Avoiding a Storm by Evaluating the Clouds: A Guide to Cloud Computing

Author: Vincent S. Cockrell

EE616 Information Engineering Graduate Project

Table of Contents

Abstract

If you are alive and have touched a computer, smart phone or other WIFI enabled device within the past two years, you have been exposed to cloud computing. Cloud computing has infiltrated the world of technology in recent years in a massive way and its popularity continues to grow.

When first introduced to this technology, one would believe that the potential for crashing networks, theft of proprietary business data, and rampant viruses would deter companies from subjecting themselves to such vulnerabilities. If you think this way, you are wrong. In an effort to cut cost as they struggle to climb out of the muck and mire caused by the recession, companies are increasingly incorporating cloud computing into their information technology strategies. The words stated by Ray Ozzie in his parting memo to Microsoft, "Go Cloud or Go Home"; appear to be ringing loudly in most large organizations.

Respondents to Earnst & Young's 13th Global Information Security Survey 2010 believe that the "reliability and security level of many cloud services is still unknown" (Earnst & Young, 2010). Companies continually deal with the challenges of protecting the information stored in-house from outside attacks. Moving from architectures built for on-premises services and secured by firewalls and threat-detection systems to mobile environments minimizes the capability of companies to secure data effectively. Government regulation that is intended to ensure data security has been made obsolete by the recent advancements in cloud service offerings. In March 2009, a meeting was held by the Federal Trade Commission to discuss security and privacy issues related to cloud computing. During that meeting, it was agreed that data management services might experience failure similar to the financial meltdown if further regulation is not implemented. Along with regulation, it is imperative that companies become educated about cloud computing and implement

the necessary precautions to ensure that the integrity and security of their information remains intact (Rittinghouse & Ransome, 2010).

This document will neither support nor oppose the adoption of cloud computing. The intention of this research is to provide information about cloud computing that will assist those interested in this technology in making well-informed decisions. Customers must become aware of the predictors that if examined properly, will indicate whether the adoption of cloud computing services will be a successful venture for their organization.

What is Cloud Computing?

Cloud computing is a broad concept. Until recent years, the term cloud was used in referring to the Internet. An outline of a cloud was commonly drawn in Network diagrams to represent the transport of data to an endpoint over the Internet. The most common perception of the cloud depicts it as a warehouse of servers, accessible via the Internet, to which our information is stored. This warehouse would therefore be owned by a third party and exist somewhere in the ether.

In essence, cloud computing consists of the fore-mentioned concepts and much more. Cloud computing has evolved into the vision that was first publicly suggested by Professor John McCarthy in 1961. He described a computer time–sharing technology in which computing power and even specific applications might be sold through a utility type business model, similar to water and electricity (Dupre, 2008). “A more tempered view of cloud computing considers it the delivery of computational resources from a location other than the one from which you are computing” (Rittinghouse & Ransome, 2010. xxvii).

Cloud Computing Delivery Models

Cloud computing services are offered in four delivery models: public, private, hybrid, and community. Before selecting a delivery model, an organization must carefully consider their specific needs, the sensitivity of the functions and data that will be the migrated to the cloud, and the distinct implications of each model. The differences between the characteristics of each delivery model will have a significant impact on the nature, content, and terms of the cloud service contract (Gilbert, 2010. 3).

The public cloud infrastructure is made available to the public or a large industry group and is owned by an entity selling cloud services (Gilbert, 2010. 3). The public cloud provides elasticity in resources and services, which allows the customers to pay for only the services that they use. In the public cloud model, resources can be deployed or removed on demand, at the customer's request. Public clouds are best used for one-time processing of data that is non-critical to a company's operations or that may be subject to unpredictable traffic demands.

The private cloud infrastructure is operated solely for an entity. It may be managed by the entity or a third party and may exist on premise or off premise (Gilbert, 2010. 3). With a private cloud, an organization has the capabilities of rapid service provisioning, elasticity of resources, network latency, and high asset utilization. The organization maintains control of its data and has the ability to implement security and compliance policies at will.

The hybrid cloud infrastructure combines public and private clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability, such as when cloud bursting for load-balancing between clouds (Gilbert, 2010. 3).

The community cloud infrastructure is shared by several entities. It supports a specific community that has shared concerns, such as the same mission or policies, or similar security requirements or compliance considerations. It may be managed by the entities or a third party and may exist on premise or off premise (Gilbert, 2010. 3). In a community cloud model, the costs are spread over fewer users than a public cloud and is therefore, more expensive. However, this model offers a higher level of privacy, security, and policy compliance, which makes it very attractive to organizations that handle sensitive information.

Cloud Computing Service Models

Cloud Computing: Implementation, Management, and Security by John W. Rittinghouse and James F. Ransome describes the vital role that Amazon.com played in the development of cloud computing. In 2002, Amazon began providing access to its systems for third-party users on a utility computing basis via Amazon Web Service, which in turn, began a revolution of sorts. Amazon Web Services began implementing its model by renting computing cycles as a service outside a user's domain, without regard to the location of that domain (Rittinghouse & Ransome, 2010. xxviii).

Amazon provided their users the ability to access technology-enabled services in the cloud, without any need for knowledge of, expertise with, or control over how the technology infrastructure that supports those services worked. This approach transformed cloud computing, allowing data to be permanently stored in remote servers accessible via the Internet and cached temporarily on client devices that may include desktops, tablet computers, notebooks, hand-held devices, mobile phones, etc. This approach is often referred to as Software as a Service (SaaS) (Rittinghouse & Ransome, 2010. xxix).

SaaS is a type of cloud computing that delivers applications over the Internet to customers using a multiuser architecture. For the customer there is no need to invest in servers or software licensing. Since service providers have only one product to maintain, costs are relatively low compared to the costs incurred with a conventional hosting model (Rittinghouse & Ransome, 2010. xxix).

Platform-as-a-service (PaaS) delivers a platform from which to work rather than an application to work with. These service providers offer application programming interfaces (APIs) that enable developers to exploit functionality over the Internet, rather than delivering full-blown applications. In a PaaS model, developers build applications designed to run on the provider's infrastructure that

are then delivered to end users via an Internet browser. The main drawback to this approach is that the services are limited by the vendor's design and capabilities (Rittinghouse & Ransome, 2010. xxx).

Communication as a Service (CaaS) providers are responsible for the management of hardware and software required for delivering Voice over IP (VOIP) services, Instant Messaging (IM), and video conferencing capabilities to their customers. A CaaS model allows customers to deploy selectively, communications features and services throughout their company on a pay as you go basis for services used. All VOIP transport components are located in geographically diverse, secure data centers for high availability and survivability. Network and capacity and feature sets can be changed dynamically, so functionality keeps pace with customer demand and provider-owned resources are not wasted. CaaS providers perform periodic upgrades or replacement of hardware and software to keep the platform technologically current, thus eliminating the expense for ongoing maintenance and operations overhead. As an added benefit, every component is managed 24/7 by a CaaS Vendor (Rittinghouse & Ransome, 2010. 30 - 31).

Infrastructure as a Service (IaaS) provides a standardized infrastructure tailored to the customer's applications. In an IaaS offering, the customers will continue to own and manage the software applications. The hosting and infrastructure management functions will be responsibility of the IaaS provider. Customers have the potential to realize a huge savings by renting data center space, servers, software, and network equipment as needed, rather than purchasing them outright (Rittinghouse & Ransome, 2010. 34 - 35).

Monitoring as a Service (MaaS) provides security monitoring to protect an enterprise or government client from cyber threats. The benefit to MaaS is the 24/7 security monitoring services, which

allows for fast responses to incidents at all hours. This service offering should definitely be considered by organizations that transport business critical or other sensitive data over the Internet and companies that are required to run batch jobs during non-business hours. (Rittinghouse & Ransome, 2010. 44)

Essential Characteristics of Cloud Computing

Cloud services exhibit five essential characteristics that demonstrate their relation to, and differences from, traditional computing approaches: on-demand self service, broad network access, resource pooling, rapid elasticity, and measured service. The Cloud Security Alliance defined these characteristics in the "Security Guidance for Critical Areas of Focus in Cloud Computing" as follows:

- On-demand self-service. A consumer can unilaterally provision computing capabilities such as server time and network storage as needed, without requiring human interaction with a service provider.
- Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs) as well as other traditional or cloud-based software services.
- Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a degree of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources, but may specify location at a higher level of

abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines. Even private clouds tend to pool resources between different parts of the same organization.

- Rapid elasticity. Capabilities can be rapidly and elastically provisioned — in some cases automatically — to quickly scale out, and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- Measured service. Cloud systems automatically control and optimize resource usage by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, or active user accounts). Resource usage can be monitored, controlled, and reported — providing transparency for both the provider and consumer of the service (Cloud Security Alliance, 2009. 14 - 15).

Figure 1 - NIST Visual Model of Cloud Computing Definition

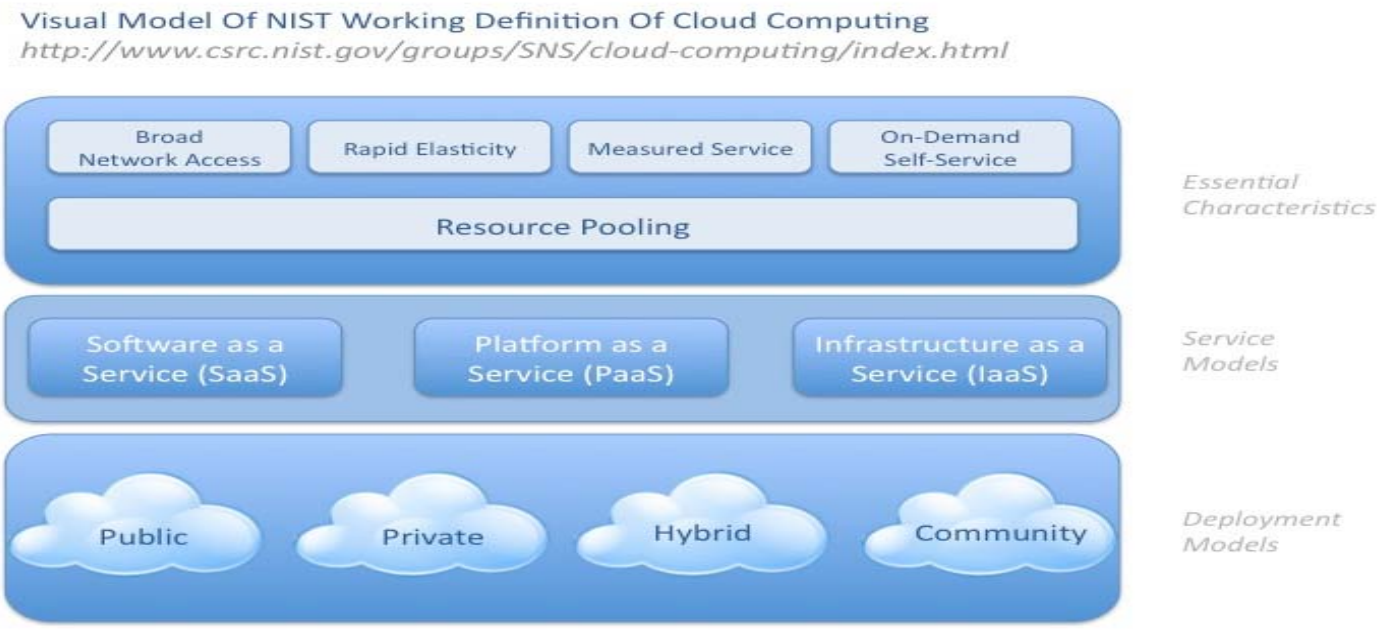

The Seven Standards of Cloud Computing

Changes in the information technology industry occur rapidly and companies that were once at the pinnacle of a particular technological movement tend to fade into the ether as hungry competitors offer bigger, better, faster, or cooler products. At this writing, Salesforce.com has proven themselves to be the leaders in cloud computing, and they are not only setting the pace, but also defining the very meaning of the cloud. SalesForce.com is setting the standard for cloud computing service delivery.

The Seven Standards of Cloud Computing Service Delivery, delivered in a white paper by SalesForce.com, describes the building blocks of the best practices that every successful cloud-computing platform should follow. The seven standards defined in this document are as follows:

1. World-class security – Provision world-class security at every level.
2. Trust and transparency – Provide transparent, real-time, accurate service performance and availability information.
3. True multitenancy – Deliver maximum scalability and performance to customers with a true multitenant architecture.
4. Proven scale – Support millions of users with proven scalability.
5. High performance – Deliver consistent, high-speed performance globally.
6. Complete disaster recovery – Protect customer data by running the service on multiple, geographically dispersed data centers with extensive backup, data archive, and failover capabilities.
7. High availability – Equip world-class facilities with proven high-availability infrastructure and application software.

World-Class Security

Security is more than just user privileges and password policies. Security is a multidimensional business imperative, especially for platforms that are responsible for customer data. Cloud-computing platforms must have detailed, robust policies and procedures in place to guarantee the highest possible levels of:

- Physical security
- Network security
- Application security
- Internal systems security
- Secure data-backup strategy
- Secure internal policies and procedures
- Third-party certification

Trust and Transparency

Cloud service providers should provide transparent, real-time, accurate service performance and availability information to their customers. Cloud-computing platforms should provide customers with detailed information about service delivery and performance in real-time, including:

- Accurate, timely, and detailed information about service performance data and planned maintenance activities
- Daily data on service availability and transaction performance
- Proactive communications regarding maintenance activities

True Multitenancy

A multitenant architecture allows for high scalability and faster innovation at a lower cost. Single-tenant systems, on the other hand, are not designed for large-scale cloud-computing success. The internal inefficiencies of maintaining a separate physical infrastructure and separate code lines for each customer make it impossible to deliver quality service or innovate quickly. Multitenancy provides customers with the following benefits:

- Efficient service delivery, with a low maintenance and upgrade burden
- Consistent performance and reliability based on an efficient, large-scale architecture
- Rapid product release cycles

Figure 2 - Multi-Tenancy (Cloud Security Alliance, 2009)

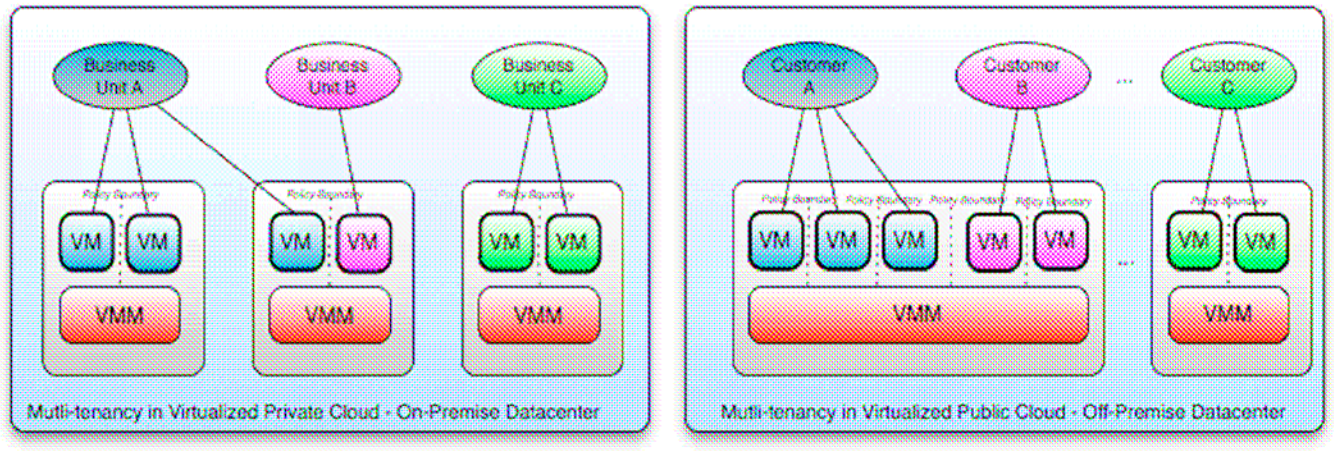

Private Cloud of Company XYZ with 3 business units, each with different security, SLA, governance and chargeback policies on shared infrastructure

Public Cloud Provider with 3 business customers, each with different security, SLA, governance and billing policies on shared infrastructure

Proven Scale

With any cloud-computing service, customers benefit from the scale of the platform. A larger scale means a larger customer community, which can deliver more and higher quality feedback to drive

future platform innovation. A larger customer community also provides rich opportunities for collaboration between customers, creating communities that can share interests and foster best practices. Cloud-computing platforms must have:

- Proof of the ability to scale to hundreds of thousands of subscribers
- Resources to guarantee the highest standards of service quality, performance, and security to every customer
- The ability to grow systems and infrastructure to meet changing demands
- Support that responds quickly and accurately to every customer
- Proven performance and reliability as customer numbers grow

High Performance

Premier cloud service providers can deliver consistent, high-speed performance globally. Cloud-computing platforms must deliver consistent, high-speed systems performance worldwide and provide detailed historical statistics to back up performance claims, including:

- Average page response times
- Average number of transactions per day

Complete Disaster Recovery

Cloud service providers should put forth effort to protect customer data by running the service on multiple, geographically dispersed data centers with extensive backup, data archive, and failover capabilities. Platforms providing cloud-computing services must be flexible enough to account for every potential disaster. A complete disaster recovery plan includes the following:

- Data backup procedures that create multiple backup copies of customers' data, in near real-time, at the disk level
- A multilevel backup strategy that includes disk-to-disk-to-tape data backup in which tape backups serve as a secondary level of backup, not as the primary disaster-recovery data source. This disk-oriented model ensures maximum recovery speed with a minimum potential for data loss in the event of a disaster.

High Availability

Equip world-class facilities with proven high availability infrastructure and application software.

Any platform offering cloud-computing applications needs to deliver very high availability.

Requirements for proving high availability include:

- Facilities with reliable power, cooling, and network infrastructure
- High-availability infrastructure: networking, server infrastructure, and software
- N+1 redundancy
- Detailed historical availability data on the entire service, not just on individual servers

(Salesforce.com, 2009)

Cloud Computing Security Risks

Although there are many benefits that can be attributed to cloud computing, companies are mainly seeking to lower the cost of their Information Technology assets. Companies are in business to make money and there certainly is truth to the adage quoted by Benjamin Franklin, "A Penny Saved is a Penny Earned," but when considering the adoption of cloud computing, saving money cannot be the only factor that is considered. The heart of a business lies within the data on which it operates and

manages. It is absolutely vital that the security and integrity of that data is protected by all means. Loss or misuse of critical applications and sensitive data could result in legal liability and significant damage to a company's image.

"Sensitive data processed outside the enterprise brings with it an inherent level of risk, because outsourced services bypass the physical, logical and personnel controls IT shops exert over in-house programs." Companies must gather as much information as possible about the service providers who will be entrusted with their data. Gartner suggests that customers "ask providers to supply specific information on the hiring and oversight of privileged administrators, and the controls over their access." (Brodkin, 2008, 1)

"Customers are ultimately responsible for the security and integrity of their own data, even when it is held by a service provider. Traditional service providers are subjected to external audits and security certifications. Cloud computing providers who refuse to undergo this scrutiny should raise a red flag to potential customers." Service providers must be transparent in their business practices to gain and maintain the trust of their customers. The absence of transparency could be a sign of unscrupulous business practices. (Brodkin, 2008, 1)

When companies use cloud services, it is unlikely that they will know exactly where their data is geographically located. Cloud service providers may even transport data between multiple countries in order to take advantage of low cost power costs during off peak hours. Gartner advises that customers "ask providers if they will commit to storing and processing data in specific jurisdictions, and whether they will make a contractual commitment to obey local privacy requirements on behalf of their customers. (Brodkin, 2008, 1)

One of a company's biggest assets is its competitive advantage. In order to protect that advantage, trade secrets, strategic planning and daily operations information must be protected from current and future competitors who may also be customers of the cloud provider. "Find out what is done to segregate data at rest," Gartner advises. The cloud provider should provide evidence that encryption schemes were designed and tested by experienced specialists. "Encryption accidents can make data totally unusable, and even normal encryption can complicate availability," Gartner says. (Brodkin, 2008, 2)

The recent earthquake and tsunami in Japan has affected businesses worldwide. The impending nuclear crisis and the destruction of its ports is crippling to the local population and the economy. The affects of the crisis has also affected companies abroad. The shortage of parts from Japan has caused several American manufacturers to decrease or discontinue production indefinitely. This is just an example of how the global market affects business. As it relates to cloud computing, companies must work with their service providers to plan for disaster recovery. Customers should know what exactly would happen to their date in the event of a disaster. "Any offering that does not replicate the data and application infrastructure across multiple sites is vulnerable to a total failure," Gartner says. Ensure that the provider has "the ability to do a complete restoration, and how long it will take." (Brodkin, 2008, 2)

An often sad but very real fact of life is that all good things must come to an end. Good companies with good ideas close their doors on a daily basis, whether it is the result of improper management, an increase in competition or a decrease in demand. It is the customer's responsibility to ensure that they are protected in the event that such situations occur. Gartner suggests that you "ask potential

providers how you would get your data back and if it would be in a format that you could import into a replacement application". (Brodkin, 2008, 2)

In addition to the security concerns raised by Gartner, cloud subscribers should also perform due diligence in evaluating the service level agreement (SLA). The SLA is a contract of guarantee between the cloud provider and the customer, which specifies the level of service that will be provided.

Cloud Computing Technologies

Cloud computing involves the sharing and coordinated use of diverse resources in distributed organizations. As per Zhidong Shen and Qiang Tong the authors of the paper entitled, The Security of Cloud Computing System enabled by Trusted Computing Technology, "Cloud computing provides a facility that enables large-scale controlled sharing and interoperation among resources that are dispersedly owned and managed" (Shen & Tong, 2010. 1). When considering the adoption of cloud computing, a company must also take into consideration the myriad of hardware, network, and infrastructure configurations that can be leveraged. Any combination of equipment and services may be utilized, but the determining factors will be the company's needs and the benefits and limitations of the technology. The following information will provide the characteristics of some of the technologies commonly utilized in a cloud computing environment.

Hardware

Mobile (Laptops, PDAs, Smartphones)

Mobile clients have security and speed concerns. Because the clients will be connecting to the cloud from various locations that may not have an optimized connection, you can not expect the speed that

a desk-bound client will achieve. All applications do not need speedy connections, and mobile users probably are not inputting gigabytes of data into the database. Mobile clients are easier to lose or misplace and the information on it can be compromised. On the other hand, if the data is maintained on the cloud and the user only has selected files on the device, only a minimal amount of data will be compromised. (Velte, Velte & Elsenpeter, 2010. 92)

Thin Clients

Thin client computers have no harddrives, no DVD ROM drives, and simply display what is on the server. If a client needs to access cloud based services only, or is accessing a virtualized server, the thin clients are a great option. They are less expensive than thick clients, are much less expensive to maintain, and use less energy (Velte, Velte & Elsenpeter, 2010. 92).

Thick Clients

Thick clients are good choices if the users need to maintain files on their own machines or run programs that do not exist on the cloud. Thick clients are more vulnerable to attack than thin clients because data is stored on the machine's hard drive. If the machine is stolen, data can be compromised. In terms of reliability, if a thin client fails, you can plug another thin client in and the user's work environment is right there. If a thick client fails, whatever data is stored on the machine, including the operating system, and all the configuration settings, is lost and a new computer will have to be configured for the user (Velte, Velte & Elsenpeter, 2010. 93).

Network

Basic Public Internet

The public Internet is the most basic method for connecting to the cloud. This type of access can be purchased from an Internet service provider (ISP) and connect with broadband or dial-up. When using this method, organizations should consider subscribing to multiple ISPs. Cloud providers should also purchase bandwidth from multiple sources. Ideally, the client would get bandwidth from one of the same ISPs as the vendor, increasing speed and reliability (Velte, Velte & Elsenpeter, 2010. 101).

The Accelerated Internet

Employing advanced application delivery features on top of your Internet connection can benefit both the service provider and the client. Cloud improvements can increase by 20 percent to 50 percent by offloading network related functions from the server. This method is mostly oriented toward the cloud service provider, but ultimately benefits the end user as well. With this method, customers should pay attention to the bandwidth charges. The accelerated method will require the service provider to install a server-side client and the end user will need to install a downloadable client (Velte, Velte & Elsenpeter, 2010. 102).

Optimized Internet Overlay

Optimized Internet Overlay allows customers to access the cloud via the public Internet, but enhancements occur on the provider's cloud. Enhancements at these points of presence (POP) include:

- Optimized real-time routing. Helps to avoid slow down, helping to make SLAs easier to obtain.
- An SSL session can be stopped so that protocols and payload can be optimized and re-encrypted.
- Some of the application can reside on the POP. This allows for better scalability, fault tolerance, and response time, usually in excess of 80 percent.
- Content that is frequently accessed can be delivered from local caches.

An Optimized Internet Overlay is costly and there is a strong vendor lock-in if the application is distributed into the carrier's network (Velte, Velte & Elsenpeter, 2010. 102).

Site to Site VPN

Customers can connect directly to the service provider using a private wide area network (normally an MPLS/VPN Connection). This method allows confidentiality, guaranteed bandwidth, and SLAs for availability, latency, and packet loss. MPLS can scale to meet changing bandwidth needs. In addition, quality of service can be written into the SLAs. On the downside, private WANS are not normally more reliable than Internet connections, especially redundant connections to multiple ISPs (Velte, Velte & Elsenpeter, 2010. 103).

Infrastructure

Virtualization

In a Virtualized environment, applications run on a server and are displayed on a client. With operating system virtualization, a single piece of hardware can be used to run multiple operating system images at the same time. Network virtualization is a method of combining the resources in a network by splitting up the available bandwidth into channels, each of which is independent from the others, and each of which can be assigned to a particular server or device in real time. Storage virtualization is the pooling of physical storage from multiple network storage devices into what appears to be a single storage device that is managed from a central console. Storage virtualization is commonly used in storage area networks (SANs). Server virtualization is the masking of server resources from server users with the intention of allowing them to share resources and utilization without having to understand the complexities of the server resources (SearchServerVirtualization.com, 2000).

Open hypervisor Standards

Hypervisors are the foundational component of virtual infrastructure and enable computer system partitioning. An open-standard hypervisor framework can benefit customers by enabling innovation across an ecosystem of interoperable virtualization vendors and solutions.

Collaboration around open hypervisor standards is expected to focus on the following areas of interoperability and performance optimization for virtualized:

- Cross-platform frameworks that govern the standardized operation and management of stand-alone virtual machine environments as well as highly dynamic, data center-scale deployment of virtualized systems
- Cooperative virtualization APIs between hypervisors and guest operating systems
- Virtual machine formats that enable virtual machine migration and recovery across platforms

(Velte, Velte & Elsenpeter, 2010. 162)

Community Source

The community source program provides industry partners with an opportunity to access VWware ESX source code under a royalty-free license. Partners can contribute shared code or create binary modules to spur and extend interoperable and integrated virtualization solutions. The idea is to combine the best of both traditional and open source development models. Community members can participate and influence the governance of VMware ESX Server through an architecture board. This approach will help to build differentiated, intellectual property-protected solutions (Velte, Velte & Elsenpeter, 2010. 163).

OVF

As a result of VMware and its industry partners' efforts, a standard has been developed called the Open Virtualization Format (OVF). OVF describes how virtual appliances can be packaged in a vendor neutral format to run on any hypervisor. OVF is a platform-independent, extensible, and open specification for the packaging and distribution of virtual appliances composed of one or more virtual machines. OVF gives customers and developers the choice to select any hypervisor based on price, preference, or functionality, and it prevents vendor lock-in. This standard packaging and

distribution format for virtual appliances will be important in accelerating the adoption of virtual appliances (Velte, Velte & Elsenpeter, 2010. 163).

Is Cloud Computing Right for Your Organization

The fiduciary duties of an organization's members extend to all stakeholders, including its customers, who entrust these individuals with the care of their information. The sensitivity of various data managed by a company may vary, but the importance of integrity in caring for that data remains the same. Before venturing into cloud computing, a company must fully understand their business, their data, and their customers. The company's responsibilities to its customers should not be clouded by the opportunity for financial gain. In an article entitled, 6 Tips for Better Cloud Computing Integration for Business in 2011, the author states "Cloud computing for business can offer substantial benefits, usually in term of increased productivity and lowered operational-related costs. However, mismanaging the cloud can be disastrous" (Ivan, 2010).

Careful consideration should be given to the sensitivity and criticality of the data that is maintained and utilized by a company's members. Data covered by the Health Insurance Portability and Accounting Act (HIPPA) should not be stored on cloud. If your company has data that is regulated such as HIPA or Sarbanes Oxley, be very careful in your plans to place data on the cloud (Velte, Velte & Elsenpeter, 2010. 25 -26).

The adoption of cloud computing may not prove to be a good venture for a company that has applications requiring specific hardware, chips, or drivers. It is unlikely that the service provider will have the precise hardware that will be needed in every case. In a cloud computing venture, the

customer is subject to the capabilities and technologies that the provider offers. Even if the provider agrees to adjust his offerings to incorporate new technology, the customer will assume great risk of potential issues as the provider attempts to mature in his knowledge of this area (Velte, Velte & Elsenpeter, 2010. 27).

The adoption of cloud computing also involves a certain loss of control. "If your applications demand complete control over everything that is running, a cloud solution may not be right for you. If you require detailed control over the amount of memory, CPU, hard drive specs, or interfaces, then the cloud isn't an appropriate match for your application" (Velte, Velte & Elsenpeter, 2010. 27).

When a company has multiple applications that need to integrate, an arrangement in which one of them runs on the cloud while the others run locally, could lead to issues with security, speed, and reliability. Often applications that run on the cloud extract, transport, and load data this is stored, or at least comingles, with sensitive information in a company's databases. The theory prevails that any interaction with sensitive data increases its vulnerability to the cloud. High speed applications running in-house that are reliant upon data from the cloud will only run as fast as the cloud allows, leading to questionable reliability. In addition, data traveling between local and remote applications increases the company's risk that it may be compromised or damaged (Velte, Velte & Elsenpeter, 2010. 28).

In cloud computing, data and applications are located on a series of servers, geographically disparate from the customer's site. With the added distance it is going to take some time to transfer the data between the two locations. If the data is required instantaneously, the cloud might not be your best option (Velte, Velte & Elsenpeter, 2010. 28).

Since cloud computing is generally billed in a utility format, companies pay for what they use. The cost begins to increase when applications are deployed that use a lot of throughput. If your company streams high-definition video over many sources, your costs are going to spike sharply. Take into account what a server, power, and all other hardware will cost. Figure in the price of management and associated IT personnel costs and then compare that with what a service provider will charge you. If it is cheaper to buy the server, it might be best to delay a move to the cloud for now. If the cost is the same, you need to ask yourself what business you want to be in (Velte, Velte & Elsenpeter, 2010. 29).

"In some cases the applications themselves are not ready to be used on the cloud. They may have quirks that prevent them from being used to their fullest abilities, or they may not work whatsoever." The application might require a lot of bandwidth to communicate with users and cloud computing is paid based on how much you use. The application might take a great deal of effort to integrate with other applications. If the application has to talk with a database onsite, it may be better to host the application locally until you can move the entire infrastructure to the cloud. This helps to avoid the service cost of having to transfer to and from the cloud. This method is also more efficient because the application can talk to the database without having to reach across the network to do so (Velte, Velte & Elsenpeter, 2010. 31).

In some circumstances, applications may not communicate securely across the Internet, which puts your data is at risk. Before pursuing a cloud-based solution, ensure that your application is compatible with a variety of web browsers and will allow the users to conduct transactions using encryption. If the application's results cannot be displayed securely, your company's information will be placed in harm's way (Velte, Velte & Elsenpeter, 2010. 33).

"Some cloud service offerings result in the creation of data, reports, research, statistics, or other combination of data for the customer. Determining who created the information or who is the author of a particular report or compilation of information might be blurred. When one party contributes some data and the other contributes a tool and perhaps other data, who owns the resulting product? Thus, it would be prudent for the parties to ensure that the rules for, and allocation of, ownership of the data be clearly delineated (Gilbert, 2010)."

Despite the opportunities presented by cloud computing for an organization, managing the cloud is a complex and challenging process. Consideration must be given to the issues of scalability and security by having your business cloud accessed by every employee you have. Ineffective and inefficient usage of cloud services could lead to more costs to your business and more risk of security troubles (Ivan, 2011).

Best Practices

To mitigate the risks associated with cloud computing, an organization must establish and continuously adhere to a well-defined set of processes and standards. These practices will govern vendor relationships, policies, and procedures to ensure that all activities involved with this venture are in the best interest of the company. The ultimate goal is to take advantage of the great benefits that cloud computing has to offer while protecting the company's image and its assets. "Remember, your company information is in the hands of someone outside your company wall and unless the correct processes and performance monitoring are in place, your critical data it is at the mercy of the cloud" (Tech Journal South, 2010).

The "Security Guidance for Critical Areas of Focus in Cloud Computing V2.1," published by the Cloud Security Alliance does a great job of outlining very significant aspects of cloud computing, of which, all companies should be aware before using cloud services. The remainder of this section will explore some of the recommendations offered in this document, which all companies should incorporate into their best practices for cloud computing.

Governance and Enterprise Risk Management

Well-developed information security governance processes should result in information security management programs that are "scalable with the business, repeatable across the organization, measurable, sustainable, defensible, continually improving, and cost-effective on an ongoing basis." Both Cloud Computing service customers and providers should develop robust information security governance designed to achieve agreed-upon goals, which support the customer's business objectives. The service model will define the roles and responsibilities in information security governance and risk, while the deployment model defines accountability and expectations (Cloud Security Alliance, 2009. 31 - 32).

Organizations should review the specific information security governance structure and processes, in addition to specific security controls, as part of their due diligence in evaluating prospective provider organizations. The provider's security governance processes and capabilities should be assessed for sufficiency, maturity, and consistency with the user's information security management processes. The provider's information security controls should clearly support these management processes while averting risks (Cloud Security Alliance, 2009. 32).

Metrics and standards for measuring the performance and effectiveness of information security management should be established prior to moving into the cloud. Organizations should understand

and document their current metrics and expectations for variances in performance that may be experienced when operations are moved into the cloud. Wherever possible, security metrics and standards, particularly those relating to legal and compliance requirements should be included in the Service Level Agreements and contracts (Cloud Security Alliance, 2009. 32).

As with any new business process, it is important to follow best practices for risk management. The practices should be proportionate to your particular usages of cloud services, whether it involves basic data processing or mission critical business processes dealing with highly sensitive information. In considering the use of cloud services for functions critical to the organization the risk management approach should include the following

- Identification and valuation of assets
- Identification and analysis of threats vulnerabilities and their potential impact on assets
- Analysis of the likelihoods of events and scenarios
- Management-approved risk acceptance levels and criteria
- The development of risk treatment plans with multiple options (control, avoid, transfer, accept).

The outcomes of risk treatment plans should be incorporated into service agreements. Asset inventories should account for assets supporting cloud services and under the control of the provider. Asset classification and valuation schemes should be consistent between user and provider (Cloud Security Alliance, 2009. 32).

The service, and not just the vendor, should be the subject of risk assessment. The use of cloud services, and the particular service and deployment models to be utilized, should be consistent with

the risk management objectives of the organization, as well as with its business objectives (Cloud Security Alliance, 2009. 33).

Third Party Management Recommendations

Customers should view cloud services and security as supply chain security issues. Due diligence should be performed in examining and assessing the provider's supply chain (service provider relationships and dependencies), to the extent possible. Customers must also examine the provider's own third party management. Assessment of third party service providers should specifically target the provider's incident management, business continuity and disaster recovery policies. The assessment should also include processes and procedures along with a review of co-location and backup facilities. This should include review of the provider's internal assessments of conformance to its own policies and procedures, and assessment of the provider's metrics to provide reasonable information regarding the performance and effectiveness of its controls in these areas (Cloud Security Alliance, 2009. 34).

Legal and Electronic Discovery

Cloud computing creates new dynamics in the relationship between an organization and its information, involving the presence of a third party: the cloud provider. This creates new challenges in understanding how laws apply to a wide variety of information management scenarios. A complete analysis of Cloud Computing-related legal issues requires consideration of functional, jurisdictional, and contractual dimensions (Cloud Security Alliance, 2009. 35).

The functional dimension involves determining which functions and services in Cloud Computing have legal implications for participants and stakeholders. The jurisdictional dimension involves the

way in which governments administer laws and regulations affecting cloud computing services, the stakeholders, and the data assets involved. The contractual dimension involves the contract structures, terms and conditions, and enforcement mechanisms through which stakeholders in Cloud Computing environments can address and manage the legal and security issues (Cloud Security Alliance, 2009. 35).

Information Lifecycle Management

One of the primary goals of information security is to protect the fundamental data that powers systems and applications. As companies transition to Cloud Computing, they quickly realize that traditional methods of securing data are challenged by cloud-based architectures. Elasticity, multi-tenancy, new physical and logical architectures, and abstracted controls require new data security strategies.

The Data Security Lifecycle is different from Information Lifecycle Management, reflecting the different needs of the security audience. The Data Security Lifecycle consists of six phases:

Figure 3 – Data Security Lifecycle (Cloud Security Alliance, 2009)

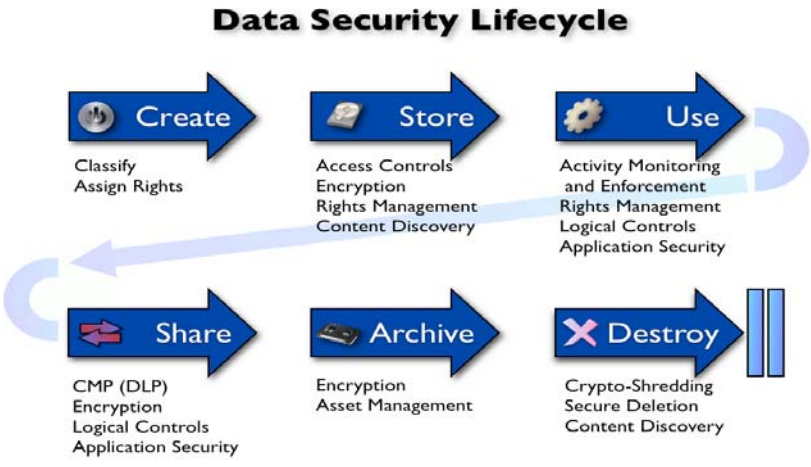

Key challenges regarding data lifecycle security in the cloud include the following:

- **Data security.** Confidentiality, Integrity, Availability, Authenticity, Authorization, Authentication, and Non-Repudiation.
- **Location of the data.** There must be assurance that the data, including all of its copies and backups, is stored only in geographic locations permitted by contract, SLA, or regulation.
- **Data remanance or persistence.** Data must be effectively and completely removed to be deemed 'destroyed'. Therefore, techniques for completely and effectively locating data in the cloud, erasing/destroying data, and assuring the data has been completely removed or rendered unrecoverable must be available and used when required.
- **Commingling data with other cloud customers**. Data, especially classified or sensitive data, must not be commingled with other customer data without compensating controls while in use, storage, or transit. Mixing or commingling the data will be a challenge when concerns are raised about data security and geo-location.
- **Data backup and recovery schemes for recovery and restoration.** Data must be available and data backup and recovery schemes for the cloud must be in place and effective in order to prevent data loss, unwanted data overwrite, and destruction. Do not assume cloud-based data is backed up and recoverable.
- **Data discovery.** As the legal system continues to focus on electronic discovery, cloud service providers and data owners will need to focus on discovering data and assuring legal and regulatory authorities that all data requested has been retrieved. In a cloud environment that question is extremely difficult to answer and will require administrative, technical, and legal controls when required.

- **Data aggregation and inference.** With data in the cloud, there are added concerns of data aggregation and inference that could result in breaching the confidentiality of sensitive and confidential information. Hence practices must be in play to assure the data owner and data stakeholders that the data is still protected from subtle "breach" when data is commingled or aggregated, thus revealing protected information (e.g., medical records containing names and medical information mixed with anonymous data but containing the same "crossover field") (Cloud Security Alliance, 2009. 40 - 41).

Encryption and Key Management

Cloud customers and providers need to ensure that their assets are guarded against data loss and theft. Today, encryption of personal and enterprise data is strongly recommended, and in some cases mandated by laws and regulations around the world. Cloud customers want their providers to encrypt their data to ensure that it is protected no matter where the data is physically located. Likewise, the cloud provider needs to protect its customers' sensitive data (Cloud Security Alliance, 2009. 60).

Strong encryption with key management is one of the core mechanisms that Cloud Computing systems should use to protect data. While encryption itself does not necessarily prevent data loss, safe harbor provisions in laws and regulations treat lost encrypted data as not lost at all. The encryption provides resource protection while key management enables access to protected resources (Cloud Security Alliance, 2009. 60).

There is the utmost need to encrypt multi-use credentials, such as credit card numbers, passwords, and private keys, in transit over the Internet. Although cloud provider networks may be more secure

than the open Internet, they are by their very architecture made up of many disparate components, and disparate organizations share the cloud. Therefore, it is important to protect this sensitive and regulated information in transit even within the cloud provider's network (Cloud Security Alliance, 2009. 60).

Encrypting data on disk or in a live production database has value, as it can protect against a malicious cloud service provider or a malicious co-tenant as well as against some types of application abuse. For long-term archival storage, some customers encrypt their own data and send it as ciphertext to a cloud data storage vendor. The customer then controls and holds the cryptographic keys and decrypts the data, if necessary, back on their own premises (Cloud Security Alliance, 2009. 60).

Encrypting data at rest is common within IaaS environments, using a variety of provider and third party tools. Encrypting data at rest within PaaS environments is generally more complex, requiring instrumentation of provider offerings or special customization. Encrypting data at rest within SaaS environments is a feature cloud customers cannot implement directly, and need to request from their providers (Cloud Security Alliance, 2009. 60).

Identity and Access Management

"Security researchers repeatedly label end users the biggest threat to enterprise security. Unlike applications that can be patched or systems that can be hardened, end users -- whether through naiveté, carelessness, or malicious intent -- continue to expose IT resources to serious security threats (DeCarlo, 2007)". Managing identities and access control is one of the greatest challenges

facing information technology. Although this can prove to be a menial task in certain cases, the proper management of user access should not be taken lightly.

Conclusion

Although cloud computing has been the hot topic in technology circles for the past few years, this is still a fairly new paradigm in information technology. Many of its nuances are yet to be explored because the technology has not matured fully and is not utilized to its fullest capacity by cloud customers. A few examples of security breaches have been publicized but the majority of them resulted from infiltration by way of phishing, malware, etc., in which users' IDs and passwords were stolen and used to access the system. Companies have experienced system downtime and valuable information has been stolen, but these events usually occur on a small scale. Security breaches will likely become common as more companies move to the cloud. Detailed planning and the implementation of stringent policies and procedures can mitigate the risks.

Companies are eager to take advantage of the benefits offered by cloud computing service providers, but they are advised to "step lightly" as they incorporate this technology into their business functions. As J.A. Zachman stated, "The cost involved and the success of the business depending increasingly on its information systems require a disciplined approach to the management of those systems." Before venturing into the cloud, companies must first ensure that they know their business. The structure, requirements, and condition of their current infrastructure, applications, and equipment must be analyzed and fully understood. What are the needs of the business? What essential requirements must be met fully in order for the business to operate properly? How critical to my business and my customers are the applications and data that I will be entrusting to a service

provider? Is it legal to allow a third party to store the information that my business utilizes and manages? These and other questions must be answered before considering a venture into cloud computing.

References

1. Rittinghouse, John W. and Ransome, James F. Cloud Computing: Implementation, Management, and Security. Florida: Taylor and Francis Group, LLC, 2010. Print
2. Gilbert, Francois. Cloud Service Contracts May be Fluffy: Selected Legal Issues to Consider Before Taking Off. New York: Aspen Publishers, Wolters Kluwer Law & Business, 2010. White Paper
3. The Seven Standards of Cloud Computing Service Delivery. SalesForce.com, 2009. White Paper
4. Brodkin, Jon. Gartner: Seven cloud-computing security risks, 2008. Retrieved from http://www.infoworld.com/d/security-central/gartner-seven-cloud-computing-security-risks-853.
5. Velte, Anthony T., Velte, Toby J., Elsenpeter, Robert. Cloud Computing: A Practical Approach. The McGraw-Hill Companies, 2010. Print
6. Tech Journal South. Six Cloud Pitfalls to Avoid in 2011, 2010. Retrieved from http://www.techjournalsouth.com/2010/12/six-cloud-pitfalls-to-avoid-in-2011.

7. Cloud Security Alliance. Security Guidance for Critical Areas of Focus in Cloud Computing V2.1, 2009. White Paper. Retrieved from http://www.cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf
8. DeCarlo, Amy Larsen. How to protect against naive, careless or malicious users, 2007. Retrieved from http://www.computerworld.com/s/article/9013618/Biggest_security_threat_Your_users
9. Zachman, J.A. A Framework for Information Systems Architecture. International Business Machines Corporation, 1987.
10. Ivan. 6 Tips for Better Cloud Computing Integration for Business in 2011. Cloud Business Review, 2010. Retrieved from http://www.cloudbusinessreview.com/2010/12/25/6-tips-for-better-cloud-computing-integration-for-business-in-2011.html
11. Ivan. How to Fail in Cloud Computing Deployment. Cloud Business Review, 2011. Retrieved from http://www.cloudbusinessreview.com/2011/03/05/how-to-fail-in-cloud-computing-deployment.html
12. Shen, Zhidong and Tong, Qiang. The Security of Cloud Computing System enabled by Trusted Computing Technology. 2010 2nd International Conference on Signal Processing Systems (ICSPS). Research Paper.
13. SearchServerVirtualization.com. 2000. Retrieved from http://searchservervirtualization.techtarget.com/definition/virtualization
14. Earnst & Young. 13th Global Information Security Survey 2010. Cloud computing: pros and cons. Retrieved from http://www.ey.com/GL/en/Services/Advisory/IT-Risk-and-Assurance/13th-Global-Information-Security-Survey-2010---Cloud-computing--pros-and-cons