

Into the Clouds...

A Look into What Cloud Computing is
and How it May Affect Data Network Cabling

White Paper — June 2010

Contents

BACKGROUND	3
------------	---

IMPACT ON IT	4
SCALABILITY/FLEXIBILITY/ADAPTABILITY/AGILITY -	4
SELF SERVICE AUTOMATION -	4
MEASURED SERVICE -	5
SECURITY -	5

IMPACT ON NETWORK DESIGN AND INFRASTRUCTURE CABLING	6
INSIDE THE DATA CENTER -	6
OFFICE LEVEL NETWORKS -	6
WIRELESS INFRASTRUCTURE -	7

SUMMARY	7
---------	---

Clearly the latest marketing focus of Information Technology related products is the “cloud”. But what does this nebulous term mean in regard to today’s IT environment, and what, if any impacts will this have on your layer one physical infrastructure?

Background –

To better understand cloud computing, it is best to look at the characteristics of cloud computing. The principal characteristic that defines a cloud computing environment is the fact that the application being utilized is delivered over the network, instead of being locally processed. An example would be using a browser to connect to a service over the Internet to create a written document, such as this article. Traditionally we have used word processing applications running on a local compute resource to achieve this capability. In the cloud model, this local processing is pushed to a compute resource located in the “cloud”, whose location and underlying infrastructure is of little to no interest to the user of this resource. The only real change to the end user is the methodology used to access this capability.

In the interest of historical accuracy, the above scenario is really describing the methodology utilized since the dawn of the PC or Personal Computer age and the extensive networking of those compute resources. In reality, prior to the PC era, document creation was typically done on a central resource likely running on a data center computer and being accessed by a remote user through some means of connectivity such as a video terminal. But, this older scenario is markedly different from the modern implementation of cloud computing.

It is also different from the evolution from centralized processing in that terminal accessed application running on a central processor, to the popular client/server model which grew out of this scenario when the PC era began. Many applications today still utilize this client/server design.

But if a client/server application involves a local user accessing a remote computer doing the “heavy lifting” from a compute standpoint, what is different about the cloud computing environment? In cloud computing the key difference is the “client” end becomes a standardized application, namely a browser. This opens the door to classes of devices to utilize powerful applications that have never had the native compute or storage capacity to perform those tasks. This would be devices like the Personal Data Assistant (PDA) or the pervasive “smart phone” which is becoming a globally available tool for the masses.

So now that we have seen the difference in cloud computing vs. traditional access methods, it is important to see the impact of cloud computing to the people who have historically provided such services, namely the Information Technology staff of your organization.

The “IT staff” at any given company might range from a single individual who was “handy” with computers, to a professional who was consulted when necessary to set up services or help alleviate problems when they occurred, to a full blown division of your company dedicated to the delivery of all information related activities. Cloud computing will have a profound impact on this IT function regardless of company size.

Impact on IT

Arguably the biggest impact of cloud computing will be on the IT staff and how it delivers information services to the end users of those services. As stated previously, IT provided services by creating application servers running in a company owned data center, a collocation data center, or leased those same compute and storage resources from a data center hosting provider.

To meet business application requirements, IT staff would receive a request from a business unit to provide a specific function. IT would then determine a suitable application, determine resource requirements and then either implement this application if resources were available, or move to acquire those resources and install and configure them, finally providing those services back to the business unit that generated the request.

In the cloud environment, the physical resources that IT uses to meet these demands are abstracted from traditional paradigms. The resources may indeed be locally owned assets, but just as likely could be provided by a ubiquitous source.

Indeed there are four characteristics of cloud computing that define the experience. These traits differentiate cloud computing from previous incarnations of application delivery. Essentially these are:

- 1) Scalability/Flexibility/Adaptability/Agility
- 2) Self service automation
- 3) Measured service
- 4) Security assurance

Let's look at each of these a little deeper.

Scalability/Flexibility/Adaptability/Agility –

Cloud computing assumes an on demand style of delivery. Many applications have peaks and valleys of usage. To handle this, most traditional IT implementations have relied on over provisioning to accommodate the peaks. This leads to inefficiency as most assets are only partially utilized. The goal of cloud computing is to create a pool of resources that are flexible enough to handle many different sorts of applications and that can be brought online or torn down (scalable) to meet demand. This flexible design is commonly based on virtual computing although virtualization is not an absolute requirement.

Self service automation –

The second key differentiator of cloud computing is the concept of self service. This is markedly different from traditional application delivery. The paragraphs above detailed how traditional services are delivered to an organization. This process can take several weeks or even months to complete. The idea in cloud computing is that application start up can be reduced to as little as minutes and certainly no more than a matter of a couple of days, depending on circumstances.

How is this possible? In our cloud delivery scenario, again we have created large pools of compute resources that are capable of delivering any number of applications “on demand”. So, when a department or individual needs to bring up an application, a properly designed cloud offering will allow for easy self service provisioning. This is especially useful as a sort of “test drive” of an application prior to spending the significant resources and time investment required to bring a new application on line with traditional techniques.

This on demand delivery of information services is analogous to modern video on demand offerings contrasted to traditional fixed program scheduling. Today's users require not only access to the data they need, but they require access when it best suits their business needs, not when it is convenient for IT to deliver it.

Other benefits are quick start capabilities. An application can be brought online quickly with a cloud provider while the lengthier process of traditional in house support is brought online. And finally, the availability of on demand capacity can allow for peaks to be handled by cloud resources allowing additional capacity when needed without expensive over provisioning.

In general, service automation will require comprehensive toolsets that allow for virtual server deployment and application delivery. Along with the tools to manage the server and application environments, it will be necessary to automate network and storage provisioning. In all of these areas, infrastructure intelligence will play an increasingly important role. Intelligence will be a critical enabler for successful cloud deployments, not only for service provisioning, but also for security and compliance requirements..

This self provisioning aspect and quick delivery of requested services is the fundamental basis of the great agility associated with cloud computing.

Measured service -

In order to provide on demand types of service, it is necessary to be able to measure what is being delivered. This might be in compute cycles or storage utilization for example. One such common metric used today by cloud providers is called VM hours, which is a measure of the hours per virtual machine utilized with the cloud offering. An effective charge back mechanism is essential to any cloud computing environment. It may have been a desirable feature with traditional application delivery methodologies, but it becomes a central component of the cloud environment.

This measurement should also include performance metrics to allow for proper Service Level Agreements to be put into place. This protects the end user and provides the service provider the means to demonstrate SLA compliance.

Security –

Any application delivery must also ensure proper security for end user information. This is not unique to cloud computing, but since public clouds will involve sharing of common resources between many clients, some of whom may even be competitors; it becomes even more essential that proper security measures can be put into place and adequately demonstrated. Without this foundation, the likelihood that a customer would adopt a cloud computing solution is greatly diminished.

Impact on network design and infrastructure cabling

Now that we have covered a little of the background behind cloud computing, we will discuss some of the implications this new architecture may have on your underlying physical infrastructure.

Inside the data center –

First of all, it is highly likely that any large cloud computing infrastructure will be based on a high degree of virtualization. This concept of virtual machines not being tied to physical servers allows for the flexible scaling required in the cloud paradigm. Along with virtualization comes the requirement for flexible and low latency network architecture. This is leading many data centers to move to a two layer design instead of the traditional three layer core/distribution/access tiers that have characterized data centers for decades. The primary reason for this is to allow for very low latency between physical servers and their corresponding virtual server instances.

Since virtualization can now support many virtual servers in a single physical host, it is highly likely that average input/output (“I/O”) levels will significantly increase from traditional levels. Today this is driving adoption of 10 Gbps Ethernet connectivity to the server, with 40 Gbps Ethernet to those same highly virtualized hosts to follow. Of course, with increasing server I/O demands, there is a corresponding increased requirement for higher and higher uplink capability from the access switches delivering this ever increasing capacity to the core switching platforms. This will drive the adoption of 40 and 100 Gbps Ethernet ports in the data center environment. Initially 40 Gbps uplinks to support 10 Gbps access capacity to hosts were used, but eventually 40 Gbps to the host with corresponding 100 Gbps uplinks to the central switching/routing fabric will be used.

This demand for higher I/O speeds requires higher performing cabling to support the infrastructure. This is why the recommendation for copper horizontal cabling in a data center is moving from Cat 6 to Cat 6A in current standards revisions. In addition fiber recommendations are moving to OM4 multimode fiber for most fiber applications.

In addition to the move to OM4 on the fiber side, it is highly recommended by CommScope to move to preterminated fiber infrastructures as well. This is because 40 and 100 Gbps will primarily be delivered across parallel fiber optic links, which are characteristic of preterminated fiber solutions. In addition to the proper Multi-fiber Push On (MPO) based connectorization, preterminated fiber will dramatically reduce installation time while simultaneously allowing for easy and rapid redeployment of fiber infrastructure when reconfiguring or refreshing a data center layout.

In the case of a private cloud infrastructure, traffic patterns within the data center are also likely to change. Since most applications are now being delivered via the network, network traffic will be shifted to the corporate Wide Area Network (WAN), potentially requiring beefed up bandwidth between that data center resource and the end user. That user may be accessing the data center cloud across normal corporate network links, but could also be using a mobile device for this purpose. In those cases it may be necessary to scale up Internet access for these mobile workers.

Office level networks –

This network delivered application implementation will certainly increase Local Area Network (LAN) traffic levels in the typical office environment that is leveraging cloud services. Moving the compute and storage requirements to the network increases not only the level of traffic, but also the requirements for reliable delivery. Packet loss in an Ethernet environment can have a dramatic, negative effect on application performance. For many applications this was a tolerable condition, but as more applications are delivered by the network, there will be more and more demands placed on the physical infrastructure to deliver that data reliably.

Desktop virtualization is another aspect of virtualization that many equate to a cloud environment. Again this type of application delivery will also increase the strain on network infrastructures and only emphasize the need for a quality underlying physical transport to insure optimal employee productivity.

Wireless infrastructure –

Keeping in mind that one of the points of the cloud computing environment was to better enable employee mobility, it is important to not forget the wireless environment when deciding to move to a cloud delivery paradigm. It may be possible in smaller organizations to limit employees to a single carrier for mobile devices, but this is usually impractical in larger organizations.

To ensure uniform performance and application availability, it will become increasingly important to make certain that wireless networks are extended throughout your buildings in a carrier agnostic manner.

Summary

Cloud computing is clearly one of the hottest topics today in IT. There are compelling financial justifications for exploring cloud computing in nearly any size organization. From rapid deployment to try-before-buy to disaster recovery to demand scaling to mobile application delivery, there are many advantages of the cloud computing architecture.

Along with these advantages, there will be architectural differences in how applications are delivered to the end users. Cloud computing will have a profound impact on how IT delivers information services to the rest of the company. In many cases this application delivery will shift outside of traditional, local data center delivery, to external cloud based resources. In other cases IT will develop internal clouds to enable faster service delivery and business agility to the many business units they support.

Along with these changes, it is inevitable that data patterns and traffic levels will also shift and the network which is already essential to business success will only become that much more important moving forward. It is imperative that network performance and reliability take on renewed emphasis to guarantee application delivery.

Management of this highly virtualized infrastructure will need to be highly automated and this will drive the need for better management tools and visibility into the physical infrastructure through the use of intelligent offerings.

While the outlook for application delivery may be “partly cloudy”, the future is looking very bright!



© 2010 CommScope, Inc. All rights reserved.

Visit our website at www.commscope.com or contact your local CommScope representative or BusinessPartner for more information. All trademarks identified by ® or ™ are registered trademarks or trademarks, respectively, of CommScope.

06/10