Testing the Cloud

Mark Sylor, New Technologist EXFO Service Assurance

INTRODUCTION

Broadband, backbone and mobile wireless service providers see the cloud as an opportunity for growth. New applications running in the cloud drive new traffic to their network. The cloud also offers new revenues from new services that can be sold to customers. The business models for carriers and how they relate to the cloud are evolving quickly, but one fact is becoming clear: to achieve success in the cloud market, carriers must actively ensure that they offer a high quality of service to cloud consumers. The best way to assess cloud quality is to test the cloud.

A cloud is a shared computing platform available over the network used to run a variety of business or personal applications. The concept is hardly new; it has roots in service bureaus, outsourced data centers and utility computing. What makes the cloud work today is the rise of the web browser as a thin client that allows individual users to run any application, the wide availability of high-bandwidth networks, and virtualization technologies for computers, storage and networking. The cost savings, expanded reach, and improved quality gained by running an application in the cloud is proving to be a business success, as shown by the growth of cloud services into a market worth billions.

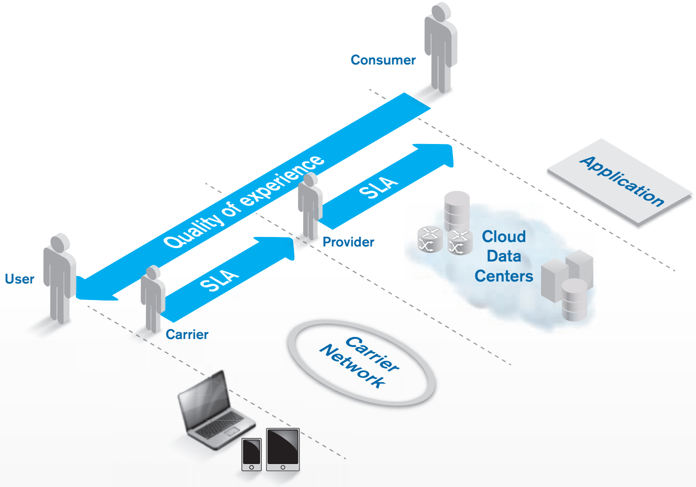

Figure 1. Cloud players: cloud users, cloud consumers, cloud providers and cloud carriers

Cloud services are sold to cloud consumers who have a business need. To meet that need, the cloud consumer deploys an application to be run in the cloud for a user community. The cloud itself is driven by cloud data centers that provide an environment for running the application. The data centers provide servers, storage and networking. User access to the cloud data centers is provided by cloud carriers. The cloud provider manages the cloud data centers and their servers, storage and networking. The carrier manages the interconnection between the user and the cloud data centers. The application may be owned and managed by the cloud consumer or the cloud provider.

The quality of the user's experience depends on both the carrier and the cloud provider. Together, they determine how well the application serves its users. The carrier manages bandwidth, latency, reachability, loss and other network key performance indicators (KPIs) that affect quality. The cloud provider manages processor utilization, storage, switch utilization and other resource KPIs that affect quality. But the KPIs that describe the quality of the service (QoS) provided to the consumer and the user are web download times, service availability, data delivery times, and other KPIs that are tied more directly to the service sold to cloud consumers. These service-oriented KPIs cannot be measured by any one actor; together, they ether is determined by the business relationship between them. If they are independent, cloud carriers offer SLA guarantees to cloud providers. Cloud providers in turn offer SLAs to cloud consumers. If the carrier owns the cloud provider, then the single organization can offer the service-level agreements (SLA). These SLAs are one factor driving carriers and providers to test the cloud.

Actor	Definition
Cloud User	A person or organization that uses and benefits from the cloud.
Cloud Consumer	An organization or person that buys services from the cloud provider to use or run an application.
Cloud Provider	A person or organization that provides a cloud service.
Cloud Carrier	A communications service provider that provides connectivity and transport between users and the cloud or within the cloud.

Table 1. Cloud actors

In this complex business environment, carriers have strengths. First, they own the network that connects users to the application running in the cloud. Second, they provide the circuits that connect cloud data centers to the Internet and each other. Third, they can provide the security and privacy customers want through dedicated circuits or virtual private networks that isolate one customer's traffic from another's. Finally, they know how to offer a high-quality service backed by the guarantees of an SLA.

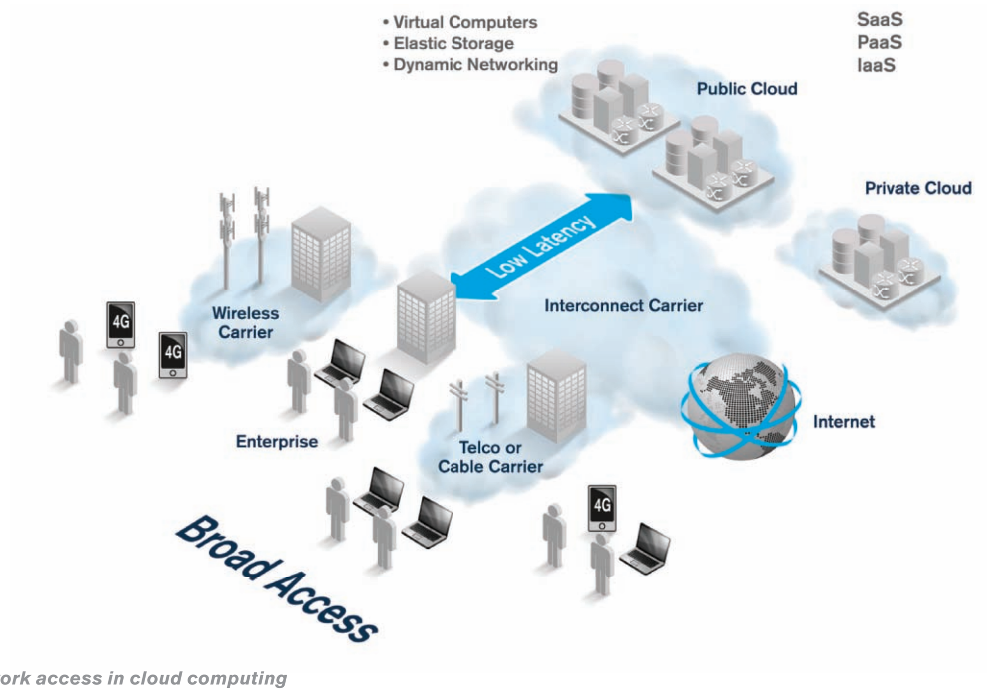

Figure 2. Broad network access in cloud computing

WHAT THE CLOUD OFFERS

Why do cloud consumers build applications in the cloud? The National Institute of Standards (NIST) has identified five essential characteristics of a cloud. These five characteristics provide a good summary of the benefits that cloud consumers want.

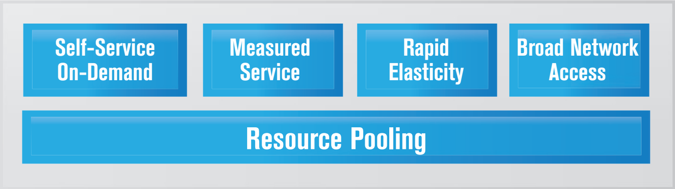

Figure 3. The essential characteristics of cloud computing (NIST)

Pooling resources to be shared among many cloud consumers yields cost savings for all. Cloud consumers pay only for the resources they use. Instead of buying a rack of servers that end up sitting idle most of the time, the cloud consumer can lease only what they need. The resources available in the cloud data centers generally include racks of central processing units (CPUs) providing virtual servers on demand, racks of disks providing elastic storage on demand, and a data center network interconnection providing addresses and network capacity on demand. Cloud consumers control the quantity of resources they lease using a self-service administrative interface, and can expand or contract the resources as needed. Usage is measured to bill the cloud consumer for the resources used. This increase in flexibility enables cloud consumers to take control of their expenses.

Applications can support a broad user community over the Internet. To support Internet scale applications, cloud data centers need high-bandwidth, low-latency connections to Internet exchanges. A high-bandwidth connection to the data center helps applications scale to large numbers of users. Still, a fast-connection distant data center won't perform well.

For most applications, latency has a bigger impact on user-perceived performance. Latency depends on distance due to the speed of light. It is therefore important to place applications in data centers located in close geographic proximity to users to ensure that data is efficiently routed between the data center and the user. Most enterprises would have a difficult time building data centers in the US, Europe, and Far East that would be capable of supporting a worldwide user base. But by renting servers in the cloud, enterprises can afford to deploy applications in close proximity to their users, wherever they happen to be located. Cloud carriers are already facing the latency problem and have placed the data centers (central offices) they operate close to users. By offering cloud services out of those centers, they can provide the lowest latency available.

There are some applications for which throughput is more important than latency. Any application that transfers large amounts of data between the application and the user will be sensitive to the available bandwidth. Streaming video and data backup services are examples of bandwidth-sensitive applications. A cloud data center with a high-bandwidth connection to the core of the Internet will be at an advantage in serving these types of applications.

Regardless of how good the cloud application is, it is of little use if unavailable to users. The cloud consumer or application owner can use the cloud to deploy the application at multiple sites that can be isolated to prevent any failure at one site from affecting the other sites. But, this introduces the technical problem of synchronizing the application's state and the user's state across the cloud. If a user stores a file in the cloud at one data center, it needs to be made readily available in the data centers serving that cloud-based storage application. Solutions to this technical problem generally require high-quality, high-performance connections between the cloud data centers.

Cloud carriers can meet the unique needs of applications by mapping the applications' traffic into specific classes of service (CoS), each with its own bandwidth and traffic-handling policies. Quite a bit of effort has been made in the standards to define some typical CoS and the appropriate performance objectives that they should meet. In particular, the Metro Ethernet Forum (MEF) has developed MEF 23.1, which includes markings and objectives for three CoS that apply to cloud carriers. Most carriers are familiar with providing CoS-aware services and are comfortable with providing CoS-aware SLAs.

While there are plenty of advantages to deploying applications in the cloud, most cloud consumers fear the loss of control, and the potential security and privacy breaches that moving to the cloud can bring. The fear of losing control can best be overcome by transparency, e.g., by providing the customer with the information they feel they need. Information about the number of users, how often they use the service, the resources consumed and the charges accrued are all examples of data that should be provided in any measured service. When combined with effective controls concerning who is allowed to use the service, the administration capabilities and the customer's ability to configure the services on demand will go a long way to addressing the perceived loss of control.

Privacy and security concerns can be partly met using virtualization techniques such as virtual private networks (VPNs) and virtual machines (VMs) to isolate one cloud consumer from another. Carriers, of course, are very familiar with providing VPNs and dedicated circuits as needed to ensure customer privacy. VMs are similar, and can be used to guard the integrity and privacy of servers and storage in the cloud. Security on the whole must be addressed as a holistic problem, and all components in the solution must be part of that solution.

TYPES OF CLOUDS

The NIST Model of Cloud Computing describes three types of clouds: infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS). These three types of service offerings are a good way to layer the services, for which there are many existing examples.

In IaaS, the cloud provider sells simple servers with little or no software. The servers may come with a virtual machine and an operating system, with attached storage and network connections, IP addresses and DNS names. The cloud consumer provides all the software needed by the application. An example of an IaaS provider is Amazon's EC2 (Elastic Compute Cloud) service. Other IaaS vendors include Savvis (owned by CenturyLink), Rackspace, IBM, GoGrid and Terremark (owned by Verizon).

In PaaS, the cloud provider sells a software platform to make application development easy. Platforms can include services like a web server, a database, shopping carts and content management tools (like Wordpress or Drupal). The value provided by PaaS is that cloud consumers can develop their applications more quickly on these middleware services than on bare infrastructure. An example of a PaaS provider is Google's App Engine. App Engine provides a web server, NoSQL and SQL databases, storage services and much more. Other PaaS vendors include Microsoft's Azure, the OpenStack platform and Force.com.

Key PaaS providers are cloud storage services like Amazon's S3 (simple storage service). Cloud storage provides a very simple web-based application programmable interface (API) with replication, which makes it easy to develop sophisticated device backup and file sharing services. Apps that "sync" mobile devices to cloud-based storage or that share photos via the cloud are enabled by cloud storage services.

In SaaS, the cloud provider sells the complete application. The cloud consumer outsources the entire application to the cloud provider, and grants users access to their specific instance of the application. Almost any application can be delivered as SaaS. An example is Salesforce.com, which provides a popular customer relationship management (CRM) application. Other examples include hosted e-mail applications from Yahoo, Google and Microsoft, ERP applications from SAP, and hosted private branch exchange (PBX) from XO Communications. Just about any application can be provided as an SaaS.

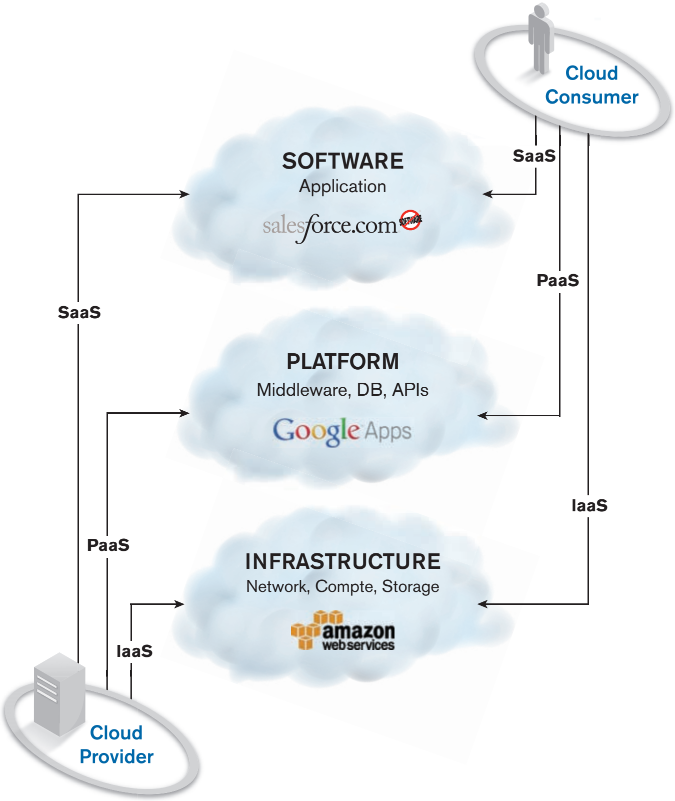

Figure 4. The three main types of cloud services: IaaS, PaaS and SaaS [NIST]

DEPLOYMENT AND OWNERSHIP MODELS

The NIST model of cloud computing describes three models for cloud deployment and cloud ownership, and in part, some basic business models designating how cloud carriers, cloud providers and cloud consumers relate to one another. The three deployment models are public clouds, private clouds and hybrid clouds.

In public clouds, the cloud provider is a separate entity selling services to many cloud consumers. The cloud provider is connected to the public Internet, and hosts applications accessible to anyone on the Internet. Public cloud providers offer access to the broadest possible community of users.

In private clouds, an enterprise buys its own equipment, and builds and operates its own data centers using the same server, storage and networking technologies as the public cloud. The users are mainly employees of the enterprise, connecting to the private cloud via a VPN for private, secure access.

In between these extremes lie hybrid clouds, in which some components of this model are public, and others are private. For example, a cloud-based banking application may have a public web-based interface to Internet users that is hosted on a public PaaS cloud, with a private connection to back-end servers running on a private infrastructure cloud that is located entirely within the bank's private network.

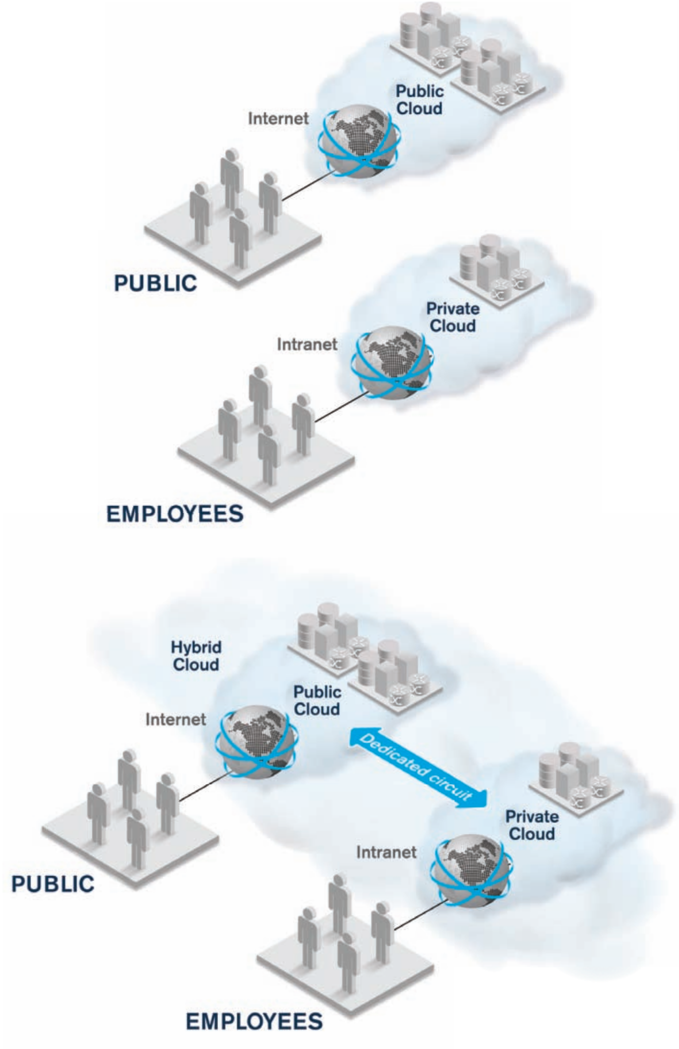

Figure 5. Cloud deployment models: public, private and hybrid

MEASURING THE CLOUD

Ultimately, the cloud consumer owns the application and therefore cares about users' quality of experience. Together, the cloud provider and carrier determine how well those expectations are met. If the network is slow, data centers are offline or the servers are overloaded, access to the application will be hampered and users will be affected. To guard against poor quality, providers must monitor the cloud's performance through measurement of availability, responsiveness and correctness. When providers have the appropriate measurements on hand, they can use this information to detect problems, predict when they will need to expand capacity, and perhaps most importantly, prove to cloud consumers that their service quality expectations and guarantees have been met.

The specific measurements needed depend on the service being offered to the cloud consumers, and the business relationship between the cloud carrier and the cloud provider. The sections below describe the measurements of quality applicable to carriers, IaaS, PaaS and SaaS providers.

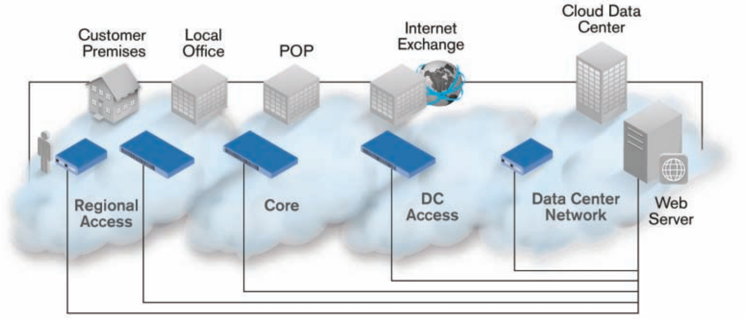

Figure 6. Testing cloud services from the carrier's network

Measuring Carrier Services

Cloud carriers have been selling data connectivity services for decades, and as such, the familiar measurements provided by testing the availability, loss, latency, delay variation and throughput of the connections all apply. The location from which the tests are conducted depends largely on how the cloud is deployed and who its users are.

For a public cloud providing services to the general public, a wireless provider might enter into a contract with the cloud provider to offer its subscribers fast access to the nearest cloud data center. The wireless provider could conduct tests from each 3G or 4G mobile switching center (MSC) to the nearest data center. The tests serve to verify the availability, network latency, packet loss and throughput between the testing MSC and the data center. For cable and wireline carriers, the same kinds of tests would be conducted from headends and central offices.

The idea is to test access to a public cloud from as near as possible to the user as is practical. Indeed, testing can be conducted from fixed locations like the home or small business using dedicated test devices. It is even possible to run tests from mobile test units or the user's phone or tablet. These examples have the advantage of being able to test the last mile, but also have the disadvantage of having to test a potentially huge number of devices. For cost reasons, most carriers would opt to test from a subset of the users, but this introduces its own set of problems involving signing up, selecting and managing a representative population of users. In any case, carriers who own the network that is used by public users to access the cloud are in the best position to measure service quality from the perspective of many users.

Carriers also sell dedicated L2 or L3 circuits to cloud providers. Dedicated circuits between distributed data centers owned by a public cloud provider are used to synchronize the application state between the data centers. Circuits from public or private data centers to Internet peering locations are used to connect the data centers to the Internet. Dedicated circuits from a private cloud data center to a public cloud data center can be used to build a hybrid cloud. VPNs that run over the Internet can be used to provide a virtual private cloud that actually runs on a public or hybrid cloud.

These circuits carry multiple CoS, each with its own QoS markings, provisioned CIR and EIR, and its own service-quality objectives. Testing and monitoring the circuit implies testing and monitoring each CoS.

Regardless of the types of circuits provided by the carrier, the same basic measurements are needed. The measurements may use different testing protocols or operate at different layers, but the metrics are similar. Carriers conduct tests to measure availability, loss, latency and delay variation (or jitter) on a 24/7 basis. They monitor utilization of the circuit and each CoS around the clock, and may also monitor which applications are using the circuit. They also perform tests to confirm that the provisioned bandwidth is actually available whenever the circuit bandwidth or configuration changes. Please note that on-demand elastic bandwidth provisioning on these circuits (a key feature of Carrier Ethernet networks running over fiber-optic cables) implies that these “turn-up” bandwidth tests need to be run on a frequent and automatic basis. The days of T1 circuits being set up and left untouched for years are long past. Indeed, changes are now so frequent that carriers may wish to periodically retest the circuit’s throughput.

While the carrier’s network performance is vital, it does not represent the whole picture. The user’s perception of the service’s quality is dependent upon both the carrier’s network and the cloud provider’s data centers. To see the cloud service, we need to test further up the stack, as described in the following sections.

Measuring Infrastructure Services (IaaS)

A cloud IaaS provider sells basic infrastructure components of a VM on which software can be run, in addition to storage for data and a local network enabling communications within the cloud data center. Each of these infrastructure resources must be managed and monitored.

There are existing tools in the marketplace that can be used to manage VM. These tools monitor the number of VMs being run, and the CPU, memory, I/O and network resources they are using. They also measure which operating systems are being run on the VMs, in addition to the middleware and applications running on those virtual machines. Similar tools are also available to manage storage, and to monitor how much storage is used and which users are consuming the disk space. There are also tools available to manage and monitor the data center network, set up VLANs and routing domains, and measure the usage of ports, VLANs, routers and switches.

These infrastructure management and monitoring systems are absolutely critical to managing the infrastructure. However, they only address part of the problem. To measure service quality, the performance of the service needs to be tested from the user’s perspective. This means testing from locations within the carrier’s network as close as practically possible to the users. The infrastructure services that should be tested include the DNS, data center network, server, application and web. Each of these services should be tested for availability and performance.

DNS provides a basic network name to address translation. As new VMs are brought online, they must be assigned network addresses. These addresses must be set to the DNS name used in the application. Users access the VM by name. Routing users to the nearest cloud data center is accomplished through DNS. A problem in DNS (either a slow response or a wrong translation) can have significant impacts on the application. A very simple test to ensure that the DNS name can be looked up quickly and that it resolves to the right address will detect these problems. Large cloud data centers are deploying new routing technologies like TRILL and IEEE 802.1aq (shortest path bridging), in addition to new software defined networks (SDN) like OpenFlow in order to deploy a more scalable, dynamic local switching and routing network within the local network. A carrier’s responsibility often ends at the switch, where a circuit plugs into the local data center network. Measuring the local switched network in the data center or the end-to-end route from the user to the server can be done with simple reachability tests, such as ping. More sophisticated network tests, like two-way active measurement protocol (TWAMP), can measure packet loss, latency, delay variation and other important measures of IP network performance and quality.

Cloud data centers are protected by firewalls and VPN servers, and use load balancers in the network. Firewalls protect the servers from attack. VPNs are used when the consumer desires private communications between the user and the cloud. Load balancers in the network front end serve a cluster of servers and balance the workload among them. Each of these network components affect the service offered in an IaaS cloud. Measuring their impact on performance can be done with differential testing, running one test of the service through the firewall, VPN or load balancer, and another directly to bypass the device.

To ensure that the virtual machines in the data centers are actually up and running an operating system, you can use simple tests like UDP Echo. Further tests can be run to verify that an application is running and listening on specific TCP or UDP ports within the OS, and that those ports are reachable by users across the network. The measurements provided include availability, and simple response and connection times.

Most applications that run in the cloud are web applications of one form or another. While an IaaS provider is not responsible for the application, it is often responsible for keeping web servers up and performing properly. To test those web servers, simple web requests must be made to the web servers running over HTTP or HTTPS. A simple test used to download a URL is capable of measuring the availability, response time, download time and throughput of the web service.

Oftentimes, tests can be combined. For example, a web service running on a cluster of servers behind a load balancer can be tested in two different ways. The tests performed through the load balancer show the service as a whole, as a user sees it. The test performed directly on each machine in the cluster tests its performance and availability. Combining the two tests allows you to isolate which machine(s) are causing a problem, helping speed up isolation.

Measuring Platform Services (PaaS)

PaaS cloud providers build on the same kind of infrastructures as IaaS providers, adding middleware that makes it easier to develop applications. Almost any middleware could be offered as a platform, but the most common are those services that are designed to make building websites or mobile apps easier. All the testing and monitoring capabilities described for IaaS are also needed for PaaS providers. However, the PaaS provider needs to go one level up in the stack to fully test the service. Three examples of PaaS services are website development services, mobile application development services and cloud storage services.

Website development services provide a basic web server such as Apache, content management systems such as WordPress, and APIs such as Google's App Engine. These tools enable the cloud consumer to develop a better website, faster. The website test is performed by downloading the complete web page and measuring its availability, response time and download time, and the throughput of the page and all of its content.

Mobile application development services simplify the process of building applications that run in mobile devices (e.g., phones or tablets), or in browser environments like Google's Chrome or Microsoft's Metro. While these services vary, many are REST APIs built on a web protocol for communications, and a web server running in the cloud. The application issues HTTP GET and PUT operations to the right URL, transmitting data encoded in a language like XML or JSON. Tests that measure the availability of the service, detect failures, and measure the service's response time and throughput simply issue the same API calls that an application would.

Cloud storage services allow applications to use persistent storage in the cloud. Cloud storage can be used for applications such as data backup on a cell phone or synchronization of files between tablets and PCs, or for secure file sharing applications. These services run over HTTP/S, and generally provide the ability to get and put files to the cloud. Because it can be time-consuming to replicate the files to other sites, a key performance metric consists of measuring how long it takes to distribute updated files to all the servers. To test these services, get and put operations are issued to the cloud storage service. Testing replication involves writing a new file to one site and then getting the file from a different site, while measuring how long it takes the new file to propagate.

Measuring Application Services (SaaS)

An SaaS provider builds on the infrastructure and platform services described above, but delivers a complete application as a service. Measuring the service means measuring the application from the users' perspective. Many application management tools measure the application server in traditional client/server architecture. This works fine when the users are close to the server or within the enterprise, but for the broad user community served by the public cloud, these traditional application management tools break down. Simply put, you have to manage the service, not the server. The specific measurements needed depend entirely on the application. The following examples illustrate this point.

For web-based software, such as a CRM application in the cloud like Salesforce.com, the user interface is entirely within the web, and therefore the web-based testing described above works just fine. In some cases, testing the application requires more sophisticated application scripting, where a series of web pages must be navigated to perform an action that should be measured. But in all of these cases, the availability and response time of the application are the primary measurements to be taken. The nature of the cloud means that most applications provide a web front end. These applications can all be tested from sites in close proximity to users via the previously described types of web-testing methods.

E-mail in the Cloud is a surprisingly large application offered as an SaaS. Gmail, Hotmail, and Yahoo! Mail are all examples of cloud-based e-mail provided to the public and enterprises. While most of these providers offer a web-based user interface, they also provide POP3, IMAP or SMTP interfaces to traditional e-mail clients, for example: Outlook on Windows and Sparrow for Mac. Providers can test the availability and response time of these interfaces, and also test delivery times by sending a message from one destination to another and measuring how long it takes to arrive.

PBX-in-the-Cloud applications host PBX functions for business voice communications in the cloud using SIP and VoIP technologies to lower the cost of ownership. These services can be tested using application-aware test tools that understand these protocols. The tests can measure voice quality KPIs such as MOS scores, and signaling KPIs like post-dial delay and post-pickup delay.

Over-the-top video services such as YouTube, Hulu and Netflix are all delivered in the cloud. Although these applications are web-based, they use specialized video delivery protocols to download the video to the user's device. Tests on these services are performed by downloading a video using the appropriate delivery protocol. The tests measure the availability of the service, the response time of the application, the download time and the throughput experienced by the user. They also detect any episodes of rebuffering in which the delivered video data falls behind the playout, and any other errors users would perceive as failures.

For SaaS testing, it's all about the application.

What's Common

Across these four service layers (Carrier, IaaS, PaaS and SaaS), there are a few common trends.

First, testing from a location close to the user and located within the carrier's network provides the best measurement of the user's experience, which is perfectly logical. The user's QoE depends on the full network path from the user through to the access network, through the backbone to the nearest cloud data center, and all the way to the server hosting the application. Carriers own the network, and therefore also own the test locations. This means that they are equipped to perform the best testing.

Due to the variety of services offered through the cloud, a variety of tests need to be performed. Any testing solution offered must be a multiplay solution—in other words, one capable of testing at all layers, from L2 to L7.

LIFECYCLE

The testing described above must be performed throughout the lifecycle of the service, as shown in the figure below.

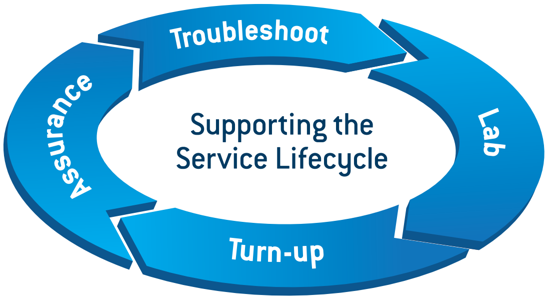

Figure 7. Service Testing Lifecycle

During development of the service, testing is conducted in a development laboratory environment in order to measure the capacity of the system under load. Cloud services, especially public cloud services, are meant to be scaled to very large numbers of users. As such, load testing the service before deployment is critical.

When a new service or a new customer is first deployed, carriers and cloud providers should test to ensure that the service will actually function for the customer. The test should verify that the capacity measures up to the actual capacity purchased by the customer, and that the promised SLA is achievable with the resources provided. Some of the challenges in performing effective turn-up testing result from the elastic capacity, and the on-demand customer provisioning of capacity that is a key feature of the cloud. Cloud consumers expect to be able to turn up new VMs, new storage, new servers with new addresses, and new DNS names backed up by new network bandwidth within minutes. To ensure that those resources are available, carriers and cloud providers need to carry out tests to confirm that they are present within that time frame.

Once a service is up and operating, 24/7 assurance testing of the service should be set in motion. These tests measure availability and performance to ensure that SLAs are met, and proactively detect developing problems.

When problems are found, the operations staff and development teams need to be able to troubleshoot the issue. On-demand testing of the service can help shorten the amount of time required to isolate, diagnose and fix the problem. These troubleshooting tests often dig deeper into the service than routine testing.

Throughout this cycle, the testing system must be well integrated into the operational support system (OSS) and business support system (BSS) that drive that drive the service lifecycle. It is just not possible to keep up with changes in a dynamic on-demand elastic cloud using manual means; therefore automation is necessary.

CONCLUSION

Cloud providers and cloud carriers have complex business relationships that are constantly evolving. Cloud carriers can use their own infrastructure to operate a cloud, partner with existing cloud providers, or broker and audit multiple cloud providers. Cloud providers can buy network interconnections from cloud carriers, sell their cloud services through carriers as resellers, or partner with a carrier to provide an integrated offering.

Whatever the business relationship, cooperation is key to delivering a quality service to cloud consumers and cloud users. While each stakeholder must independently manage its own resources (e.g., the circuits, switches, servers and storage that form the cloud infrastructure), it must also cooperate with measuring the actual service delivered to its customers. Customers may buy the services with the objective of saving money, but if the service fails to meet their quality expectations they will switch to another provider or take the application back in-house.

The best way to confirm whether the service is really working is to test it. Carriers and providers want to perform tests as close as possible to the customer, which means conducting testing from within the carrier's network by testing the cloud and cloud data centers where the application is running.

Tests performed at the appropriate points in the service lifecycle are capable of accomplishing a great deal. During development, the service under load should be tested in order to determine whether it scales. During customer provisioning and turn-up, tests should be conducted to verify whether the service is up and running, and whether it has the anticipated capacity and performance. During normal, daily 24/7 operations, the service must be periodically tested in order to measure SLAs, manage end-to-end performance and proactively detect problems before users do. During troubleshooting, testing must be performed to isolate and diagnose any problems, and once again when the service is fixed to verify its success.

The tests performed are dependent upon the specific services being offered. The test should assess the service in conjunction with the needs of the user. If the service is delivering e-mail, tests designed to ping a web server won't be of much help. The table on the next page lists some of the tests that could be used by carriers, IaaS, PaaS and SaaS providers, and the measurements that they take.

Layer	Service	Test and Measurements
Carrier	Ethernet L2 EVC	Ethernet OAM Test: Circuit Availability, Loss, Delay, Variation EtherSAM Test (Y.1564): Throughput, Loss, Delay, Jitter (at load)
	IP Connectivity	Ethernet OAM Test: Circuit Availability, Loss, Delay, Variation EtherSAM Test (Y.1564): Throughput, Loss, Delay, Jitter (at load)
Iaas	DNS	DNS Resolution Test: DNS Availability, Response Time, Accuracy
	Switched LAN	Ping Test: Reachability, Loss, Latency
	Operating System	UDP ECHO Test: Availability, Loss, Latency
	VPN	VPN Connection Test: VPN Availability TWAMP: Reachability, Loss, Delay, Variation
	Firewalls	UDP and TCP Port Test: Port Availability, TCP Connection Delay
	Application Servers	TCP Port Test: Port Availability, TCP Connection Delay
	Web Server	Web URL Download Test: Web Server Availability, Response Time, Download Time, Throughput
	Web Load Balancers	Web URL Download Test: Web Server Availability, Response Time, Download Time, Throughput
Pass	Website Development	Web Page Download Test: Website Availability, Response Time, Download Time, Throughput (for all page content)
	App Development	Web Request Test: Availability, Response Time, Download Time, Throughput
	Cloud Storage	Cloud Storage Test: Availability, Response Time, Upload and Download Time, Throughput Cloud Replication Test: Replication Time
Saas	Web Application	Scripted Web Test: Availability, Success Failure of script, Total Script Time, Each Step Time
	Cloud E-mail	Email Tests for SMTP, POP3 and IMAP: Email Availability, Response Time Email Delivery Test: Message Delivery Time
	Hosted PBX	VoIP Call Tests: Availability, MOS, Post Dial Delay, Loss, Latency, Jitter, Buffer Overflow (and Underflow)
	Over-the-Top Video	Video Download Test: Availability, Download Time, Rebuffering, Video Quality

Table 2. Service tests

Cloud consumers utilize the cloud because it saves money. They like that they can buy only what they need, and elastically expand and contract their purchase on demand. They can also deliver better quality of experience to their users by choosing a cloud provider that has data centers in close proximity to users. The cloud can provide higher availability due to the simple fact that cloud providers are experts in operating large data centers. But, there is no denying the fact that cloud consumers fear the loss of control inherent to using the cloud. To help consumers overcome this fear, the cloud carrier and cloud provider must provide information about the cloud, especially information on users' quality of experience. Measuring QoE is dependent upon the cloud carrier working with the cloud provider to test the cloud. And together, they must make what they find available to their users. When it comes to user experience, the cloud must be transparent, not murky.

EXFO Headquarters > Tel.: +1 418 683-0211 | Toll-free: +1 800 663-3936 (USA and Canada) | Fax: +1 418 683-2170 | info@EXFO.com | **www.EXFO.com**

EXFO serves over 2000 customers in more than 100 countries. To find your local office contact details, please go to **www.EXFO.com/contact**.

Printed in Canada 12/12

Assessing Next-Gen Networks