FLATIRONS | SOLUTIONS

Trusted People, Proven Processes

Cloud Computing

De-Mystifying the Cloud

Jeff Vandervort
Senior Systems Engineer

Contents

Introduction

Cloud Computing is clearly the latest industry buzz word. With the amount of hype, one would expect to find plenty of technical information to support the claims, but nothing could be farther from the truth. Research has found that most of the information available tends to focus on the business proposition and actually gives very little insight into the underlying technology. A 2009 Federal survey by the

Merlin Federal Cloud Initiative shows the majority believe Cloud is the key to saving money, but there remains a general lack of understanding of what Cloud is and more importantly isn't.

This has created a dichotomy of views about Cloud computing in many organizations. For executives, the term "Cloud Computing" tends to conjure images of blue skies and fluffy white clouds. For IT managers, Cloud can represent a tumultuous proposition. While offering potential improvements to adopting organizations, Cloud Computing as an Enterprise solution is poorly understood and requires informed, objective analysis of the current computing environment and utilization, the performance, regulatory and security requirements, and associated costs. Absent a thorough analysis, an adopting organization will likely find the Cloud to be a stormy (and costly) place.

This paper attempts to objectively demystify "Cloud Computing", explore some of the related claims and provide an objective, non-vendor viewpoint on Cloud risks and opportunities.

Background

Cloud symbology has historically been used on network diagrams to represent the internet without having to enumerate its complexity. Cloud Computing simply extends the abstraction of complexity concept up to the data center layer.

Anyone that has used the internet has been exposed to some form of what today has been packaged into a single term; Cloud Computing. Google, Yahoo and Microsoft web-mail are commercially available on demand services that are now part of a group of services called "Software-as-a-Service". On demand hosting services, also known as "Platform-as-a-Service", have also provided cheap and easy way for individuals and small businesses to quickly establish a web presence. Simply chose an environment, select a domain name, and the client is ready to start developing their web site.

Most acknowledge Amazon as the pioneer in what is commonly understood as Cloud Computing. Using 3Tera's Applogic framework, Amazon was able to aggregate its excess data center capacity into a single web based solution they labeled Elastic Compute Cloud (EC2). Unlike traditional outsourcing, Amazon is able to extend control of data center resources to the client, as well as implement a cost model based on hourly usage. This Cloud formation is what is now commonly called "Infrastructure-as-a-Service"

Most Amazon EC2 users were initially enthusiast and weekend developers, but the response was so great they found they had to expand their data centers to keep up with demand. Amazon's popularity has encouraged others to develop their own offerings. Over the last year these offerings have begun to show up on the market and many have been directly targeting public and private enterprise customers.

With the benefits of each "as-a-Service" offering being associated with the single term "Cloud Computing", it has become extremely difficult for IT decision makers to sort through the hype and find real value.

De-Mystifying Cloud

Standards organizations, including the Open Cloud Consortium and National Institute of Standards & Technology (NIST) are attempting to develop a standard definition for Cloud Computing, but to date no single definition has been widely adopted. The problem is a Cloud can be implemented in many different ways, so depending on the configuration being discussed the definition can be quite different. Most definitions however do include common terms such as ubiquitous access, pooled resources, on demand self-service, metered usage and abstraction.

For Cloud, ubiquitous is often used in the context of platform and location independence. Access to Cloud resources is generally web based, but where they can be accessed really depends on the type of Cloud. Pooled resources and on demand provisioning are characteristics of the underlying virtualized environment of a Cloud. Self-service and abstraction allow users to get a server up and running with a few mouse clicks and very little technical knowledge. What really sets cloud apart though is the metered usage. Cloud provides a means to measure what services are used and for how long. This allows providers a means to bill customers only for the resources used.

The 3 basic service types are IAAS (Infrastructure-as-a-Service), PAAS (Platform-as-a-Service), and SAAS (Software-as-a-Service). The 2 common resource configurations are public and private, and the 4 location types are internal, external, hybrid, and community. A Cloud can have any combination of the type, configuration and location, and can even be configured to have more than one from each category. These combinations are being described by some as "Cloud Formations"

Infrastructure-as-a-Service (IAAS) is basically a means of delivering Data Center hardware to customers. The customer is responsible for the server operating system and applications. Amazon EC2, RackSpace and Terremark's V-Cloud are examples of IAAS providers. Platform-as-a-Service (PAAS) refers to a pre-configured development environment and includes the hardware, operating system, and some form of application development tools. Microsoft Azure and Google App Engine are an example of a PAAS service. Software-as-a-service is a pre-configured environment that provides specific functionality, such as e-mail the customers. Google Mail and SalesForce are examples of SAAS.

Public Clouds are hosted by a 3rd party in a shared data center environment, where a Private Cloud is often characterized by an environment that is dedicated to a single group or organization. Internal Clouds use an organization's data center resources, where external clouds use data center resources provided by a 3rd party and hybrid Clouds use a combination of internal and external resources. A community Cloud simply means 2 or more organizations are pooling internal cloud resources.

Cloud may seem like a revolutionary new technology but, in reality, many underlying Cloud technologies are currently used in many Enterprise Data Centers today. Virtualization for example maximizes server utilization and provides a means to rapidly provision new services. Several virtualization vendors have web based monitoring and control interfaces that offer Cloud like functionality. Metering may seem unnecessary for internal clouds but that may not be quite true. Metering can offer finer granularity on how IT resources are consumed, utilization rates and better visibility into TCO of each service provided.

The most common question asked about Cloud Computing is "what can I put into the Cloud?" Since Cloud is a virtual environment, basically anything that can be virtualized can be moved to the Cloud. Virtualization technologies used in cloud are industry standard, so applications that are already running in a virtual environment most likely can be easily converted and migrated. Most challenges come when compensating for the differences in attached storage devices and making changes to the network configuration, so even though it's possible, it can be technically challenging.

Compatibility with non-virtualized systems really depends on how the application was written, how it was configured, and if there are dependencies with other applications and or servers. Cloud architectures also vary from vendor to vendor, so compatibility with one vendor cloud may not necessarily guarantee universal compatibility. A thorough evaluation of each application and its dependencies must be made before compatibility can be determined.

For Software-as-a-Service Cloud services such as email, compatibility is not as big of a concern. Vendors such as Google can migrate email accounts from Microsoft Exchange and Lotus Notes fairly easily. What they may not be able to do is replicate the functionality of any dependent applications tied to the enterprise email system.

Cost versus Traditional Alternatives

Most real world use cases and testimonials tend to focus just on ROI without fully explaining how the calculation was derived or even what cloud formation was used. Since most of the early adopters are testing the waters with services such as outwardly facing web servers, non-mission critical applications and short duration projects, the number of testimonials about experiences with business critical services and real world TCO / ROI numbers are very limited.

With prices hovering around $0.10 per hour per server instance for IAAS, Cloud Computing seems almost too good to be true. $0.10 typically only buys a minimally configured, single CPU server. Is ownership really that much more expensive? For comparison, we will use the Dell R410 server. The R410 has 16 processor threads (2 quad core Zeon processors with hyper threading). Each thread is equivalent to 1 Cloud server instance. Configured with enough RAM and disk space to configure the equivalent of 16 base Cloud instances, the R410 would cost as low as $7,000. Assuming a 4 year lifecycle and 8760 hours per year, this translates to 1.2 cents per hour / per instance to own. Running the same 16 instances in the Cloud would cost over $56,000 for the same 4 year period.

Even though this comparison doesn't include the facility, infrastructure, manpower or utility costs, the $0.10 hourly rate for Cloud isn't the bottom line price either. Cloud Computing models are a-la-carte, which can further mask the true cost of Cloud solutions. Besides the server, Kilobytes of bandwidth consumed (inbound and outbound), Gigabytes of data stored, IP addresses, Load balancers, and user access/transaction fees are all extra charges.

These might not sound like a big deal, but Cloud costs are much different than traditional data center costs. Traditional Network Bandwidth for example is determined by the maximum amount that can be consumed and that cost is divided by the total number of servers to calculate TCO. With Cloud each server is charged for the actual bytes consumed. While this does mean a customer is not paying for unused bandwidth, it doesn't guarantee a minimum throughput.

For an individual or small business with no existing infrastructure and relatively low utilization rates, Cloud can definitely be a cheaper than building from scratch. For an Enterprise running servers 24/7 the savings might not be as easy to achieve. Simply moving a few servers to the cloud probably won't significantly reduce the existing ownership costs. Most of the network infrastructure and associated utility costs will have to be retained to support the remaining servers and users. With Infrastructure-as-a-service, the burden of maintaining the operating system and application remains with the client, so for the most part, the manpower supporting data center services will probably need to remain intact too.

With so much written about cloud and an explosion in vendor offerings available, it is becoming increasingly difficult to wade through the mounds of conflicting information to make critical investment decisions. There are some real benefits with Cloud, but there are also many risks that must be understood and challenges that should be considered before committing to a cloud based solution.

Risks

Contrary to popular belief, the number of organizations moving to Cloud is still relatively small. The majority actually say they are still in the learning / evaluation stage and intend to proceed cautiously. Cloud is a relatively new concept and is still evolving. The lack of security, configuration and interoperability standards have many potential customers simply sitting on the side lines waiting for the best of breed vendors to emerge.

Security is by far the number one concern. Most vendors are hesitant to disclose their security practices, but will instead offer proof of 3rd party auditing. Although understandable, the lack of visibility into the security processes can cause anxiety for any security manager. Vendors will also often cite the financial incentive to maintain a high level of security, but that same financial incentive could also deter a vendor from fully disclosing incidents that a customer would not be otherwise aware of.

Privacy concerns stem from several highly visible Google incidents where personal information was inadvertently shared with other users. Although Google stated only a small percentage of users were affected, the potential for future incidents associated with storing data in the Cloud are still a major concern.

Some of the inherent benefits of Cloud can also be vulnerabilities. For high availability, Cloud providers often have more than one facility and a customer's data can potentially be replicated to many different locations around the globe. Some vendors sub lease from 3rd party vendors, and those relationships may not be disclosed to the customer. To maintain high utilization rates, resources are shared with other customers. This could mean multiple customers could be sharing the same physical processor and hard drive. The very nature of Cloud means a customer could potentially be sharing resources with a hacker.

The hardware itself can potentially introduce vulnerabilities. In 2008, the FBI discovered over 300 counterfeit Cisco Routers had been installed on public networks. These routers had code embedded in the chips that gave hackers backdoor access to routers and installed networks. The investigation found the Routers were purchased at discount prices outside Cisco's partner network. Depending on the relationship with the Cloud vendor, disclosure of supply channels used by the vendor may not be possible.

Performance can also be a concern. Major companies have invested billions of dollars building out giant data centers to support the anticipated flood of Cloud customers. Since most of these data centers are sitting relatively empty, early adopters will find performance levels more than acceptable. As these data centers begin to fill up, performance can become an issue. Depending on the SLA, customers may have no recourse if performance degrades below acceptable levels.

Other concerns not discussed here include regulatory compliance, investigative support, Cloud sprawl, the potential for future costing increases once committed, proper disposal of storage media, and retrieval of data once a relationship is terminated.

Service Level Agreements

SLA's simply ensure there is a common understanding about services, priorities, responsibilities, guarantees and warranties, and establish minimum and targeted performance levels. Standard customer SLA's may appear straightforward and accommodating, however executing them can prove difficult. Amazon defines an outage as the inability to launch a replacement image in one of their availability zones within 5 minutes. The customer is responsible for proving the outage occurred and reimbursement is generally limited to just the period of non-availability. Although many claim 99.999% availability, some SLA's state a customer will only be compensated for outages once up time drops below 99.995% over a rolling 12 month period.

Most will find performance points and measures different for Cloud than traditional SLA's for IT services. With Cloud, something as simple as defining what constitutes an outage can be extremely confusing. Some SLA's state an outage occurs once it is reported by the customer, others once it is verified to be the provider's issue. With IAAS for example, the client responsibility starts at the operating system. This can make it difficult to for a client to prove the outage wasn't due to issues with the client's operating system or installed applications.

Summary

This paper touches on many key aspects of Cloud, but only scratches the surface of the many challenges associated with Cloud adoption. There are many benefits to Cloud and they should be capitalized on, but only when and where it makes sense. There are also many risks and concerns that need to be understood and addressed before undertaking a Cloud project.

Cloud Computing may be one of the latest data center technologies, but it's definitely not the only technology arena that data center managers must consider today. Evolving server communication and storage technologies coupled with data center virtualization tools offer significantly higher server utilization rates. Personnel, infrastructure, sunk costs, and flexibility required for example must also be factored into any enterprise data center solution.

Utilizing a solid methodology that assesses the existing environment, accounts for business goals and user needs, and at a minimum, fully addresses the critical areas addressed in this paper is a must. Flatirons Solutions and its team of partners will assist your data center managers and IT investment owners make informed decisions about their infrastructure needs for today and for tomorrow. Using proven methodologies, our experts work side-by-side with IT investment decision-makers in large organizations to sift through the many options available and identify the implications of each to their organization and stakeholders.

About Flatirons Solutions

Flatirons Solutions, an Inc. 500 company, provides consulting, systems integration, and systems & software engineering services to Fortune 500 companies and government agencies. A leading content management solutions provider specializing in XML-based publishing, Flatirons has provided enterprise-wide solutions – using both DITA and DocBook – in industries such as

aerospace, transportation, publishing, manufacturing, financial services, insurance, and healthcare. Flatirons Solutions also actively participates in both DITA and DocBook OASIS technical committees. Established in 2001, Flatirons Solutions is a privately-held company headquartered in Boulder, Colorado, with offices in Washington D.C. and Ft Worth, TX. For more information visit Flatirons Solutions on the web at http://www.FlatironsSolutions.com.

4747 Table Mesa Drive, Suite 200
Boulder, Colorado 80305

(303) 544-0514
info@FlatironsSolutions.com

www.FlatironsSolutions.com