White Paper

THE INS AND OUTS OF CLOUD COMPUTING

and Its Impact on the Network

April 2010

Rev. A 04/10

SPIRENT

1325 Borregas Avenue
Sunnyvale, CA 94089 USA

Email: sales@spirent.com
Web: http://www.spirent.com

AMERICAS 1-800-SPIRENT • +1-818-676-2683 • sales@spirent.com

EUROPE AND THE MIDDLE EAST +44 (0) 1293 767979 • emeainfo@spirent.com

ASIA AND THE PACIFIC +86-10-8518-2539 • salesasia@spirent.com

Spirent Communications, a leader in networks, services and devices testing, offers Spirent TestCenter™ Virtual, the industry's first solution specifically designed to holistically validate the performance of all elements of the data center and cloud computing environments including virtual machines, servers and storage devices.

The Ins and Outs of Cloud Computing

and its Impact on the Network

CONTENTS

INTRODUCTION: DEFINING THE 'CLOUD'

Ask 10 people to define cloud computing, and you'll likely get 10 different answers. But while the technical intricacies might seem complex, the basic notion behind cloud computing is fairly simple: It's a new IT provisioning and support model that provides on-demand network access to a shared pool of computing resources. That shared resource pool is called a "cloud."

The cloud can be located in one or more private data centers, in facilities operated by a third-party service provider, or across a combination of the two. From the cloud, users dynamically access compute cycles. Users experience what seems like a limitless supply of application, storage, network and compute resources that are simply "out there" in cyber space, available on demand, without having to explicitly request, purchase, install and provision them. There are economic and productivity business drivers that make it desirable for users to be able to tap IT resources dynamically in this way and to pay only for what they use.

Internal IT staff or a third-party service provider manages the pool of hardware, software and operational resources behind the scenes, so that capacity continually scales. Key technology enablers help them and play a significant role in maintaining the elasticity of the cloud. Among them: virtualization software and hardware, broadband networks, and cloud management and benchmarking tools.

All this sounds appealing, yet there are network performance and security ramifications associated with cloud computing. Enterprises, government agencies and service providers that operate cloud environments – as well as equipment vendors that build the network infrastructure gear on which clouds run – must take these considerations into account in their designs. Technology enablers such as cloud performance and security testing tools can help them do that, as well as help them measure whether the metrics and results expected from the cloud are actually being delivered.

The remainder of this paper will describe the drivers behind cloud computing; the various cloud types and services; and some of the network performance, availability, security, and scalability ("PASS") challenges associated with overhauling the traditional data center into a cloud model.

DRIVERS BEHIND 'CLOUD MANIA'

The primary goal of most corporate and government entities investigating cloud computing is economic: that is, they want to lower their data center total cost of ownership (TCO). Moreover, in the process, they also hope to improve user IT service levels and reduce provisioning times in an effort to boost user productivity and overall time-to-market competitiveness. Dramatic improvements in each of these areas are being made possible by virtual machine (VM) technology and by new usage-based billing models that work in conjunction with the cloud.

Equipment and Power Savings

Historically, keeping up with IT and Web demands has called for new capital infrastructure investments. But having to continually buy additional equipment to deploy a new application is growing cost-prohibitive, particularly given that most servers operating in a traditional manner have fairly low (often under 20 percent) utilization rates and waste money and energy. This method of deploying a new resource also delays its availability as well as user productivity and overall organizational competitiveness while users wait for servers, applications, storage drives, and network connections to be provisioned.

Cloud environments, by contrast, enable enterprises to quickly and easily tap computing, storage, software, development, and network resources on demand. This elasticity is often made possible through the use of virtualization technology, which decouples operating system and application software from an underlying hardware platform. Virtual software can be moved around the cloud from server to server – or even from cloud to cloud – to wherever the compute cycles happen to be available on the physical infrastructure at the moment.

This setup makes consolidation of servers and, in some cases, even entire data centers possible. Utilization rates of physical servers can be pumped up to 65 percent or greater, significantly reducing the number of them needed, as well as the real estate required to house them and the carbon footprint needed to power and cool them.

Cloud Characteristics & Attributes

- Dynamic user self-service
- Broad network access
- Resource pooling for efficiencies
- Rapid elasticity/service provisioning
- Measured services; usage-based billing

Pay-As-You-Go Billing

Another characteristic of cloud computing is its utility billing model. With internal chargeback systems or service provider billing systems, cloud services can use a mix of flat monthly subscription fees and/or usage-based charges. With many third-party services, enterprises and government agencies can offload the capex burden from their internal budgets and instead "pay as they go" for what they consume. This model is similar to how consumers today pay the public water and electric utility companies for their usage each month without negotiating binding contracts and complex service-commitment terms.

Virtualization Technology

While cloud computing and virtualization technology are not the same thing, they are often closely related. It is technically possible to build a cloud without the benefit of virtualization technology, particularly in the case of software as a service (SaaS), one of three primary types of cloud services described in more detail later in this paper. However, the flexibility and automation associated with virtualization are key enablers for creating highly scalable, dynamic cloud environments that meet enterprise-class needs.

Virtualization decouples software from physical hardware, and it does so by using a special piece of software often called a hypervisor, or "virtualization layer." This virtualization software or firmware usually sits between the physical machine and the operating system and allows a single physical server (or other hardware such as a storage drive or piece of network equipment) to appear to be many separate devices.

There are ways to scale services in traditional data centers today, such as load balancers (often called application delivery controllers) that dynamically direct user application requests to the servers best able to handle them. However, many organizations would be hard-pressed to create and maintain clouds, which promise the elasticity to dynamically scale up and down, without virtualization, and the related management and orchestration tools that automate how, when, and where to move VMs around to meet compute demands.

Similarly, data center networks and network gear with more, high-capacity ports play a role in providing the constant realignment of network resources, as theyare accessed increasingly across LANs, SANs, and WANs, and workloads are balanced among VMs across multiple servers.

There are a number of reasons, that cloud computingis currently in vogue, Specifically, this model can:

- Reduce capital expenditures
- Save on recurring costs by using pay-as-you-go billing models instead of lengthy contracts
- Cut or eliminate provisioning times, allowing users to be as productive and competitive as possible
- Decrease power consumption and its associated carbon footprint and monthly utility bill

PRIMARY CLOUD SERVICES

There are several types of clouds (see box), as well as three primary cloud services as defined by the National Institute of Standards and Technology (NIST). These cloud services are infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS).

Let's take a quick look at each.

Infrastructure as a Service (IaaS)

Generally offered as an unmanaged service, IaaS consists of server, storage, and local network resources made readily available by a third-party on an outsourced basis. Services are billed based on the level of each resource actually consumed during the billing period.

The provider gives a business customer a generic hardware foundation on which the business customer installs and runs its own operating systems, applications and storage data. Business customer employees then simply dip into the resources in the cloud. This approach saves enterprises from making upfront and ongoing investments in the underlying foundation of servers, storage devices, and local network connections as the business grows and requirements change.

Platform as a Service (PaaS)

This type of unmanaged cloud service serves the Web application development and hosting community, including enterprise programmers, large independent software vendors (ISVs) and entrepreneurs. The tools, application programming interfaces (APIs), protocols, operating system platforms, and storage necessary for developers to create and host new Web applications are available with PaaS.

There are four IT cloud models:

- Private Cloud
  A resource pool shared within a single organization The private cloud can be delivered in two ways: 1) an enterprise or government agency could build its own cloud and host and manage it internally; 2) the organization could use an outsourced service, whereby a third-party houses and manages resources on physical infrastructure components dedicated to that organization only.
- Public Cloud
  Offered as an outsourced service by a third-party. Business customers share physical infrastructure for optimum economies of scale.
- Hybrid Cloud
  A combination of private cloud(s) and public cloud service(s). The private/public mix can change dynamically as customer requirements change.
- Community/Federated Cloud
  The sharing of private clouds among entities with similar interests. Within a single organization, for example, engineering could potentially "burst" into the production IT cloud to temporarily use services at peak times. Or cloud resources could be shared across organizations; for example, public safety personnel could share emergency-response clouds for disaster relief efforts.

Consumers of PaaS also pay by usage, circumventing massive upfront capital and software licensing investments that can be a barrier to innovation. Developers can also store their code and data in the cloud, using and paying for only as much storage space as needed, and then using the cloud as a large-scale channel for distributing their software to consumers.

Software as a Service (SaaS)

In this type of cloud service, software applications become available in the form of a network-based service, allowing users to access them across a WAN. SaaS is an alternative to consumers or businesses buying software licenses, loading application software on an internal computing infrastructure and remaining in charge of software patches and version updates.

In other words, SaaS is a low-cost way for businesses to procure rights to use software as needed by using on-demand licensing. Businesses get the benefits of commercially licensed use without the associated complexity and potential high initial cost of equipping every device with application software.

SaaS is usually considered a fully managed service. The provider maintains both the software and hardware environment on behalf of the customer. Enterprise IT personnel are relieved from having to do any software customization or version tracking, patching and updating.

SERVICE MODEL ARCHITECTURES

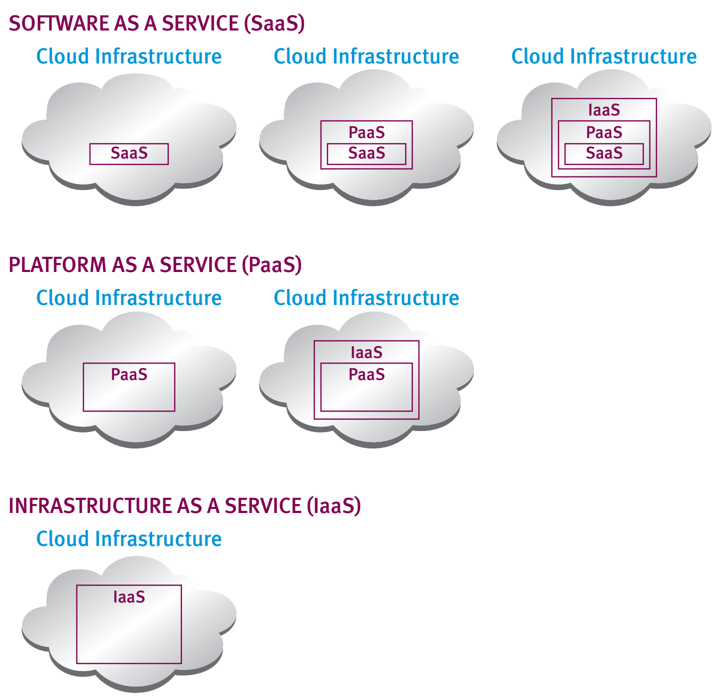

Cloud services are interrelated in that some types use the others as foundational building blocks.

IMPACT ON THE NETWORK

It goes without saying that LANs, SANs, and WANs, including the Internet, are likely to require upgrades to support the increase of elastic and real-time network traffic. The move to dynamic data centers with virtualization, places greater demands on the network infrastructure.

Traditionally, a single physical server connects to a switch access port. In a virtualized cloud environment, however, potentially more than 64 virtual servers, including a virtual switch with as many virtual ports, reside on a single physical server and use that single switch access port. In addition to the increased complexity, typical access network over-provisioning ratios of 1:20 are changing to 1:5 or higher performance with higher utilization achieved through virtualization. As a result, network access, aggregation, core, and equipment backplanes must scale to accommodate the condensed traffic loads.

Without a properly virtualized environment, the network just sees a physical connection to a computer or a server and doesn't have the ability to see individual VMs on that computer or server. As a result,it is important that networking equipment makers innovate or establish partnerships to make the capabilities of their physical networks available within virtualized clouds. For example, appliances are appearing with virtual service control planes for provisioning VMs.

The industry is probably most familiar with the practice of network virtualization, which allows for segmentation and partitioning for security and quality of service (QoS) reasons. Partitioning in a VLAN or virtual private network (VPN) is done primarily to keep customer resources segregated and confidential. Network segmentation can also take place to give certain resources higher priority – such as real-time applications with low tolerance for delay and jitter.

CLOUD CHALLENGES: PASS

IT clouds can affect the fundamental and critical data center and network attributes of performance, availability, security, and scalability, known simply as "PASS." Let's take a brief look at each.

Performance

There are several factors to consider regarding performance in a cloud environment, particularly one that is virtualized. The network, as mentioned, comes into play, in that it requires the capacity to handle dynamic access across a large number of VMs. The network must be able to recognize VMs and be in a position to shuttle data back and forth between VMs and users, regardless of where the VMs and users are located at the moment.

Traditional over-provisioning methods of fixed resources – physical servers, storage drives, network switches–no longer apply in the virtualized environment. The cloud designer must take this into account by ensuring enough VM instances are provisioned to make dynamic access possible for all users .

Another factor to consider is the impact of new security mechanisms, such as virtualized firewalls (described in the upcoming "Security" section) on compute performance.

Availability

The traditional ways of providing local redundancy must also be reconsidered in a virtualized cloud environment. Today's blade servers often integrate LAN and SAN network interfaces (and sometimes even network switches), into single devices that can support 1,000 or more VMs. These physical foundations for VMs can become a single point of failure if appropriate approaches to VM load balancing, automated resource scheduling and live migration to other hardware are not built into the design.

Security

Traditionally, data center designers have positioned security appliances, such as firewalls, intrusion detection systems, and VPN termination systems, in strategic physical locations, such as at the WAN edge where requests and traffic from the Internet can be filtered and decrypted. However, geographic locations of physical servers have less meaning in a virtualized cloud, as users might be tapping resources from VMs located on one of any number of servers or even data centers.

Moving to a virtualized data center requires two-tier approach to security. The physical server hosts must be secured, as do the VMs. In a virtual cloud environment, security teams will need to deploy security tools that account for devices that might change IP addresses frequently in response to shifting workloads. This might ,for example, result in creating virtual firewall and intrusion detection instances that run alongside each VM and move with the VM as it moves.

Scalability

The promise of cloud elasticity with infinite scale may sound realistic in theory. However the elasticity of the physical infrastructure has finite limits. Addressing this risk requires not only having spare room for additional servers, storage, power, and cooling capacity. It also requires a well-thought-out network infrastructure where aggregation and core interconnects do not become the bottlenecks of the elastic demand and scale that the cloud promises.

IMPORTANCE OF CLOUD TESTING

Test and measurement solutions have evolved to better assess the performance, availability, security, and scalability of cloud computing with the addition of virtual test appliances that open up visibility of each of these attributes throughout the entire data center infrastructure. Cloud test solutions including a virtual appliance are important to :

- Holistically validating the performance of all elements of the data center and cloud computing environment
- Benchmarking the performance of new virtual switches, firewalls, load balancers and other virtual appliances
- Quickly determining which virtual or physical component(s) is impacting performance
- Measuring the impact of application reliability on dynamic resource scheduling with live migration of virtual servers
- Testing the vulnerability of virtualized security devices by emulating real-world attacks and threats

CONCLUSION

There are many compelling reasons for enterprises and government agencies to build or buy cloud services, which represent a whole new computing model and fundamentally change many aspects of IT services.Key advantages of this environment include lowering the total cost of ownership of a data center and gaining dynamic access to resources as needed, without the laborious provisioning wait times that impede worker productivity and organizational competitiveness.

However, in moving to the cloud, there are ramifications for the network, which is now required to deliver better and faster services. Only with cloud-optimized test and measurement solutions can operators and network equipment makers gain visibility into the cloud and measure its impact on the fundamental PASS attributes: performance, availability, security, and scalability.

SPIRENT®