



Ensuring High Service Levels in Cloud Computing

Keys to Effective Service Management

Table of Contents

Executive Summary.....	2
Introduction: The Cloud is Dead, Long Live the Cloud	2
What is Cloud Computing?.....	2
Application Platform Clouds—The Most Established “Water Vapor”	2
Private Clouds—No Cloud at All?	2
Infrastructure Clouds—The Most Disruptive Cloud Type	3
Five Things to Consider Before Committing to an Infrastructure Cloud Deployment	3
1. The Cloud is a Third Party Service—SLAs and Clear Accountability will be Vital	3
2. Performance and Availability will be outside Your Control—Monitoring will be Key	3
3. The Cloud is Relatively New—Try Before You Buy	3
4. Security of the Network is not Guaranteed—Assume the Worst.....	4
5. Availability of Cloud Components is Irrelevant—Service Availability and Performance Are the Only Valid Metrics	4
Implications of the Cloud for Service Level Management	4
Support for Highly Dynamic Environments—Polling for Data Will no Longer Cut it.....	4
Focus on End User Experience Monitoring—The Only Real Measure of Service	5
Seamless Integration between Multiple Service Provision Centers—Consistent Data Model with Flexible Service Modeling.....	5
Simplicity of Deployment and Use—Zero Touch Monitoring Set Up	5
Maximum Coverage of Technologies and Application Service Delivery Capabilities	5
MTM Technologies’ Aware360 ^{v2} : Service Level Management for the Cloud	6
About MTM Technologies	6

Executive Summary

For many organizations, the move to cloud computing is imminent. This paper offers an overview of some key issues organizations should be aware of before they make the plunge. Featuring a range of practical insights, the paper looks at some of the most pressing challenges cloud computing presents in terms of service level management, and it offers some tips for overcoming these obstacles.

Introduction: The Cloud is Dead, Long Live the Cloud

Cloud computing is getting a lot of coverage in the media and a lot of space in the slide decks presented in analyst briefings. Seemingly, however, each pundit has his or her own way of both defining and evaluating the merits of cloud computing. On one end of the spectrum, some see cloud computing as quite simply the beginning of the end of computing as we have known it. On the other end, the whole notion is being written off simply as a load of “tech hype,” the latest in a wave of media-generated frenzy that in the end won’t signify anything of lasting import. The truth, as always, is probably somewhere in the middle.

The purpose of this white paper is neither to extol cloud computing’s virtues nor to write it off, but rather to focus on what is happening today, and what some of the practical implications for today’s enterprises and service providers will be. Today, the reality is that many organizations are either using or delivering cloud computing services or have near-term plans to do so. For these organizations, there are several considerations and implications, and in this white paper, we’ll outline what those are. First, however, we’ll start with some definitions. (Feel free to skip if you’re getting tired of being told what cloud computing is.)

What is Cloud Computing?

Cloud computing means many different things to many different people. However, in order to clarify what we’ll be discussing, and what we won’t, it is important to provide a basic definition of terms. Quite simply, we define cloud computing as follows:

Any computing service that is provided outside the customer premise and that is provided on a ‘pay-as-you-go’ basis.

This is a very loose definition, but since most people already have their own understanding of cloud computing, it should suffice. Today, cloud computing falls into the three categories outlined below.

Application Platform Clouds—The Most Established “Water Vapor”

The most commonly used forms of cloud computing at the moment are social networking applications, services that are largely focused on the consumer, rather than on the business. In the business world, service desk and CRM systems are being offered by third parties on a subscription basis. Collectively we will describe these as “application platform clouds.” These are the predominant examples of cloud computing and they have been for a few years now. While they do meet the definition of cloud computing, most of these services emerged under another name, namely “software as a service” (SaaS) offerings. The leading providers of these services are transforming their offerings to begin providing generalized application platforms that allow customers to develop and deploy their own cloud-based applications.

Private Clouds—No Cloud at All?

There is also another form of cloud computing referred to as the “private cloud”, which is really a description of a highly virtualized, on-premise data center that is behaving as if it were that of a public cloud provider. Given the buzz around cloud computing, this term is likely to get more popular and replace what most would have described in terms of “virtualization” up until the recent past. Referring to these next generation data centers

as “private clouds” does run the potential for confusion, however. A key tenet of the cloud is the fact that the services are consumed on a pay-as-you-go basis. Also, the private cloud has a maximum resource capacity, unlike public clouds, which are theoretically unlimited in terms of resources available to the individual customer. In any event, private clouds are expected to proliferate, and in some cases may serve as a stepping stone to an organization’s use of public clouds. This will also potentially lead to the need for private/public cloud integration—including what is becoming known as “cloud-burst”: scenarios in which the private cloud has a temporary need for extra resources and bursts out in to the public cloud to obtain those resources.

Infrastructure Clouds—The Most Disruptive Cloud Type

More recently, some very large, highly branded organizations have launched pay-as-you-go computing infrastructure components such as servers, data storage, and network entities. These services are delivered in billable units of consumption, in a manner very analogous to such traditional public utilities as electricity or water providers. These offerings can be described as “infrastructure clouds.” These public infrastructure clouds have the potential to usher in the most disruption in the IT industry, and are likely to have a progressively larger impact on the private data center and its evolution.

In the following sections, we’ll be focusing on infrastructure clouds, outlining key up front considerations as well as some of the most important implications for service level management.

Five Things to Consider Before Committing to an Infrastructure Cloud Deployment

Based on both industry reports and interactions with our customers, it looks like 2009 is the year in which many organizations will be making some kind of entree into the world of cloud computing. For those organizations looking to take this step, here are five key things to consider, before any move is made.

1. The Cloud is a Third Party Service—SLAs and Clear Accountability will be Vital

While cloud computing is a new, rapidly evolving paradigm, in the end it is fundamentally another third party service, and must be managed as such. The business representatives negotiating these agreements will need to effectively negotiate commercial, technical, and legal terms. Toward that end, it will be important to establish concrete, measurable service level agreements (SLAs). While these set an effective groundwork, the degree of reliance an enterprise places on a cloud provider can’t be overstated: if a vendor’s cloud goes down for two hours, some organizations could stand to lose millions—lost revenue that SLAs simply won’t cover. Yet without clear SLAs and an effective means for verifying compliance, the damage from poor service levels will only be exacerbated.

2. Performance and Availability will be outside Your Control—Monitoring will be Key

The assets and infrastructure that make up the cloud service are not your organizations. They are owned by the cloud providers, which will most likely have a focus on maximizing their revenue, not necessarily optimizing the performance and availability of your services. Establishing sound monitoring practices for the cloud services from the outset will bring significant benefits in the long term. This has been true of any outsourced service provision, from the traditional managed service provider model to the application platform cloud model. Outsourcing delivery of service does not necessarily imply that you can outsource the monitoring of that service. Besides, today very few cloud providers are offering any form of service level monitoring to their customers. In short, they are providing the service but not proving that they are providing that service.

3. The Cloud is Relatively New—Try Before You Buy

Even if services are being offered by older, more well-established vendors, the reality is that cloud service

offerings are still relatively new, with very little track record. Often, service providers have a history of providing services that bear very little resemblance to the practice of service management. It is true that most leading cloud service providers have a history of providing some form of highly utilized consumer service based on massive computing facilities, but that is no reason to assume that they will provide the service that your business will need, either immediately or in the years to come. Consequently, piloting the services from multiple vendors is likely to be the best starting point.

4. Security of the Network is not Guaranteed—Assume the Worst

The use of cloud computing brings the security of interconnecting networks in to extreme doubt. Moving from a highly secure data center model to an Internet-based cloud model will require a complete review of network security. In essence, you have to assume that the network is totally insecure and that the only security mechanisms in place are those that come prepackaged with the infrastructure components in use. This has implications for all of the applications that you host on the cloud platform and may require a complete review or even a rewrite of those applications. In addition, it is important to ensure that the systems management solution doesn't become the weakest link in the security chain, and offer hackers an opportunity to compromise your applications and data. That's why it's important that the systems management solution also comply with the same security model as other components.

5. Availability of Cloud Components is Irrelevant

—Service Availability and Performance Are the Only Valid Metrics

This is the most significant systems management paradigm shift of them all. Most people's systems management experience has led them to associate systems availability with service availability, that is, if a network component was running normally, the services running across that network component were also running normally. There has been an implicit association of these two in most traditional performance and availability monitoring solutions. However, the "pay-as-you-go" nature of cloud computing breaks the link between component and service performance: typically, organizations pay for capacity or throughput, rather than specific components. Plus, the highly dynamic nature of the computing infrastructure that exists in the cloud makes traditional CMDB (or simple list) based systems management virtually impossible to implement. All the traditional server and network reporting that shows 99.999 up-time will become secondary and probably irrelevant for future service level management and reporting. What this means is that synthetic transaction monitoring—that is, generating, monitoring, and reporting on simulated service requests—will be of paramount importance.

Implications of the Cloud for Service Level Management

Cloud computing represents a fundamental paradigm shift for providers and customers alike, and service level management is no exception. Following are a few of the most critical implications cloud computing will have on service level management.

Support for Highly Dynamic Environments—Polling for Data Will No Longer Cut it

The advent of cloud computing will create a significant challenge to traditional service level management solutions, which have relied on knowledge of the infrastructure components and their relationships to perform their functions correctly. This applies particularly to centralized, server-based solutions that rely on a schedule of polling activities to monitor the underlying infrastructure. Given that there is no way to predict which cloud computing infrastructure components should be accessible at any point in time, the only reliable mechanism for monitoring cloud components for performance data will be through on-board agents that gracefully register with distributed collection points and, more importantly, gracefully de-register when the cloud computing

components close themselves down. This mechanism will require secure authenticated communication between agents and collection points, all configured in a highly resilient topology.

Focus on End User Experience Monitoring—The Only Real Measure of Service

As stated above, the linkage between component availability and service availability is finally broken in the cloud computing environment. That implies a need to monitor service performance and availability in a more direct manner that more accurately mirrors how real users access their applications. The cloud computing model also means that there is no convenient point in the network from which to monitor these applications. Solutions for the cloud will need to be simple to deploy across multiple monitoring points, including beyond the bounds of the corporate offices. There is no boundary any more between the corporate and private location when it comes to accessing services. The ubiquitous consumer to ubiquitous provider computing model poses serious challenges to monolithic, high-touch service level management solutions. Highly distributable, light weight monitoring components will be required to get a high level and comprehensive view of the end user experience.

Seamless Integration between Multiple Service Provision Centers —Consistent Data Model with Flexible Service Modeling

The use of cloud services will not be a singular activity. Organizations are likely to use multiple service provision models, including the traditional data center, distributed department computing, SaaS-based CRM systems, and multiple virtual data centers residing in the cloud. We can refer to each one of these as a service provision center (SPC). The ability to monitor all of these SPCs and bring all the disparate data into a normalized monitoring database will be a challenge for the majority of legacy service level management solutions as they weren't designed with that in mind. Most traditional solutions assume a single SPC with a secure and reliable network connecting all the components back to a monolithic systems management server. Also, the logical association between the data across multiple SPCs needs to be very flexible as you may find that a critical application service is being delivered from a number of SPCs simultaneously. There will also be a need to model a service and its underlying applications and infrastructure in a highly dynamic way. Key to this will be the normalization of performance data in a logically consistent form and the ability to build highly dynamic service models on top of this data.

Simplicity of Deployment and Use—Zero Touch Monitoring Set up

Traditionally, there was an accepted process cost associated with adding monitoring to IT infrastructure. In the cloud computing paradigm, that accepted cost will no longer be acceptable. Systems management solutions will need to match the dynamic, on-demand model of the cloud. These solutions will need to be “zero-touch”—that is, when an infrastructure component is instantiated, no human intervention should be required for that component to also become part of the monitored system. Toward that end, the monitoring solution will have to be embedded inside the cloud components and activate and register themselves at the time of component instantiation.

Maximum Coverage of Technologies and Application Service Delivery Capabilities

To be viable, a service management solution must adapt to the “unknown unknowns” of the cloud's evolution. Many solutions have struggled to keep up with the pace of technological change during the virtualization phase of the last few years. The pace of change is not going to slow and the surviving systems management solutions of the future will be those that have been built with adaptive capabilities in mind. Strong application programming interfaces that allow rapid development of new technology monitoring will be key to success in the cloud computing world.

MTM Technologies' Aware360^{v2}: Service Level Management for the Cloud

Given the interest and activity surrounding cloud computing today, it's no surprise that many vendors are jumping into the fray. When selecting vendors, it is important to understand the difference between those vendors that are simply slapping a cloud label on an existing offering, versus those that truly understand the dynamic, paradigm-changing nature of cloud computing. With its cloud-proven solutions, MTM Technologies is a vendor that has the products and expertise that can help deliver effective service level management in today's cloud computing environments. Only Aware360^{v2} offers the ease of use, zero-touch setup, robust infrastructure coverage, and sophisticated end user response monitoring capabilities that are well-suited to cloud computing environments. In addition, Aware360^{v2} supports many cloud service providers' APIs and resources.

About MTM Technologies

Virtualization, Unified Communications, and Managed Services are the focus of the MTM Technologies drive to address specific customer challenges. One of only twenty-six companies in the United States to earn Cisco Master Unified Communications Specialization, MTM Technologies is also the Citrix Multi-Location Partner of the Year – an honor earned for the past two consecutive years. We create unprecedented value and opportunity for customers, strengthen network security, free employees to work from anywhere, and significantly enhance business agility. For more information, visit www.mtm.com.

Corporate Headquarters
MTM Technologies
1200 High Ridge Road
Stamford, CT 06905

Boston
Chicago
Cincinnati
Columbus
Dallas
Dover
Houston

Los Angeles
Milwaukee
New York
Philadelphia
San Francisco
St. Louis
Wilmington