Cumulonimbus Computing Concerns

Information Security in Public, Private, and Hybrid Cloud Computing

Daniel Adams

A Senior Thesis submitted in partial fulfillment
of the requirements for graduation
in the Honors Program
Liberty University
Spring 2015

Acceptance of Senior Honors Thesis

This Senior Honors Thesis is accepted in partial
fulfillment of the requirements for graduation from the
Honors Program of Liberty University.

_____
Mark Shaneck, Ph.D.
Thesis Chair

_____
Terry Metzgar, Ph.D.
Committee Member

_____
David Wang, Ph.D.
Committee Member

_____
Brenda Ayres, Ph.D.
Honors Director

_____
Date

Abstract

Companies of all sizes operating in all markets are moving toward cloud computing for

greater flexibility, efficiency, and cost savings. The decision of *how* to adopt the cloud is

a question of major security concern due to the fact that control is relinquished over

certain portions of the IT ecosystem. This thesis presents the position that the main

security decision in moving to cloud computing is choosing which type of cloud to

employ for each portion of the network – the hybrid cloud approach. Vulnerabilities that

exist on a public cloud will be explored, and recommendations on decision factors will be

made for which specific types of systems to harbor inside a private cloud. Picking the

best location for each system allows risk to be managed and sensitive information to be

protected while at the same time providing a cost effective option.

Cumulonimbus Computing Concerns

Information Security in Public, Private, and Hybrid Cloud Computing

Purpose

As personal data is being moved into and processed by computing systems on the Internet, security is paramount. For most people, data that is the key to their identity, financial, and social wellbeing is in the hands of numerous companies and these companies' computer systems. These entities range from banks to insurance companies to stock markets to even retail stores. All these companies have a duty to protect their customers by protecting their data.

To fully understand the issue at stake, this thesis will first dive into the issue of what defines a cloud and then examine how organizations would benefit from moving to the cloud in the first place. Defining what constitutes a cloud is important because then a company can realize how to reap the benefits of cloud computing without necessarily deciding to utilize a pre-packaged, public cloud offering. Once it is established that certain benefits can be gained from using the cloud for computing resources, this paper will raise some cautions about security risks in the cloud world, especially with regards to public clouds. Then the conclusion will give suggestions about how to both reap the benefits from the cloud model while at the same time how not to substantially increase risk in the area of security.

This thesis will take the position that the main security decision of cloud computing is which type of cloud, namely deployment model and service model, will be used. The decision can eliminate whole categories of security risks but at the same time increase responsibilities of the organization. Achieving the correct balance with this issue

can be accomplished through the use of hybrid clouds with strong segmentation between the public and private components. Also, this paper will expound upon some of the best practices to attempt to mitigate the security concerns if public cloud is the choice taken.

<div align="center">**Background**</div>

The starting point in this discussion of cloud computing security prompts the explanation of what *is* cloud computing, exactly. To begin, this thesis will look at the aspects of cloud computing as defined by the NIST and then will move on to how these aspects can be beneficial to companies leveraging the cloud.

NIST Definition

The National Institute of Standards and Technology defines cloud computing as follows:

> Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. (NIST, 2011)

This definition captures the essential characteristics of cloud computing while not excluding legitimate cloud computing manifestations that might not come to mind right away. The first thing most people think of in relation to the words "cloud computing" is along the lines of Amazon Web Services' EC2, which stands for Elastic Compute Cloud. Other forms of cloud computing are vastly different in terms of business model and deployment, but they still fit the essential definition. For example, a business could make an on-premise private cloud in which they purchase all the hardware themselves instead

of renting space from a cloud provider. As long as the 5 qualifications of on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service are met, then the deployment qualifies as cloud computing (NIST, 2011). The importance of having the correct picture of what actually defines a cloud is that a company can evaluate if they can provide those money-saving and agility-providing qualities themselves instead of having to outsource to a possibly less secure organization.

On-Demand Self-Service. This includes provisioning and configuring computing resources, whether they be virtual machines, storage pools, or network topology and routing. One of the benefits of the cloud is fluidity with regards to the ability to deploy certain programs on specific computing and network architectures. This characteristic of cloud computing is what realizes that benefit.

Broad Network Access. Cloud computing provides services which are accessed and used over the network. Other than datacenter maintenance, there is not a need to physically be in the same location as the servers making up the cloud. Whether the resource is an RDP connection, database connection, or HTTP session; they are all accessed over the network. This quality provides enormous flexibility in how a cloud is implemented while keeping the same end-user view.

Resource Pooling. One of the major reasons to move to cloud system usage is the increased utilization and server density gained by resource pooling. This drives down initial purchase cost, ongoing electrical usage cost, and datacenter size (Pearce, Zeadally, & Hunt, 2013). Resource pooling also easily accommodates vastly different sized instances of servers and applications in terms of capacity. For example, a development

server doesn't need nearly the resource allocation that a production version of the same server averaging 1,000 concurrent users would require.

Rapid Elasticity. Resource pooling goes hand in hand with elasticity. Because the cloud is virtualized, boundaries of where one physical server starts and another ends are not hard constraints in provisioning systems. This gives the possibility for servers and applications to be dynamically resized.

Measured Service. Cloud services are provisioned for performance management and are priced (if renting from a cloud provider) by the amount of resources allocated to a specific host, network, or business account. In order to effectively accomplish this, a metering service must exist to both control the amount of resources allocated and to record the resources used. Even if the company owns the cloud rather than rents it, measured service is still needed to accurately provision the desired resource allocation.

Reasons to Move to the Cloud

The characteristics of a cloud computing system discussed in the previous section translate into major benefits in the areas of cost, flexibility, and performance. This section will look at the issue of moving to the cloud more from a business approach as well as from an operations view. The business focus will be primarily on issues such as the following: "how can we keep costs down?" and "will we be able to add this new web feature marketing wants?" Operations' viewpoint is more concerned with moving applications between development and production networks with differing security approaches, continuous integration, and similar issues. The flexibility the cloud offers can be a win for both teams.

Cost Savings. The primary reason some companies move to a cloud model is for cost savings due to higher efficiency usage of computing resources. Virtualization allows more than one virtual server for each physical server, thus increasing utilization while at the same time allowing flexibility for the times when one virtual server's resource needs spike. Dell states in their report, "Evaluating the Useful Life of a Server," that servers with the capabilities for virtualization do cost slightly more (around 22%) due to improved components to handle the additional computing needs of the virtual servers, but make that cost up with savings in other areas. The author states that "their higher energy efficiency and smaller footprint, however, meant 27% less overall facilities cost, and even modest virtualization improvements reduced the number of servers by half" (Nolle, 2012).

Cloud services need not be necessarily on-site or completely off-site utilizing a hosting provider, but a detailed look at total cost of ownership (TCO) is necessary when evaluating which is better. According to the same report, the median cost of a server is $1500-1600, but the power/cooling and maintenance costs were $790 and $940 respectively over the server's four year lifespan. These add up to $1730, about $200 more than the cost of the server itself. These costs, plus the additional personnel salaries for a datacenter and virtualization management team, drive up the operational price of the servers even further.

Third-party cloud service providers manage all these aspects and charge the subscriber a fixed rate for the use of certain amounts of provisioned virtualized resources. Because cloud companies have put vast effort into streamlining and automating processes as well as splitting those overhead costs among many tenants, they can offer services at a

more cost-efficient rate than that of most other companies trying to do the job themselves.

Additionally, large startup costs can be avoided and small amounts can be paid along the

way. "Cloud computing will have a profound impact on the cost structure of all the

industries, turning some of the fixed costs into marginal costs of production" (Mitchell &

Meggison, 2014). That idea brings the discussion to the next aspect, agility.

Agility. The second reason companies find moving to the cloud attractive is for

increased agility in their operations. This can happen in several ways. One is flexibility in

using computing resources rather than being required to own all computing resources and

only have some utilized at certain times of peak business. Another is giving up both

control and management of other areas of the software stack. In that case, a company

decides to use Software as a Service (SaaS) or Platform as a Service (PaaS) instead of

Infrastructure as a Service (IaaS). An example of using SaaS is a company signing a

contract with Microsoft to use Outlook.com email rather than hosting and managing

email services on their own. More detail about these will be provided in the section below

titled "Service Models."

Most pieces of software designed to be accessed by many people at the same time

will need some sort of scalability. This is especially true of modern web apps like Gmail

where the number of concurrent users might be in the millions. Even the most powerful

single server build could not handle that load, so horizontal scalability (as opposed to

vertical scalability) is what must be implemented. With horizontal scalability, new servers

are added as the load increases, and the load is spread out among them. Clearly the cloud

is the perfect methodology for realizing this goal. A virtual machine can be cloned,

provisioned, deployed, and added to the main load balancer in an automated fashion, and

almost instantly a new level of capacity is added to the application. Digital Ocean, a

public cloud provider aimed at developers, touts, "Deploy an SSD cloud server in 55

seconds" (2015). This type of quick deployment can be used for load following

applications.

Besides the time-based agility cloud services provide, third-party cloud hosting

providers can also provide resource elasticity beyond what a company has in their

datacenter on a regular basis. For example, a small online retailer trying to make a big

marketing push on Black Friday might need four times the normal resource allotment.

Financially it doesn't make sense to incur a large expense to scale up infrastructure for

less than one week each year. At non-peak times those computing resources are still

costing the company money but are going to waste. A company leveraging the cloud

would pay to rent those systems only for the one day sale and then scale back to their

normal business level for the rest of the year. This is a much better solution. Now that the

benefits of cloud computing have been shown, the options relating to cloud types will be

examined.

Types of Cloud Computing

The position of this thesis is that the most fundamental security decision when

using the cloud is picking which of the many types of cloud offerings to use for a specific

purpose. Decisions in the area of service model and deployment model can mitigate

entire categories of risks; but at the same time, they require certain responsibilities to

mitigate other risks. At a very high level, the service model determines how much of the

software stack a client company controls, and the deployment model choice dictates how

much of the physical computing hardware the company oversees. Both of these control

decisions are security related issues.

Service Models

There are three service models that cloud computing can follow. From a technical

perspective, most types of applications can be deployed in any of the three methods. It

then becomes a business decision which one a company chooses to use depending on

their in-house expertise, the amount of human resources they can dedicate to the problem,

and the amount of security and control needed around the data associated with the app.

Moving from one side of the spectrum to the other, there is a trend of more to less control

over the underlying architecture. For specific purposes this can either be a good or a bad

thing. Only needing to configure the application, or even having to write custom code, is

a much more focused scope than managing the server-level infrastructure on which the

application runs. Figure 1 below gives an understanding of how each of these models

differs from the others.

Figure 1: Cloud Hosting Models was removed for open source publishing. See figure 1-2 in *Windows Azure Hybrid Cloud* (Garber, Malik, & Fazio, 2013).

Infrastructure as a Service (IaaS). Infrastructure as a Service is possibly the

most well-known iconic type of service when someone thinks of "the cloud." Basically it

is provisioned virtual servers that the customers can log into and configure however they

would like within certain limitations. Functionally IaaS acts much the same way as a

physical server sitting in the company's datacenter. Patch management, application

runtime environment installation and setup, and application installation are all done by

the customer. The portions managed by the vendor are only the datacenter virtualization

aspects, networking, and usually OS imaging on deployment. All three of these service

models will have tradeoffs in control vs. time and expertise spent doing management. IaaS is fully on the control and customization side; the customer is in charge of managing most everything about the virtual host.

The book *Windows Azure Hybrid Cloud* by Danny Garber, Jamal Malik, and Adam Fazio presents a good analogy for each of the service models by comparing them to methods of transportation. Starting off, buying one's own car is the traditional corporate datacenter approach. The company is in charge of research and procurement of specific models of servers, networking gear, and datacenter equipment (racks, cooling systems, etc.). One step away from that is leasing a car, the IaaS cloud option. With a lease, you have basically the same responsibilities as if the car were your own, but it isn't. Another parallel with IaaS and leasing is the asset depreciation. Just as cars have a limited lifetime, servers do as well. Leasing a server for two years will take away almost half of its usable life, so your lease cost covers their replacement fund, server hardware maintenance, network connection, and people to manage the hardware layer of infrastructure (i.e. replacing failed hard drives in a SAN.) The point of the cloud is to manage the provisioning in an automated fashion as much as possible so salaried humans don't have to do the job.

Platform as a Service (PaaS). Platform as a Service is the level on which developers work. Referring to Figure 1 again, it includes the network/storage layer, the virtualization layer, the operating system layer, as well as the runtime layer. This allows custom code to be easily deployed to the cloud server. The hosting provider's cloud management solution takes care of routing and isolating the tenants, managing the underlying runtime engines, as well as sometimes supporting dynamic, load-based

scalability. The customer has complete control over the code he or she deploys, but that is

basically the only operation that is done on the infrastructure: deploy. The specific other

components of the system are provided, whether that is standard filesystem storage, SQL

databases, NoSQL databases, network I/O. IaaS generally allocates a specific amount of

allocated hardware (i.e. number of CPUs, amount of RAM, disk space) that is the client's

to do whatever he or she wants, where PaaS usually provides whatever services are

needed and charges a small charge for a certain number of operations. For example, with

Google App Engine, NoSQL database operations are $0.06 per 100k read or write ops

and $0.18/GB per month (Google, Inc., 2015). Examples of these PaaS providers are the

aforementioned Google App Engine, Red Hat OpenShift, Heroku, and Windows Azure.

The analogous operation to Platform as a Service in the transportation world is a

rental car. A person can pick it up and drive wherever he or she wants (deploy whatever

code), but at the same time doesn't have to worry about doing routine engine

maintenance and fixing emergency problems (like routine patching & configuration and

conflicts due to those actions). The cloud service provider manages those lower levels of

infrastructure in their own automated way. PaaS is the middle ground on the spectrum of

control versus handing off duties.

Software as a Service (SaaS). Software as a Service is the highest level approach

to cloud applications. A client will contract with a provider to have the provider's

software made available to the client, and the client will pay for the amount of usage that

will be required. Normally this is on a per-account or per-user basis. Many people will

sometimes think of these services as just "websites" when they, in reality, are Software as

a Service cloud apps. An example of an application that would fall into this category is

email hosting on Microsoft's Outlook.com. This approach takes most all of the technical

management concerns away from the client company and lets the hosting company take

care of these aspects. The client company's responsibility is to customize the software

settings made available by the SaaS provider. In the case of cloud-based email service,

those client-controlled aspects would be the Active Directory settings, group

organization, etc.

To complete the analogy, Software as a Service is similar in usage to a public

transit system. With public transit, the individual user does not have much control over

the scheduling or routes – he just has to make do with the limited flexibility that is

provided. He is left to choose which of the available routes and departure times works the

best for him. On the flip side, the ordeal is extremely simple – there are no worries about

how to repair the bus if it breaks or even any of the challenges of driving in a crowded

city.

Deployment Models

Any cloud service model discussed above can be realized on most any

deployment model explained in this section. Having the wide variety of combinations

available allows for picking the solution that best fits the situation. What is crucial in

cloud security is to select the right type of service and deployment model to fit the need

while still being secure. The remainder is following the security best practices associated

with the chosen service model.

Public Cloud. The general idea of a public cloud is each provisioned virtual

machine is on the same network as everybody else – every machine is exposed to the

Internet. Later this paper will discuss ways a public cloud can be secured, but in general

this is the least secure form of external party cloud hosting. On the flip side, it also has the lowest cost for a given amount of computing power.

Public cloud providers are often huge, owning hardware that would allow tens of thousands of virtual machines to run at the same time. These massive amounts of resources are available to whichever company or individual person would like to rent them. Consequently any given entity's "neighbors" in the cloud could be comprised of companies in the same industry, small startup companies, hobbyists, as well as hackers. Allowing hackers to get so close to important infrastructure is generally a very bad idea, but there are certain situations where the security requirements can still be met, and the benefit of reduced cost wins.

Figure 2, from the NIST's Cloud Computing Synopsis and Recommendations, depicts the groups of clients using the cloud resources as well as the public cloud logical setup. In a public cloud security has to be more weighted toward the host itself since the network environment around it is more hostile. This diagram shows how there can be a boundary controller (like a VPN concentrator) that would allow a secure location like a business branch office to have a more privileged level of access to the cloud computing resources than the level of access attainable from being beside the specific cloud-based host within the cloud network. Also the diagram calls out the maintenance the cloud hosting provider does on their own hardware (with the "new hardware in" and "old hardware out" arrows.) This would otherwise be functions the business would be required to perform but now of which are relieved.

Figure 2: Public Cloud was removed for open source publishing. See section 4, figure 1 from NIST Special Publication 800-146 at http://csrc.nist.gov/publications/nistpubs/800-146/sp800-146.pdf (NIST, 2012).

Private Cloud. The complete opposite side of the spectrum from a public cloud is a private cloud. This is a cloud for only one client organization, and that organization has full control over the cloud structure. Because there are no other clients using it, they can control all aspects of the network, firewalls, and other protections to their liking. Also, the cloud is under only the owning company's physical control, which reduces the possibility of certain types of attacks. This type of cloud can be made the most secure of any, but if the company's IT is not mature, they are getting rid of the cloud service provider, which might be a small last line of defense or alert in the case of an attack.

Private clouds are in essence a very high level of virtualization flexibility in a company's own datacenter. As a result, some of the benefits of the cloud are not displayed in as magnified a way as with a public cloud. This is because, unless the company is extremely large, the 1) management overhead is about the same as if cloud was not used, and 2) the datacenter will have a more limited cap in the amount to which it can scale at a moment's notice. In a public cloud there are so many tenants that it is very unlikely that all will experience a peak usage surge at exactly the same time. The public cloud company must have enough to handle its largest tenants' peak usage, so it will have a respectable amount of free resources at any given time. Other clients benefit from this availability. One the other hand, companies with their own datacenter might not have the money to buy extra hardware "for just in case" when there's not an immediate business reason to buy it.

That being said, private clouds not only have the advantage of the owner being able to configure just the way the owning organization would like but also the fact that the traffic on the internal network of the cloud can be more trusted because physical

control of the network is maintained. It can be argued that the computers still need to be just as secure as if they were on the public Internet if they are sitting on a private network with firewalled access, but in reality, no CIO would feel comfortable letting a hacker just walk in and get on the internal database subnet of the network 24/7. One last detail about private clouds is that they can either be on premise or located at a trusted, audited 3rd party hosting provider location.

One more twist on a private cloud is now the existence of "virtual private clouds" (Amazon Web Services, 2015). These are run in a multi-tenant public cloud datacenter like a normal public cloud, but the network is logically separated from the Internet. Rather than attaching a VM directly to the Internet, it is put on a private IP network range of which the client organization has control over subnets, ACLs, and specifically configured chosen public Internet access. This type of cloud has the separation from the Internet available in a private cloud, but still does not maintain physical control over the hardware, which leaves the possibility of tampering open.

Community Cloud. Community clouds are an interesting concept that tries to combine the best aspects of public and private clouds into a single entity. Public clouds are an "everyone is welcome" party, where private clouds are just the one organization, but that organization does not have access to some of the benefits that come about when many companies share the same resource pool. Organizations with similar goals and security postures will band together to make a shared data center. This is especially common when several organizations have the same legal regulatory requirements. Several credit card processing companies could come together and make a PCI compliant datacenter, or a couple of healthcare organizations could build a HIPAA compliant

community cloud. In this model, it is financially possible to have larger non-utilized

pools to handle spikes in need but at the same time be able to trust the co-tenants.

Hybrid Cloud. While the community cloud concept combines aspects of public

and private clouds into one physical datacenter, the hybrid cloud takes a different

approach to achieving the benefits of both public and private clouds. In a hybrid cloud,

each server or application is analyzed and given a verdict whether its security

requirements dictate that it needs to go in a private cloud or if it is acceptable to put it in a

public cloud. In the end, some companies might be split with more in public and others

might keep more in a private cloud, but both can gain the advantages wherever possible.

Those apps and servers that were put in the public cloud can take advantage of the cost

savings, and those in the private cloud can have the benefit of being in a more secure,

known environment. The main concept of this thesis is how to effectively use a hybrid

cloud setup to achieve the best balance between security and cost savings offered by the

cloud.

<div align="center">

Cloud Security Concerns

</div>

Security issues in the cloud can be divided roughly into two categories: technical

issues like the quality of the hypervisor doing logical provisioning of resources on the

same physical host, and human issues such as problems caused by lack of due diligence

in overlooking the details of a system's design or operation. Both of these can lead to

avenues of attack with the same end result of causing huge problems in the information

security CIA triad.

Technical

The technical concerns for security mainly have to do with the ability to eavesdrop on various portions of the data flow in a typical business scenario. The three types of eavesdropping covered in this section are disk access (data at rest), processing access (data "in use"), and network access (data in motion). Any of these three can cause a severe breach of the CIA triad.

Disk and Data Access. One of the major differences encountered when using a cloud hosted at a third party site is that the servers and network are no longer under the client company's physical control. In a co-location situation, there might be thousands of virtual servers at a data center all sharing from the same or a few number of physical data stores. Because there are more virtual machines than data stores, logical provisioning is used to allocate virtual disks to the VMs. All of this data is able to be examined by some person or program with sufficient access. The steps necessary to gain access and manipulate hypervisor disk structure, including planting persistent backdoors, are described in ERNW's February 2013 newsletter (Kuhn, 2013). Pulling of a VM escape attack is hard due to the fact that hypervisors have a relatively small code base that can be more thoroughly tested than larger operating systems. Kuhn also points out that it is harder to gain access to some files such as the /etc/passwd file because the hypervisor keeps those files in an exclusive access locked state.

Besides just attacks from other tenants' virtual machines, there is also a possibility of a much easier attack through someone with credentials to access the hypervisor itself. Access to the hypervisor is at a level below the guest operating system and can easily read or even modify the entire virtual disk contents. Tampering with the data on the

servers could cause major problems with the client's business. However, with that there is

a possibility that from some other reference perspective there will be a realization that the

data has been modified and can then take action. On the other hand, just "spying" type

read-only access might never be detected.

The impact of this type of unauthorized access is huge! Confidentiality,

Availability, and Integrity can all be breached in the given situation. Confidentiality is the

first one to go with a more basic level of exploit: the revealing of information to the

unauthorized party and possibly disclosure to an even broader group. Availability can be

compromised if the data can be deleted or corrupted. Having backups can minimize the

downtime, but still will not contain all the information lost. Integrity is an issue when

data is modified to say something other than what it was supposed to say. This is a large

problem, especially if the target system is an important database.

Cloud backup and archival services are another portion of the cloud ecosystem

where data at rest can be attacked. Companies will have data generated either from their

machines in the cloud or even servers on site that will be fed into a cloud backup

solution. The businesses can then take advantage of the virtually unlimited storage at a

good price the cloud offers, and they also don't have to maintain an offsite location to

store the backups, but the main advantage is the redundancy. Amazon Glacier, one of

these data archival services, designs for "average annual durability of 99.999999999%

for an archive" (Amazon Web Services, 2015). This is more than even mid to large sized

companies can do with their own storage and location resources.

The next section will discuss mitigations to the problem of data leakage. The most

common solution is to use encryption, but even then it must be done the correct way, and

there are certain situations where it will not be feasible given another conflicting

requirement.

Data Processing. The attack against data while it is being processed is a similar

situation in terms of how it can be achieved and the results. It would take more effort to

glean a reasonable amount of information than just from viewing the results from disk,

which makes it seem like an attack that would not happen as frequently. However, the

counterbalance is that protecting data during processing is much more difficult than

protecting it while it is sitting out in storage. In general the data must be in a plaintext,

expanded form to be processed.

There has been work done on this problem of the necessity of data to be in

plaintext form in order to be processed. Homomorphic encryption is a type of encryption

that allows data to be processed while it is in an encrypted state where the processor does

not have the capability of decrypting it. At this point, it is mostly all theoretical work.

Much research and time will be needed before homomorphic encryption is a business

practical option. Craig Gentry's doctoral dissertation is on this subject, which was

presented at ACM's 2009 Symposium on Theory of Computing. (Gentry, 2009) One

application of this and similar research is the ability to search and process encrypted data

sets without giving anyone that was able to watch the processor perform the operations

the actual data that was being manipulated.

Network Sniffing. Like the previous two attacks in this section, the main

opportunity for this attack is in a datacenter that is either hosted with a third party

provider or where multiple companies are co-located in the same facility. Traffic going

between two machines might be intercepted and monitored for information. While private

and virtual private clouds have logical separation between the traffic of different clients,

flat public clouds can expose every machine straight onto the Internet. At that point

anyone along the network could sniff transmissions.

This problem, of the three, is by far the easiest to solve. Data can be received into

the server in an encrypted form, decrypted, processed, and encrypted again on the server,

and then sent back out. Even in a fully on-site corporate environment, best practice is to

encrypt traffic on the network whenever possible. TLS, SSH, and VPN are musts in

today's business world!

Human/Managerial Concerns

Another whole domain of risks exist when using hosted cloud computing. The

technical risks mentioned above are less and less possible from a purely technical exploit

basis now that hypervisors have become more refined and standard in most all business

areas. The results of those attacks can still happen, though, with misuse of authorized

access rather than gaining unauthorized access through exploits. This section will explore

some of those management concerns that can lead to a breach.

Unknown Risk Profile. In order for a company to decide on adopting a particular

strategy, they must first weigh the pros and cons in order to see if the move is

worthwhile. If it fulfills the requirements set in place and would be a positive change,

then the company has to do a risk assessment to see what risks are present. There is

always some sort of risk in a plan, but it might be more or might be less than the previous

methodology's risk level. Some common terms related to risk are defined below (Threat

Analysis Group, LLC, 2010):

-          Asset: what an entity is trying to protect

-          Threat: what an entity is trying to protect against

-          Vulnerability: a weakness or gap in protection efforts

-          Risk: the intersection of assets, threats, and vulnerabilities, A+T+V = R

In cloud computing, assets and threats are relatively clear. The problem is that there are unknowns about the vulnerabilities in the mix. These vulnerabilities can be categorized and information found out about necessary conditions, but data is missing that tells the story of how and how often each one happens in the real world. In quantitative risk analysis, the cost of an event multiplied by the probability of the event occurring gives you an expected loss average. This method can already be hard to use since one has to put a dollar amount on a not directly cash related item (i.e. loss of customer data), but it is even further complicated when the probability of an event happening is not known accurately. Because the cloud is a relatively new occurrence, the historical data on threats like these is scarce. This is a challenge IT policymakers have to deal with when deciding whether to move into the cloud.

Malicious Insiders. Rather than an attack themselves, malicious insiders are an attack vector. The people that manage the datacenter have a privilege level that is necessary to do their job, but could possibly give them access to client data they should not have access to. Their access can be used to perpetrate the attacks listed in the technical section. For example, network administrators regular sniff traffic along the networks to troubleshoot connection issues and re-route or re-design the network for higher performance. Rather than just looking at the traffic related to the problem at hand,

the administrator could also see other traffic going by and disclose that data if he so

desired.

Defenses against malicious insiders are generally a combination of the separation

of duties and least privilege concepts. These mean that each administrator has a

designated role and can only perform duties within his vertical "silo" (separation of

duties) and that privileges within his "silo" of responsibility are only given to the height

that is needed to perform job duties (least privilege). An example of this would be a

network admin given privilege to sniff traffic in order to troubleshoot, but not allowed to

modify traffic in transit.

Insufficient Due Diligence. At the end of the day, it is still humans that are in

charge of datacenter operations. They set the policies, make the base configurations, and

set the on-demand changes. This is a very large responsibility, given that an error in

configuration could impact the security of hundreds of client companies. Audits are very

important as they can point out gaps in the high level policies or specific low level

implementation. Responding to the audit findings and closing holes is a very important

process in the ongoing improvement of the datacenter.

The Cloud Security Alliance includes insufficient due diligence as one of the

factors in their "Notorious Nine Top Cloud Computing Threats of 2013." It was ranked

8th out of nine, but it was the only one where the graph of actual risk vs perceived risk

had the actual risk over tripling perceived risk. This shows that organizations don't

consider this threat as closely as they need to. Moving to the cloud requires adding

another group of people to the trust circle that has control over the information systems.

Thus, it is important that the cloud service provider be worthy of that trust.

Here is an example of how it is the human factor that fails in many attack scenarios. Target suffered a breach in November of 2013 in which many customers' credit card information was stolen. Fireeye, an intrusion detection system gave an alert that a threat "malware.binary" was found traversing the network. The security team didn't perform incident response actions on the alert because it was a generic low information one. John Strand of Black Hills Information Security said, "Target is a huge organization. They probably get hundreds of these alerts a day" (Finkle & Heavey, 2014). Even though this incident wasn't fully the security team's fault, it shows how in the end, human decisions are what make or break situations.

Restrictions on Penetration Testing. The last topic that will be discussed in this section are the additional hardships that are attached to testing information security in a public cloud. Penetration testing is an important step in the development lifecycle before deploying to production. A successful penetration test does not ensure an application is completely secure, but finding security vulnerabilities during the test does necessitate sending the application back and fixing it before it is deployed into active customer use. A challenge with cloud providers and penetration testing is that sometimes the line is blurry at what point the cloud service provider's and the leasing customer's responsibilities switch over (Jones, 2013). This is especially true with PAAS and SAAS because just as much or more of the full hardware and software stack is run by the hosting provider than the customer.

Amazon and other cloud providers require advance notice and approval before allowing penetration tests to occur on VMs they host (Amazon Web Services, 2015). These hosting providers want to know that attack traffic floating around their networks is

authorized, and they want to ensure any attack traffic does not interfere with other

companies' machines. Also, AWS makes sure other customers are not forced to endure

lower performance from one organization's penetration test. Amazon says that their

policy "does not permit testing m1.small or t1.micro instance types. This is to prevent

potential adverse performance impacts on the resources you may be sharing with other

customers in a multi-tenant environment." Although understandable requirements, these

might limit what a company is able to do penetration testing wise with their servers.

Security Approach

This thesis has presented the some information on cloud types and security

challenges faced by organizations moving IT services to the cloud. This concluding

section deals with methodologies, mitigations, and management to work toward the best

security possible in the cloud.

Cloud Type Choice

The first and probably most important aspect of the security plan in cloud

computing is the choice of cloud type. There exists a large spectrum of options with

cloud type, and all aspects of security, usability, and price must be taken into account

when making a decision. The choice ranges from an on-site, private cloud where the IT

staff of the owning company run everything, the most secure if the company knows what

they are doing, to a public cloud where the security is a combination of inherently worse

and unknown, but the cost is very low and management is offloaded. A risk management

study needs to be done to see which service model and deployment model should be

used. Specific situations may differ, but Figure 4 shows the conclusions of this research

on the best security choice cloud types for specific tasks.

Type	Private Cloud	Community Cloud	Public Cloud
Security Summary	Most secure	Similar to private, slightly less secure	Most unknowns, highest possibility for insecurity
Possible Uses	• Database with PII, payment, or login information • Protecting encryption keys (code signing, etc.)	• Web apps, customer portals • When several companies in the same market sector come together	• Public-facing website • Non-sensitive user services (i.e. news feed aggregation and delivery)

Figure 4: Cloud Type & Use Summary

      The deployment model choice affects what types of unauthorized access attacks are possible threats, but service models are chosen depending on the expertise of the IT employees a specific company has. If an organization is more lacking in the server hardening area, it would be a better solution to deploy code for a website on a PAAS. Then the PAAS product would take care of configuring server security instead of the less than knowledgeable company having to do it themselves. On the other hand, another company that has a very deep understanding of server hardening and containerization might want to configure their own servers because they trust themselves more than they trust another company's configuration. The service model can be a security related decision if the organization either knows a great deal or very about securing the levels below the application in the software stack. If the company does not have a mature server security team, it would be better to let the service provider control that aspect, but if that's where they excel and need an additional level of security, doing it in-house is better. In most cases, expertise levels will be similar or an edge to the cloud company, so it is just a management decision of whether or not the company wants to outsource that specific duty.

When choosing a deployment model, it is not necessary to use only one model and be locked in to the benefits and risks of just that one. Generally the best strategy today is to employ multiple different types of clouds depending on the risk assessment of the component being served. According to Ziff David B2B, the majority of companies utilize more than one cloud type in their IT strategy: 29% are public cloud only, 7% are private cloud only, and 58% use a hybrid cloud approach (2014). For example, user data with credit card number should probably be stored in a high security private cloud or a third party PCI certified community cloud, but the non-portal portion of a public website could be easily hosted in a public cloud (IBM Global Technology Services, 2013).

An important aspect to remember in any network, and especially in a hybrid cloud, is to divide the network into different security zones. High security areas should be segmented off from lower security, less trusted ones. In the connection between public and private clouds, the public cloud should be treated at a level similar to a DMZ. The public cloud portion should have heavily firewalled access to only the resources necessary on the private cloud side, and those resources should be abstracted from the critical infrastructure, (i.e. application APIs with strict sanitization instead raw database connections.) Having good segmentation in place allows companies to safely use both public and private clouds in concert to each do what they do best.

Making the right decisions in the area of architecting public and private cloud portions of infrastructure can result in an optimal balance of cost efficiency and security. The most cost efficiency can be gained with the public cloud, and full control over security can be retained by the private cloud. Matching the company's risk profile with

their desired risk appetite can determine where their computing needs would best be placed.

 Public Cloud Host-Based Security

Private clouds provide the most assurance of security for organizations that have the expertise necessary to manage it, but at some point, either due to cost benefits or lack of management expertise, a public cloud will be utilized. This section deals with protections that should be maintained in a public cloud environment, addressing the security issues presented in the section entitled "Cloud Computing Concerns."

Host-Based Firewalls. In a traditional data center or on-site private cloud, a reasonable amount of confidence could be assumed about the physical and logical protection of the network. In a public cloud, this might not be the case. Consequently, more security has to be baked into the hosts than might be needed in a traditional network.

Host-based firewalls are always a good idea in addition to network-based firewalls in an internal corporate network, but in a public cloud setting, they are imperative. Networks that just rely on network-based firewalls are like an egg with a crunchy shell and soft inside, but in a public cloud sometimes that outer shell isn't a possibility or is more limited in functionality. Host-based firewalls like Windows Firewall and iptables or firewalld on Linux can protect a single computer from all other computers on the network, whether those computers are friendly or malicious.

Secure Service Configuration, Least Privilege. Trying to only let allowed traffic get to the host is a good first step, but hackers will eventually be able to squeeze malicious traffic through the allowed rules. This is where other internal host hardening

techniques come in to play. The first of these is locking down each service configuration

so that it can perform all its duties without enlarging the attack surface any more than

necessary. One important aspect of this is permissions. These can be in the form of

database user permissions, filesystem permissions, or a newer variant: containerization.

Containers will be used more and more in the future as a way of deployment management

as doing a good job of security partitioning among different processes and services

(Docker, Inc., 2014-2015).

OS Hardening and Integrity Monitoring. Now that the network entry points

and services have been secured, one last aspect is the operating system itself. There are

multiple ways to secure the operating system. One of these is taking steps to harden the

kernel with anti-exploitation techniques. This is different depending on the operating

system. Linux will have options like Address Space Layout Randomization (ASLR) and

SELinux. These can be changed from their default state by using kernel flags. Microsoft

has a free product called EMET, which stands for Enhanced Mitigation Experience

Toolkit. It is a "shim" module that attempts to watch for and stop exploitation attempts

before they reach the kernel (Microsoft, 2015).

Another method for detecting break-ins and backdoors is deploying file integrity

monitoring software such as Tripwire (2015). File integrity monitoring ensures that the

configuration file is in the desired state in which it was left rather than being changed by

an attacker to allow malicious activity. A database of the mandated correct configuration

files and hashes is stored, and servers in production are compared on interval to the

master spec. If an inspection finds servers that are misconfigured, an incident can be

opened to see if it was a deployment mistake or an intrusion attempt, and the root cause

can be dealt with.

Crypto Protections. Because one of the possible attack vectors in a public cloud

was virtual disk access, an organization will do what is possible to protect their data

through cryptography even if the disk can be read. Two major states in which data needs

to be protected are in motion and at rest. Both are targeted by various attack vectors in

cyber space. Cryptography can be used to protect the data in both situations, but it is

implemented a different way (Lippard, 2013).

Traffic cannot be assumed to be safe from sniffing or modification in transit

between different servers on a public cloud or between servers and clients. When

information leaves the host, it must be encrypted. This follows best practices and is not a

difficult task to perform. TLS 1.2 is a good protocol for securing traffic between hosts,

and IPsec is used for tunneling between private networks. SSL 3.0 is now publicly known

as insecure due to the POODLE vulnerability and should not be used (United States

Computer Emergency Rediness Team, 2014).

Bruce Schneier states that encryption was originally invented to protect data in

motion – a message being transmitted between two parties (2010). For data in motion,

session keys can be created on the fly and transmitted through a key exchange protocol

for usage in just that conversation. Data in rest has to use the same key over a period of

time, which means the key must be stored somewhere where it is accessible to decrypt

the data at the time it is needed. The key must be present on or near the system that is

actually using the data, but it must be protected since it is what is being relied on to keep

the data confidential. A common, cost-effective solution for secure data storage and

archival is to encrypt the data while on the private cloud, and then transfer it out to the

public cloud for actual storage. Public clouds are generally cheaper and more

geographically redundant than a small to mid-sized company's private cloud, so the

backups are actually less likely to be lost out on a public cloud, but following this model

they are unreadable to even someone with direct disk access.

Third Party Hosting Company

Now the focus will move from securing each individual host to the last

recommendation for a cloud security approach: picking the right third party hosting

company. Having a trustworthy organization as your public cloud provider is a must. A

good company in this role will do their best to ensure the possible threats discussed in the

"Cloud Security Concerns" section do no happen. The two things to look for are a good

technical configuration that can ensure your virtual machines are logically separated from

other clients' ones, and a good ranking on a comprehensive audit.

Technical Configuration. This point of critique concerns the specific technical

configuration with which the virtual machines will be deployed, however in most all

cases, a hosting provider will not give a client their underlying infrastructure's source

code and configuration settings. Therefore it is somewhat hard to judge in how secure of

a manner the datacenter is configured. Probably the best method to get an estimation of

the technical skill of the company is to look at the number of servers it hosts on a day-to-

day basis, its employment size, and the level of qualification that is required for

employment at that organization.

Having a highly competent staff is a good sign, but the end goal is strong AAA

(authentication, authorization, and accounting) and following the least privilege model

for all cloud access and management. Services should be designed in a way that allows only authorized people to have certain access under the least privilege model. This applies to both the hosting company setting up their hypervisors management as well as the company putting their IT in the cloud. Trust is a necessity in computing: using any hardware component or piece of software requires a level of trust in the people that made it. However more reliance should not be put on trust than has to be. Systems can be made more secure by using designs that do not require as many trust assumptions.

An example of this in the real world is Lavabit, the encrypted email provider used by Edward Snowden. Moxie Marlinspike, a secure protocol researcher, published this analysis of Lavabit's method of operation. (Marlinspike, 2013) Even though it was a semi-widely acclaimed encrypted email service, there were some very fundamental trust assumptions that did not need to exist. These should actually be considered flaws in design. The most glaring one was that the private keys for each user were stored on the same set of servers as the actual email! Every time the user logged in with his passphrase, the server would take the passphrase it was given in plain text, decrypt the private key, and then use that to decrypt the email to show the user. All that needed to happen for the email service to be compromised was to permanently store the user's passphrase the next time he logged in. This design had basically no customer protection. As long as the company respected the user's trust and did what they promised to do, everything would be fine, but at any moment they could turn on the customers since the design offered no mandated security.

Audit. Most large hosting companies should be doing routine, comprehensive audits of their datacenters. These audits range from examining physical security to

checking authentication for datacenter personnel to analyzing redundancy capabilities in

Internet and power. In hosting contract negotiations, getting the hosting company to show

the results of their recent audit history can shed light on if that company would be a

suitable provider for cloud needs. If it has been consistently missing the mark on

important protections and hasn't taken steps to improve its position, there is no guarantee

it will fix the problems once a contract is signed.

With more rigorous audits, a datacenter can be certified to meet a specific set of

guidelines necessary for legal compliance in a field. For example, the provider Online

Tech has audited and certified its datacenters to be HIPAA compliant. (Online Tech,

2012) This is just what would be needed to help a small doctor's practice make a choice

as to which datacenter they would want to host their office servers. Even though hosting

companies will not share their exact details of operation, an audit can provide insight to a

potential client on whether a specific provider is a quality choice.

Final Thoughts

Moving a company's IT to the cloud can bring about huge benefits in terms of

cost and agility. The problem is that public clouds with outsourced management and

physical location have an inherent lack of security controls from the client. These leave

the possibility of major vulnerabilities that could impact the confidentiality, integrity,

and/or availability of important company data. The best solution is to use a hybrid cloud

and split the sensitive and non-sensitive infrastructure between private and public cloud

respectively with a closely controlled and monitored connection between the two.

Finally, extra care should be taken on the public cloud side to mitigate some of the

weaknesses of the public cloud model.

References

Alali, F. A., & Yeh, C. (2012). Cloud computing: Overview and risk analysis. *Journal of Information Systems, 26*(2), 13-33.

Amazon Web Services. (2015). *Amazon Glacier Product Details.* Retrieved from http://aws.amazon.com/glacier/details

Amazon Web Services. (2015). *Amazon virtual private cloud.* Retrieved from http://aws.amazon.com/vpc/

Amazon Web Services. (2015). *Penetration testing policy.* Retrieved from http://aws.amazon.com/security/penetration-testing/

Badamas, M. A. (2012). Cyber security considerations when moving to public cloud computing. *Communications of the IIMA, 12*(3), 1-18.

Bhadauria, R. C. (2014). Security issues in cloud computing. *Acta Technica Corviniensis Bulletin of Engineering*, 159-177.

Cloud Security Alliance. (2013). *The notorious nine: Cloud computing top threats in 2013.* Retrieved from https://cloudsecurityalliance.org/download/the-notorious-nine-cloud-computing-top-threats-in-2013/

DigitalOcean Inc. (2015). *DigitalOcean vps features.* Retrieved from https://www.digitalocean.com/features/technology/

Docker, Inc. (2014-2015). *Docker security.* Retrieved from https://docs.docker.com/articles/security/

Finkle, J., & Heavey, S. (2014, March 13). Target says it declined to act on early alert of cyber breach. Retrieved from http://www.reuters.com/article/2014/03/13/us-target-breach-idUSBREA2C14F20140313

Garber, D., Malik, J., & Fazio, A. (2013). *Windows azure hybrid cloud.* Indianapolis:

    John Wiley & Sons, Inc.

Géczy, P., Izumi, N., & Hasida, K. (2013). Hybrid cloud management: Foundations and

    strategies. *Review of Business & Finance Studies*, 37-50.

Gentry, C. (2009, September). *A fully homomorphic encryption scheme.* Retrieved from

    https://crypto.stanford.edu/craig/craig-thesis.pdf

Google, Inc. (2015). *Google app engine pricing.* Retrieved from

    https://cloud.google.com/appengine/#pricing

IBM Global Technology Services. (2013). *Private cloud in a hybrid cloud era: The*

    critical choices driving business value and agility. Retrieved from

    http://public.dhe.ibm.com/common/ssi/ecm/rl/en/rlw03025usen/RLW03025USE

    N.PDF

Jones, G. (2013). Penetrating the cloud. *Network Security*, 5-7.

Kaufman, L. (2010). Can public-cloud security meet its unique challenges? *IEEE*

    Security & Privacy, 8(4), 55-57.

Kesavaraja, D. (2011). Novel cloud intelligent defensive system against cyber attacks

    using penetration testing. *International Journal of Advanced Research in*

    Computer Science, 88-94.

Kuhn, F. (2013, February 21). Exploiting virtual file formats for fun and profit. *ERNW*

    Newsletter 41.

Lim, I., Coolidge, E., & Hourani, P. (2013). *Securing cloud and mobility: A practitioner's*

    guide. Boca Raton, Florida: Auerbach Publications.

Lippard, J. (2013, November 8). *Encryption may be the solution, but make sure you understand the problem.* Retrieved from http://blog.earthlinkbusiness.com/ security/encryption-data-in-use-data-disposed/

Marlinspike, M. (2013, November 5). *A critique of Lavabit.* Retrieved from http://www.thoughtcrime.org/blog/lavabit-critique/

Microsoft. (2015). *Enhanced mitigation experience toolkit.* Retrieved from http://microsoft.com/emet

Mitchell, R., & Meggison, P. (2014). Strategies for integrating cloud computing concepts. *Journal of Applied Research for Business Instruction, 12*(2), 1-6.

NIST. (2011). *The NIST definition of cloud computing.* Gaithersburg, MD: National Institute of Standards and Technology Special Publication 800-145.

NIST. (2012). *Cloud computing synopsis and recommendations.* Gaithersburg, MD: National Institute of Standards and Technology Special Publication 800-146.

Nolle, T. (2012, April 2). *Evaluating the useful life of a server.* Retrieved from http://www.dell.com/learn/us/en/04/smb/evaluating-the-useful-life-of-a-server

Online Tech. (2012). *HIPAA compliant data centers.* Retrieved from http://www.onlinetech.com/images/stories/downloads/whitepapers/hipaa-compliant-data-centers.pdf

Pearce, M., Zeadally, S., & Hunt, R. (2013). Virtualization: Issues, security threats, and solutions. *ACM Computing Surveys, 45*(2).

Ren, K., Wang, C., & Wang, Q. (2012). Security challenges for the public cloud. *IEEE Internet Computing, 16*(1), 69-73.

Schneier, B. (2010, June 30). *Data at rest vs. data in motion.* Retrieved from

      https://www.schneier.com/blog/archives/2010/06/data_at_rest_vs.html

Threat Analysis Group, LLC. (2010). *Threat, vulnerability, risk – Commonly mixed up*

      terms. Retrieved from http://www.threatanalysis.com/blog/?p=43

Tripwire, Inc. (2015). *Advanced cyber threat detection. increased security. Tripwire.*

      Retrieved from http://www.tripwire.com/

United States Computer Emergency Rediness Team. (2014, October 17). *SSL 3.0 protocol*

      vulnerability and POODLE attack. Retrieved from https://www.us-

      cert.gov/ncas/alerts/TA14-290A

Witt, C. (2011). HIPAA versus the cloud. *Journal of Health Care Compliance, 13*, 57-58.

Ziff David B2B. (2014). *It's not public versus private clouds: It's the right infrastructure*

      at the right time. Retrieved from

      http://public.dhe.ibm.com/common/ssi/ecm/en/xbl03047usen/XBL03047USEN.P

      DF