud Computing

The final case, in which SDN and cloud computing are used together, is not the most common today, at least for private clouds, but it probably will be over the next few years. The reasons vary depending on your organization (an average company or a public cloud provider).

Service Provider Reasons

Service providers need the ability to deploy resources to companies (cloud computing—that is their business model after all) in a fast, transparent way. They also need the ability to have network visibility across the entire network to handle problems, traffic bottlenecks in real time, and so forth. SDN can provide the needed omniscience and do so in a secure manner. In addition, with hundreds or thousands of customers, the need to deploy virtual switches and routers to link a customer's equipment throughout and between data centers requires quick provisioning and centralized command and control.

Average Companies

Most companies don't have service provider needs, but still may want to offer more options of things that can be deployed (such as firewalls or load balancers) to their customers (end users, developers, email and database administrators, and so forth.). As the catalog becomes larger and the computing resources for a single-use case spread out across one or more data centers, the ability to deploy and manage all of the equipment becomes more challenging. If you then add the needs for security (such as Intrusion Detection or Prevention Systems (IDS/IPS), firewalls, and antivirus deployments) and reporting, the complexity and chances for making a mistake increase exponentially. Further complicating the scenario is the need for load balancing as projects grow and consume more resources, thus the need for rapid scalability and provisioning of devices becomes more pronounced. The need for automated provisioning of virtual networking equipment almost becomes a requirement, and with all the new virtual devices (especially in conjunction with the existing physical devices), centralized command and control becomes a must. This is the sweet spot of cloud computing, SDN, and NFV.

Another reason that SDN and NFV are likely to be deployed is the growing adoption of Agile, DevOps, and other philosophies in the data center to become more nimble and better able to meet customer demand. It is important to link the development environment with the production environment, allowing patches and upgrades to be deployed quickly, so having environments that are as similar as possible becomes a goal. SDN and NFV can help in this regard too, building fenced networks for development, quality assurance, production, and so forth that are very similar, resulting in fewer errors.

Conclusion

As stated previously, while most organizations have not yet fully (or in many cases even partially) adopted cloud computing, to say nothing of NFV and SDN, the trend is growing in all but the smallest businesses. People will go to the cloud and create the VMs needed to get the job done if IT can't do so in a timely manner. Security can be compromised or the company may be slow to adopt change and could get left behind, neither option being good for IT. These changes are coming and to keep your job, you'll need to be familiar with these concepts and adapt or get left behind.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge through training.

Cloud Essentials

Cloud and Virtualization Essentials

SDN Essentials: The Future of Networking

SDN Planning Workshop: Demystifying Vendor Solutions

Visit **www.globalknowledge.com** or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

John Hales, VCP-DCV, VCP-DT, VCAP-DCA, VCI Level 2, is a VMware instructor at Global Knowledge, teaching most of the vSphere classes that Global Knowledge offers, including the Horizon Suite, vCAC, and vC OPS classes. John is also the author of many books, from involved technical books from Sybex to exam preparation books, to many quick reference guides from BarCharts, in addition to custom courseware for individual customers, including Global Knowledge's new SDN Fundamentals and SDN Planning Workshop courses. His latest book on vSphere is titled *Administering vSphere 5: Planning, Implementing and Troubleshooting*. John has various certifications in addition to the VMware certifications previously mentioned, including the Microsoft MCSE, MCDBA, and MOUS, the EMC EMCSA (EMC Storage Administrator), and the CompTIA A+, Network+, and CTT+. John lives with his wife and children in Sunrise, Florida.