Ruckus
Simply Better Wireless.

white paper

Introduction to Cloud Computing

The Future of Service Provider Networks

The service provider industry is in the midst of a major inflection point as networks move away from vertically integrated architectures that bundle proprietary hardware and software together and toward solutions based on cloud computing. These solutions are more flexible, more cost effective, and enable a much faster time-to-market for new applications. Cloud computing makes use of a heavy dose of Software Defined Networking (SDN) and Network Function Virtualization (NFV). SDN is focused on the separation of the control plane from the data plane, while NFV is focused on the decoupling of applications from proprietary hardware by the use of virtualization technology. This concept has already swept through the IT world, and it is now starting to revolutionize the SP industry. Virtualization technology enables a wide variety of SP applications to run on commodity hardware in a much more cost efficient manner.

Cloud computing isn't just another name for virtualization, but it does build on virtualization. The following definition of **Cloud Computing** comes from NIST.

> *"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."*

From the Ruckus and broader industry perspective, cloud services almost always involve some form of hardware virtualization. The latter refers to the process of abstracting the hardware layer (compute, storage, and memory) from the operating system and the application that runs on that OS. The service provider industry has picked up on this concept and extended it to enable Network Function Virtualization (NFV). The software layer that enables NFV is called a hypervisor.

Network Function Virtualization (NFV) offers a new way to design, deploy and manage networking services. NFV decouples the network functions such as Routing, WLAN Control, GGSN, NAT, etc., from proprietary hardware appliances, so they can run in software. NFV is designed to consolidate and deliver the networking components needed to support a fully virtualized infrastructure including compute, storage, and memory. It utilizes standard IT virtualization technologies (aka hypervisors) that run on commodity hardware to virtualize network functions. It is applicable to both data plane and control plane functions.

A **hypervisor**, also known as a virtual machine manager, is a program that allows multiple operating systems to share a single host. Each operating system appears to have the host's compute, storage, and memory all to itself. However, the hypervisor is actually controlling underlying resources, allocating what is needed to each operating system in turn and making sure that each operating system cannot disrupt the others. This is done using the concept of a virtual machine.

The **virtual machine (VM)** provides the virtual compute, storage, and memory resources required for the user's application. In a managed service environment each customer can be assigned their own instance of an application that runs on their VM. Another option would allow multiple customers to share an instance of an application by operating in a multi-tenant environment within a single VM. The latter is a much more cost-effective way of deploying virtualized services.

In Figure 1 we see the hypervisor software sitting atop physical resources that include compute, storage, and memory. The hypervisor creates a series of virtual machines and each virtual machine can provide services for a different application. The applications and OS that run on a VM have no idea that they are actually running on a hypervisor, it is all abstracted away.

The following are the most popular hypervisors in the industry:

- The VMware vSphere hypervisor is the clear leader in this market. Most of VMware's deployments are in private clouds for large enterprises. (65%)
- Microsoft has a hypervisor called Hyper-V (25%)
- KVM (kernel-based virtual machine) is an open-source hypervisor with a lot of industry traction. Distributions are available from 3rd parties that will provide support such as Red Hat and many others. (8%)
- Xen and Virtual Box (4%)

Cloud Services

There are several different instantiations of a cloud service including public clouds, private clouds, and hybrid clouds. A **public cloud** is a service that allows customers to share network infrastructure. The customers pay for compute capacity (defined as CPU, memory, and storage) on an as-needed basis. If the customer's application requires greater amounts of processing power, the hypervisor can dynamically provide this resource and the user is charged on a usage basis. As the user's demand drops, the hypervisor will release those resources. As long as the underlying resources are available, the user can access more and more compute capacity, as their circumstances require. Public clouds are usually the domain of cloud service providers.

In some situations, it might not be acceptable to assume that additional compute capacity will be there when needed. In these cases a **hybrid cloud** could be provided which sets aside resources in a public cloud for a specific customer. This approach will be more expensive, but it does guarantee that those resources will be there when required. The third option is a private cloud. In this situation the hardware resources are in the data center of a specific company and those resources are only for the benefit of the users of that company. This means that various departments can access resources as required based on their needs.

FIGURE 1: Network Function Virtualization

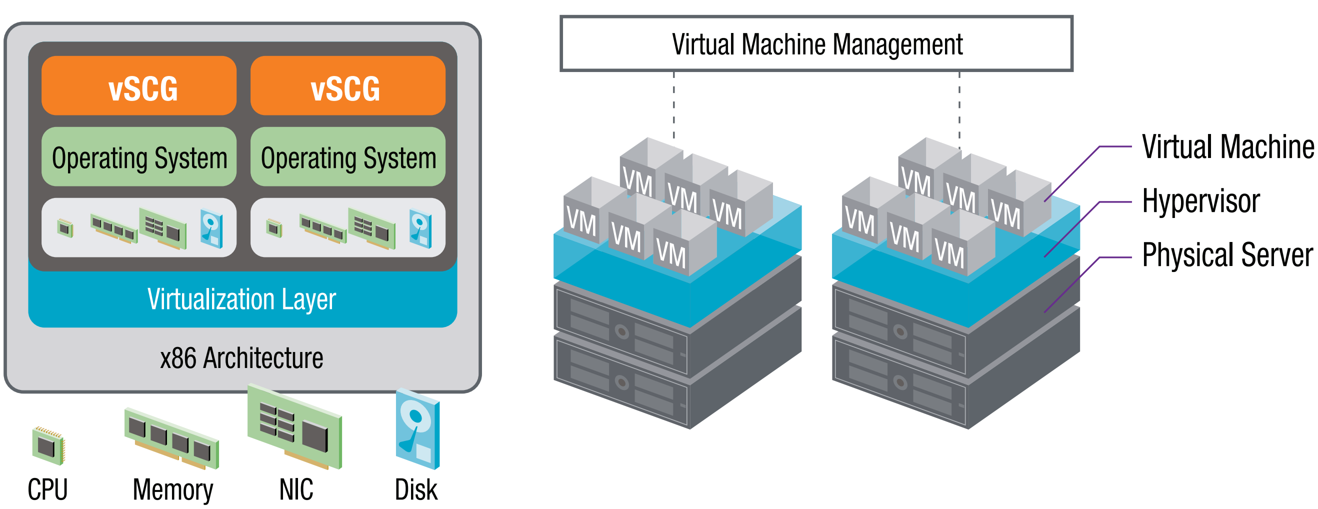

The key capabilities of a cloud computing architecture include:

- **Pricing** – allows the user to pay for only the compute resources that they need. As their usage goes up so does their cost.
- **Resource abstraction and pooling** – In a cloud environment all customers share compute capacity, but are unaware of this arrangement as it is abstracted away by the hypervisor. They just know that they get capacity when they need it.
- **Network centric** – Cloud services sit in a data center and are accessed via the network.
- **Simple, fast provisioning of resources** – It becomes possible to very quickly turn-up new applications in this environment.
- **Rapidly and elastically provisioned resources** – allows applications to dynamically acquire more compute resources as their circumstances dictate. This is one of the most compelling features of a public cloud, as it allows applications to easily deal with sudden bursts of traffic without always having to plan resources for maximum load.

A Public Could Service can be offered in one of three different ways:

Infrastructure as a Service – a service made popular by Amazon Web Services (AWS) and others. It includes the hardware infrastructure and virtualization layer (aka the hypervisor). The users of the service runs their OS, application, and libraries on top of the service provider's hypervisor. The OS will make requests to the hardware layer for support. The hypervisor intercepts these requests and then relays them to the underlying hardware elements. Multiple users can run their software on a host at the same time with no knowledge of the others. The hypervisor shares the available hardware resources as required by the various applications.

Platform as a Service – provides an additional layer of services a top an IaaS infrastructure, which typically includes middleware such as application servers and databases.

Software as a Service – has been made popular by applications from companies like Intuit and many others. SaaS is almost always used in B-C applications, where the user is able to access and use a particular application via the Internet.

A Deeper Look @ Network Function Virtualization

Network Function Virtualization is a key building block in a cloud computing architecture. The following figure looks at NFV in more detail.

In Figure 2 we see the ETSI (European Telecommunications Standards Institute) view of network function virtualization. We'll start by looking at the virtualization layer in the lower left corner. This is where we find the hypervisor, which sits atop the physical infrastructure that includes compute, storage, and

FIGURE 2: Network Function Virtualization

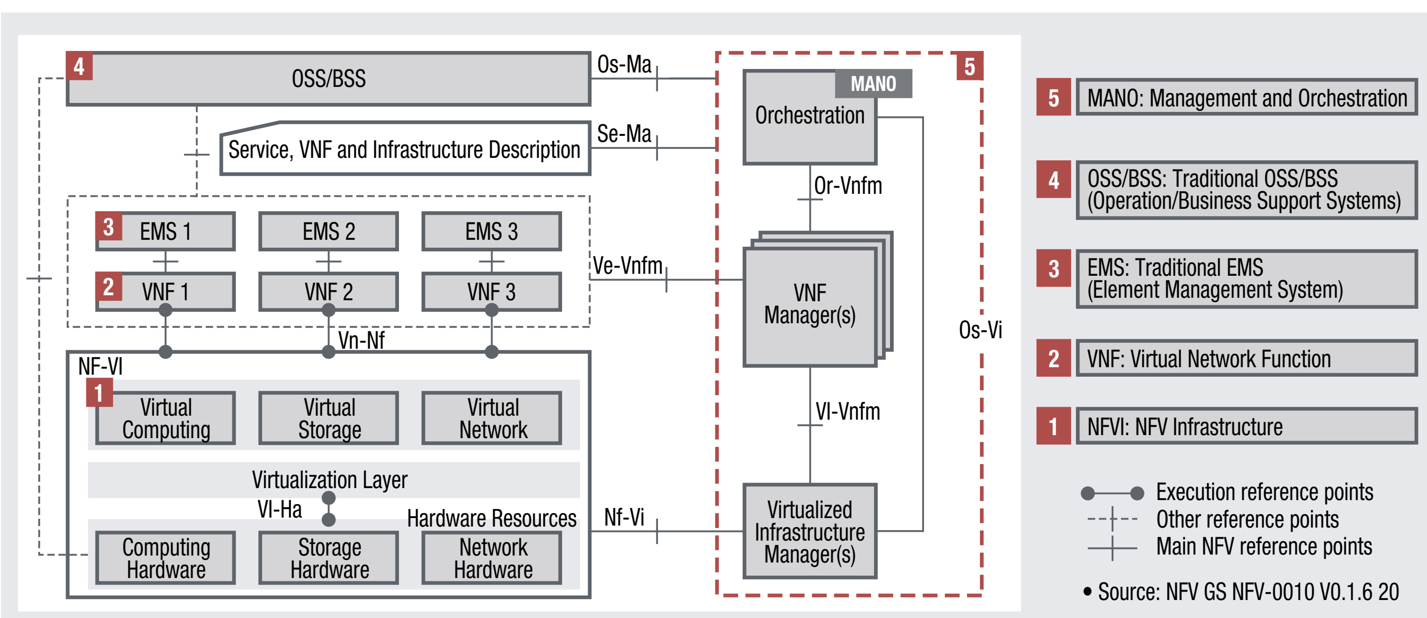

FIGURE 3: OpenStack as the Basis of a Cloud OS

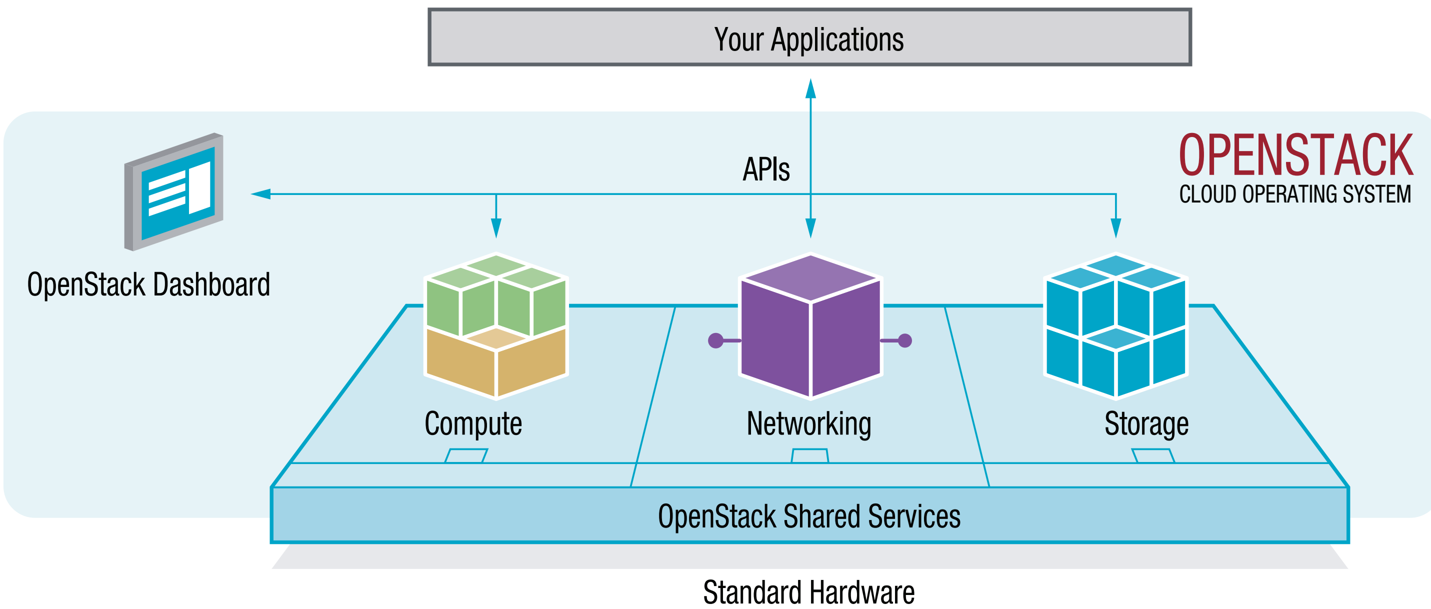

memory resources. The virtualization layer must be managed and to that end we have the virtualized infrastructure manager. The open-source community has kicked-off a major effort here called OpenStack to do the virtualization layer management. OpenStack will often be deployed with KVM, which is an open-source hypervisor, and together they can be thought of as a cloud operating system. Proprietary solutions from vendors like Microsoft and VMware will implement something similar for the virtualization layer management function in their architectures. Sitting a top the hypervisor is the customer's application and OS. These applications will typically have their own EMS systems, which can in turn work with upstream OSS/BSS systems. The virtualized applications also need a manager and in this drawing it is referred to as the Manager of Managers. It provides management, orchestration and provisioning of the underlying applications. This would, amongst other things, include a coordination process when service chaining is involved. The Manager of Managers function would typically be provided by a 3rd party that would make use of APIs to manage a variety of different applications from different vendors.

OpenStack can be thought of as an open-source cloud computing operating system. Predominantly acting as an Infrastructure as a Service (IaaS) platform, it is available under the terms of the Apache License. The project is managed by the OpenStack Foundation, a non-profit corporate entity established in September 2012 to promote OpenStack software and its community. There are many OpenStack distributions.

Practical Aspects of NFV

A service provider who follows the NFV design principles will implement one or more Virtualized Network Functions (VNF). A VNF by itself does not automatically provide a usable product or service. To that end, it is necessary to introduce the notion of service chaining, where multiple VNFs are used in sequence to deliver a service.

NFV presumes and emphasizes the widest possible flexibility as to the physical location of the virtualized functions. This means a service provider could locate NFV in all possible locations including national data centers, regional data centers, the customer premises, and anywhere else that makes sense.

Software Defined Networking

SDN, or Software-defined networking, is a concept related to NFV, but they refer to different domains. SDN is focused on the separation of the control plane from the data plane, while NFV is focused on decoupling network functions from the underlying compute layer.

Figure 4 shows how SDN would relate to NFV. An SDN controller would be used to control the underlying forwarding (aka data) plane while also chaining together various functions to provide a service. It could operate under the control of a Cloud OS like OpenStack. The protocol that would be used between the SDN Controller and network forwarding elements could be OpenFlow, or it could be something else. SDN's impact on Wi-Fi access networks is still to be determined. However, some of the concepts are certainly applicable such as the handling of all control plane functions by the vSCG, while all data plane functions are handled by a separate forwarding plane.

FIGURE 4: Software Defined Networks and NFV

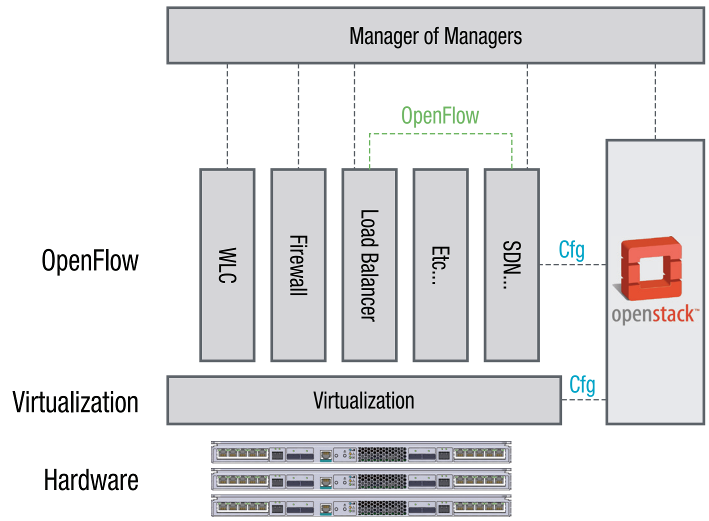

Ruckus Wireless, Inc.
350 West Java Drive
Sunnyvale, CA 94089 USA
(650) 265-4200 Ph \ (408) 738-2065 Fx

Ruckus Simply Better Wireless.

www.ruckuswireless.com