Concordia University College of Alberta

Master of Information Systems Security Management (MISSM) Program

7128 Ada Boulevard, Edmonton, AB

Canada T5B 4E4

A comparative case study on Cloud Service Providers, their Service Level Agreements, and loss of availability due to security breach: Amazon EC2 and S3, Microsoft Windows Azure Compute and Storage.

By

Ghebreslassie, Fitsum

A research paper submitted in partial fulfillment of the requirements for the degree of

Master of Information Systems Security Management

Research advisor:

Dale Lindskog, Assistant Professor, MISSM

A comparative case study on Cloud Service Providers, their Service Level Agreements, and loss of availability due to security breach: Amazon EC2 and S3, and Microsoft Windows Azure Compute and Storage.

By

Ghebreslassie, Fitsum

Research advisor:

Dale Lindskog, Assistant Professor, MISSM

Reviews Committee:

Dale Lindskog, Assistant Professor, MISSM

Ron Ruhl, Director and Associate Professor, MISSM

Concordia University College of Alberta

Master of Information Systems Security Management (MISSM) Program

7128 Ada Boulevard, Edmonton, AB

Canada T5B 4E4

A comparative case study on Cloud Service Providers, their Service Level Agreements, and loss of availability due to security breach: Amazon EC2 and S3, and Microsoft Windows Azure Compute and Storage.

By

Fitsum Ghebreslassie

Fghebres@student.concordia.ab.ca

A research paper submitted in partial fulfillment of the requirements for the degree of

Master of Information Systems Security Management

26 July 2013

Research advisor:

Dale Lindskog

Assistant Professor, Information Systems Security Management

Concordia University College of Alberta

Table of Contents

Abstract

Even though Amazon and Microsoft are promising 99.9 to 100 percent service uptime for their cloud compute and storage services, they don't take responsibility of service downtime that could result from Distributed Denial of Service (DDoS) and outside malicious attacks. The objective of this paper is to explore the actual security controls of these Cloud Service Providers (CSP) in order to find out if they can minimize service downtime from those attacks and make their promised service uptime percentage more appealing to customers. For this purpose those CSPs' cloud Service Level Agreement (SLA) and security practices are explored in relation to availability.

Keywords- CSP, SLA, Cloud Availability, Cloud Computing, Cloud Storage, Security Control, DDoS, Outside Malicious Attacks

1. Introduction

SLA is a contract between a network service provider and a customer that specifies, usually in measurable terms, what services the network service provider will furnish. [1] The focus of cloud SLA is service availability. Amazon and Microsoft are promising their customers at least 99.9% service uptime, for their compute and storage services. This service uptime excludes service interruptions from DDoS and outside malicious attacks. [2,3,4] The presence of high level of threat to service availability from those attacks makes this promised service uptime percentage somehow not very appealing information to cloud customers. [5] This paper will explore the security practices of Amazon and Microsoft, their cloud storage and compute services, in order to

find out if they can mitigate service downtime from those attacks and make the offered service uptime percentage a more appealing information to the customer.

This paper is organized in four parts. In the first part cloud SLA exclusions and service availability determination are explored. The second part explains compliance and best security practices. The third part explores the security controls of Amazon and Microsoft, that can mitigate service downtime as a result of DDoS and outside malicious attacks to their cloud compute and storage services, and evaluation against some best security practices. Finally the fourth part will include conclusion and future work which can be done.

2. SLA, SLA Exclusions, and Availability

This study is focused on non-negotiable cloud SLA. For a typical customer cloud SLA is non-negotiable and it includes CSPs' promised service availability, service availability determination, and exclusions to service availability. [6] The following table summarizes cloud SLA of Amazon and Microsoft: cloud compute and storage services.

	Amazon		Microsoft	
Cloud product and service	Elastic Compute (EC2)	Simple Storage Service (S3)	Windows Azure Compute	Windows Azure Storage
Availability computation	Annually	Monthly	Monthly	Monthly
Guaranteed availability	At least 99.95%	At least 99.9%	At least 99.95%	At least 99.9%
Availability period interval consideration	5 minutes	5 minutes	none	1 hour
Exclusions to service availability	Factors outside of Amazon's reasonable control(demarcation point)	Factors outside of Amazon's reasonable control(demarcation point)	Factors outside of Microsoft's reasonable control(demarcation point)	Factors outside of Microsoft's reasonable control(demarcation point)

Table 1 : Availability related components of cloud SLA of Amazon and Microsoft: cloud compute and storage services. [2,3,4]

Determination of EC2 Annual Uptime Percentage and S3 Monthly Uptime Percentage exclude downtime resulting directly or indirectly from any Amazon EC2 or S3 cloud SLA exclusions. Amazon's cloud SLA exclusions include among others; unavailability caused by factors outside of Amazon's reasonable control, including any force majeure event or Internet access or related problems beyond the demarcation point of Amazon EC2. [2,3,4]

The following equations show service uptime percentage determinations: [2,3,4]

- Amazon EC2 Annual Uptime Percentage = 100% - (% of 5 minute periods during the Service Year in which Amazon EC2 was in the state of "Region Unavailable")
- Amazon S3 Monthly Uptime Percentage = 100% - (the average of the Error Rates from each five minute period in the monthly billing cycle)

- Windows Azure Compute Monthly Uptime % = (Max. Connectivity Minutes - Connectivity Downtime) / Max. Connectivity Minutes
- Windows Azure Storage Monthly Uptime % = 100% - (the average of the Error Rates from each one hour period in the monthly billing cycle)

To determine service uptime percentage both CSPs use availability time interval, for instance Amazon uses 5 minute interval which means that service downtime for less than 5 minutes would not affect the CSP promised service uptime percentage. [7] Error rate is the total number of error status received from CSP's servers divided by the total number of cloud service storage requests of customers during the predetermined availability time interval. [2,3,4] For instance in the case of Amazon S3 customers should count and divide the number of 405 error messages received by the number of requests made to S3 within 5 minutes. [8] The question is what would be the effect to service availability if the received 405 error messages are as a result of DDoS and outside malicious attacks. Given the conditions stated in cloud SLA this question might not be relevant to the CSPs but is relevant to a typical customer who expects maximum cloud uptime (at least up to the CSPs promised service uptime). To answer this question the likelihood of those attacks can be determined by exploring the security practices of the CSPs, those practices which can mitigate service downtime from those attacks, which will be explained in the sections that will follow.

3. Compliance and Best Security Practices

Cloud security and security compliance are shared responsibilities between CSPs and customers. [9] Compliance assure cloud customers that CSPs are following best security practices. CSPs and

customers share the responsibility of compliance to industry and regulatory jurisdictions, this helps to minimize cloud service downtime due to DDoS and outside malicious attacks.

Both Amazon and Microsoft are offering their customers a number of security options, with the objective of assisting their customers to comply with industry standards and different regulatory jurisdictions. For instance they are providing customers the option of choosing among different regulatory jurisdictions (geographical availability zones) and other different security options like encryption, for their cloud computing and storage services. [10,11]

Microsoft is signatory to Safe Harbor and some portion of its operations is ISO/ IEC 27001: 2005 certified, yet Microsoft is working towards getting more security certifications. And Amazon designed and managed its web service infrastructure based on the following IT security standards: ISO 27001, ITAR, FIPS 140-2, DIACAP and FISMA, FedRAMP , SOC 1/ SSAE 16/ ISAE 3402 (formerly SAS 70), SOC 2, SOC 3, PCI DSS Level , and other initiatives including CSA and MPAA. [10,11]

Both Microsoft's Windows Azure Compute and Storage, and Amazon's EC2 and S3 are implementing their cloud security controls in such a way to comply with the above mentioned security standards. And some of those security controls which will be explained in the next section are expected to mitigate cloud service downtime from DDoS and outside malicious attacks.

4. Cloud Security Controls, Evaluation Against Some Best Security Practices, and Availability

Even though Amazon and Microsoft are not taking responsibility over service downtime as a result of DDoS and outside malicious attacks, they do have some security controls in place. Both the CSPs and their customers are expected to participate in cloud security as cloud security is a shared responsibility. The focus of this paper is security practices of CSPs, therefore customers are advised to know and implement their part of cloud security as it complements the availability of cloud service. The CSPs are participating through their compliance to security standards and their assistance to customers in their endeavor to comply with security standards. For instance the CSPs are giving their customers the option to store their data in multiple geographic availability zones, [10,11] and it is up to the customer to take advantage of it in order to improve the likelihood of data availability in the event of geographically available zone failure.

These two CSPs are actually implementing good security practices, as discovered from the exploration and evaluation of their security controls which will be explained, which improve their service availability even if they are not taking responsibility of service failure from DDoS and outside malicious attacks. What this actually mean to customers is those CSPs are implementing good security practices but there is no guaranty on service availability because of the lack of responsibility.

In the following sub-sections DDoS and outside malicious attacks mitigating security controls of the two CSPs are explained and evaluated against some best security practices.

4.1 Outside Malicious Attack Mitigation

Amazon and Microsoft are implementing the following security controls: access control, encryption, isolation, integrity, and penetration testing which can protect their cloud compute and storage services from outside malicious attacks. Customers should know that the difference in the two CSPs' security design, the need for more cloud security testing and security standardization because cloud is not yet mature, and differences in customer specific security goals often makes it hard to evaluate the security controls of the two CSPs and choose one CSP over the other.

4.1.1 Access Control

Access to Windows Azure storage is restricted and originally only least privileged access is granted to legitimate customers. Customers need to have key (storage account key) associated with their storage account in order to access the storage service they are subscribed to, and access is authenticated through either Windows Azure Portal Live ID or Service Management Application Programming Interface (SMAPI). [10]

In a similar manner to windows Azure storage access, access to Amazon S3 is restricted by default and originally only the subscriber to S3 has the right to access. The option of including metadata and setting access permissions to uploaded files gives customers the flexibility to grant access to user or group of users. Identity and Access Management (IAM) policies, bucket policies, Access Control Lists (ACLs) and query string authentication can be used to control access to S3. And the option of access logging serve the purpose of auditing.[11]

Accessing Windows Azure Compute takes a series of security steps. First customers need to connect to Windows Azure Portal through Live ID authentication, then the Fabric Controller (FC) takes over the task of controlling and managing the customer's Virtual Machine (VM) through bi-directionally authenticated Secure Socket Layer (SSL) connection. The FC carry out its task through deployed agents; Fabric Agent (FA) manages the root operating system of the physical machine hosting the customer VMs through instructions received from the FC, and Guest Agent (GA) manages the guest operating system of the customer VM through instructions received from FA. And for the purpose of FC security FA can't initiate connection to FC. [10]

Customers are required to provide credentials (X.509 Certificates or Amazon EC2 Key for authentication) associated with their accounts before accessing Amazon EC2 service. These credentials are based on public key cryptography, where the public key is placed by Amazon on the customer EC2 instance and the private key is to be stored and kept confidential by the customer. Amazon EC2 also provides customers with the option of creating users and groups, and managing their instances' access privilege and authentication. [11]

Some of the best access control security practices that customers can use to evaluate CSPs are: least privilege access principle, auditing and logging, and use of multi-factor authentication. [12] Both Amazon and Microsoft are following the principle of least privilege, by default they don't grant customers administrative access to their VMs. [13] They are also implementing auditing and logging. [14] Amazon has been providing the option of multi-factor authentication to its customers and recently Microsoft started to do the same. [15] Customers have to know also Microsoft's Windows Azure FC can be a security advantage or disadvantage, Because of this

Microsoft is determined to secure it, at this time it is too early to analyze FC as it yet needs to be tested widely.

4.1.2 Encryption

Amazon provides customers two options of data protection. One option is to encrypt data stored-at-rest in S3 through Server-Side Encryption (SSE), which uses 256 bit Advanced Encryption Standard (AES). The key used for SSE is managed and protected by Amazon itself. And the other option is Amazon S3 client-side encryption which allows the customer to encrypt data before uploading to S3, and the customer is to manage the encryption and encryption keys. [11]

Just like Amazon's options of S3 data encryption, Microsoft's Windows Azure storage provides the option of encrypting data stored and data in transit. All communications between FC and customers' compute and storage nodes is encrypted using SSL. This communication uses Microsoft CA issued certificates. Certificates and private keys in transit are protected through SSL encryption, and stored using password encryption in a secret location on FC. For more security, all temporary copies of certificates and keys are destroyed once the intended command is executed on the customer VMs. And to prevent FCs exposure escalation, FCs which don't support SSL encryption are placed in a separate Virtual Local Area Networks (VLANs). [10]

Amazon EC2 supports the option of connecting using Secure Shell (SSH) network protocol and Remote Desktop Protocol (RDP). In order to encrypt traffic to and from EC2, customers can enable port 22 for SSH and port 3389 for DRP (which are disabled by default) connections on their EC2 instances. [16]

All the previously mentioned access control authentications to Amazon S3, Amazon EC2, Windows Azure storage, and Windows Azure compute are encrypted.

Some of the best encryption security practices are: to encrypt data in motion and at rest, and to encrypt and store the encryption keys separately from the VM and the data. [12,17] Both Amazon and Microsoft are encrypting or assisting customers to encrypt data in motion and at rest. Microsoft encrypt and store encryption keys in a secret location on FC, customers should know that the importance of the FC can make it a potential target to outside malicious attacks but Microsoft seems to be determined to secure it. Amazon encrypt the encryption keys of stored data with regularly rotated master key, encryption keys are stored in hosts that are separate and distinct from those used to store data. [18] This evaluation shows customers that both CSPs are implementing the pre-stated best encryption security practices in their own cloud specific security designs.

4.1.3 Isolation

Amazon isolates EC2 virtual instances from each other and from their physical host. The only open source type-1 hypervisor (Xen hypervisor) is used to isolate VMs' CPUs, virtual memory and virtual hard disk. To further enforce the isolation all packets must pass through firewall residing within the hypervisor layer, which is between the physical network interface and the instance's virtual interface, thus an instance’s neighbors have no more access to that instance than any other host on the Internet and can be treated as if they are on separate physical hosts. [11]

Amazon also isolates the encryption key and master key (the key used to encrypt the encryption key) of the server side encrypted data-at-rest in S3 by placing them on different hosts. [18]

In a similar manner to Amazon's EC2, Microsoft's Windows Azure uses hypervisor that is specifically designed for use in the cloud (Windows Azure Hypervisor) together with root operating system to isolate customers' compute and storage nodes in their resource utilization and from communicating. This way root VM is isolated from guest VMs and guest VMs from each other. [10]

To secure the FCs Microsoft places the FCs that control and manage customer storage and compute nodes in separate host than the nodes. The FCs that manage compute nodes are different from the FCs that mange storage nodes. FCs use deployed agents to manage nodes in order to avoid direct communication. FCs manage root operating systems through fabric agents (FAs) and the FAs manage the guest operating systems that run on the customer compute and storage nodes through guest agents (GAs). [10]

Microsoft Windows Azure also uses VLANs to isolate trusted and un-trusted (for example the FCs that don't support encryption in their communication) components of the network. [10]

One of the best isolation security practice is CSPs making sure customers' virtual machines are isolated from each other. This prevents outside malicious attackers from easily propagating their attack from a compromised VM to other VMs. As explained Both Amazon and Microsoft Windows Azure are following good isolation security practices. But customers should know that selecting one out of the two CSPs by evaluating their isolation security practices is still hard because of lack of detailed information on their security configurations.

4.1.4 Integrity

Amazon S3 regularly verifies the integrity of stored data using checksums. Combination of MD-5 checksums and Cyclic Redundancy Checks (CRC) are used to detect stored data corruption, if data corruption is detected redundant data is used to repair it. Amazon also use checksums on traffic to and from S3 to detect corruption of data in transit. And cryptographic hashes SHA-1 and SHA-2 are used as an assurance to the integrity of its S3 and EC2 authentication system. [11]

Microsoft achieves Windows Azure storage and Fabric integrity through its design of access restriction, as already explained under the section access control. [10]

When we see Windows Azure compute, Microsoft incorporates integrity in its design of the three local Virtual Hard Drives (VHD) attached to each VM. The D: drive contains guest operating system, E; drive contains FC constructed image based on customer provided package, and C: drive contains configuration information, paging files, and other storage. Virtual drives D: and E: are read-only drives, and drive C: is read and write drive. The D: drive's guest operating system update and the E: drive's update with a new application image are implemented as VHD with delta files. The delta drives for each read only drive are discarded every time the VHD is updated. Each instance has its configuration file in its read and write VHD C: drive, in order to update this configuration file a customer must be authenticated through the Windows Azure Portal or SMAPI. The configuration file update is done after passing through a series of hardened operating systems, the FC then FA (which is in a hardened root operating system) and at last passes through GA (which is in a hardened guest operating system) to reach the instance's virtual C: drive. [10]

One of the best integrity security practices is the use of checksums and replication. [7] Amazon is following this practice to check data integrity and restore corrupted data, replication is explained in more detail under the section redundancy, which is more focused on reactive approach to integrity. But Microsoft is following a different path to reach the same end, a design focused proactive approach to integrity. The choice depends on the client, which approach of integrity suits best his/her specific availability security needs.

4.1.5 Penetration Testing

Penetration testing improves cloud availability through its identification of vulnerabilities, CSPs use the identified vulnerabilities information to harden the cloud system from outside malicious attacks. Both Amazon and Microsoft are securing their cloud services through performing penetration testing on regular bases, and they provide customers with the option to do their own penetration testing of the cloud resources they are subscribed to. Customers are required to fill out penetration testing request form (requires the customer to read and agree to terms and conditions specific to penetration testing and to the use of appropriate tools for testing) and wait for authorization from the CSPs before performing the test. [19,20]

What clients should know before performing penetration testing is the delineation point (the scope of testing) between service provider and tenant. For instance the tenant is responsible for the virtual machine (and everything that runs within it) in an Infrastructure as a Service (IaaS) cloud service model, for application and API or Graphical User Interface (GUI) in Platform as a Service (PaaS) cloud service model, and partially for API/GUI in Software as a Service (SaaS) cloud service model. Therefore clients should harden their domains themselves from outside

malicious attacks after performing application or other domain appropriate penetration testing. [21]

Cloud customers can evaluate the two CSPs' penetration testing with the objective of minimize service downtime, as shown in the following table.

Evaluation criteria	**Amazon**	**Microsoft**
Does the CSP perform internal penetration testing?	Yes	Yes
Does the CSP allow third party penetration testers to perform internal penetration testing?	Yes	Yes
Does the CSP publish results of its penetration testing?	No	No
Does the CSP allow clients to perform penetration testing?	Yes	Yes
Does the CSP allow clients to delegate third party penetration testers?	Yes	Yes

Table 2:Cloud Pen Testing Evaluation of Amazon and Microsoft [19,20,22,23]

As shown in the table those CSPs are following good security practice by participating in penetration testing, and engaging customers and third party penetration testers. For security reasons both Amazon and Microsoft do not publish the results of their penetration testing, but those results can help customers in their evaluation and choice of those CSPs.

4.2 DDoS Attack Mitigation

Amazon and Microsoft are implementing the following security controls: packet filtering, redundancy, and load balancing which can protect their cloud compute and storage services from DDoS attack. Customers should know that the difference in the two CSPs' security design, the need for more cloud security testing and security standardization because cloud is not yet mature,

and differences in customer specific security goals often makes it hard to evaluate the security controls of the two CSPs and choose one CSP over the other.

4.2.1 Packet Filtering

Amazon placed firewalls and other boundary devices at external and within key internal boundaries of its network, and updates those devices regularly with Amazon's Information Security approved ACL policies. A default firewall policy of deny all inbound traffic is applied for all EC2 instances. And customers are given the option to manage incoming traffic to their EC2 instances based on protocol, port number and source IP address (individual or Classless Inter-Domain Routing [CIDR] block). [11]

Amazon has strategically placed a limited number of access points to the cloud to allow for a more comprehensive monitoring of inbound and outbound communications and network traffic. These customer access points (API endpoints) allow customers to establish a secure Hypertext Transfer Protocol (HTTPS) communication session with their S3 or EC2 instances within Amazon Web Services (AWS). [11]

When we see Microsoft's Windows Azure, packet filtering is performed at both the Windows Azure network's edge load balancers and at the root operating systems. This prevents traffic not addressed to VMs, traffic to protected infrastructure endpoints, and broadcasts. And VLANs are segmenting networks and forcing traffic to pass through appropriate routers. This VLANs include untrusted VLANs of customer nodes, and trusted VLANs of FCs, supporting systems, and other infrastructure devices.[10]

Some of the best packet filtering security practices that customers can use to evaluate CSPs include: network and host level packet filtering. Both Amazon and Microsoft are implementing network packet filtering and assist their customers to manage their VMs' packet filtering. [10,11]

4.2.2 Redundancy

Redundancy isolates the failure of one availability zone or geographic availability region from the other at times of DDoS attack, it ensures service continuity. Amazon offers the option of placing EC2 and S3 instances within multiple geographic regions (East US [North Virginia], West US [North Carolina], West US [Oregon], EU [Ireland], Asia Pacific [Sydney], Asia Pacific [Singapore], Asia Pacific [Tokyo],and South America [Sao Paulo]) and across multiple availability zones within each geographic availability region. Each availability zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Amazon also ensures redundancy in telecommunication, arrangements are made for the availability of Internet service from multiple Internet Service Providers (ISPs) to each region. [11]

Customers should know that when they create Amazon S3 bucket or EC2 instance Amazon requires them to specify geographic availability region, then the services (customers are subscribed to) are made redundantly available across multiple availability zones within the customer specified region. But unlike S3 out of the ten EC2 instance families (families are created based on memory size, processing speed, storage size...etc.) some might not be available in specific geographic regions. [11,24]

And Microsoft's Windows Azure achieved redundancy in a similar manner to Amazon's cloud redundancy. Microsoft uses two approaches to replicate and geo-failover customers' data stored in its Windows Azure storage: Geo Redundant Storage (GRS) and Locally Redundant Storage (LRS). GRS is equivalent to Amazon's geographic availability regions and LRS to Amazon's multiple availability zones within each geographic availability region. GRS is turned on by default and replicates data between two locations hundreds of miles apart within the same region (North Central US [Chicago, IL], South Central US [San Antonio, TX], East US [Virginia], and West US [California], North Europe [Dublin, Ireland], West Europe [Amsterdam, Netherlands], East Asia [Hong Kong, China], and Southeast Asia [Singapore]). LRS replicates storage three times within the same data center if a customer turn off geo-replication. Both approaches are used to replicate Windows Azure storage blobs (service that store unstructured binary and text data) and tables (service that store non-rational structured data), but not queue (service that store messages that may be accessed by a client and providing reliable messaging between role instances) and customers' applications deployed to Windows Azure compute. [25,26,27]

Customers should know that both Amazon and Microsoft are following good security practices by make their cloud services redundantly available, but their regional and zone availabilities differ. Even though both CSPs don't have presence in some continents Amazon's datacenters are present in many continents compared to Microsoft's Windows Azure. And both CSPs have their differences in the number of availability regions they have at a continental level.

Customers evaluation of the two CSPs' offered redundancy can vary based on their specific needs. For instance if a customer wants to restrict his/her cloud assets within North America (because of his/her preferred regional jurisdiction) with the highest number of availability regions (to

maintain high level of availability at times of availability region directed DDoS attacks) the choice will be Microsoft's Windows Azure, but if the customer wants to achieve the same objective in South America then Amazon will be the choice.

4.2.3 Load Balancing

Load balancers minimize cloud service downtime at times of DDoS attack by evenly distributing traffic to multiple instances (within a single or multiple availability zones) and by stopping traffic from being forwarded to instances which are already overwhelmed by the attack. Load balancing complements redundancy through its functions of minimizing resource consumption and implementing fail over. [28]

Customers can use Amazon's elastic load balancing option to automatically balance traffic across multiple Amazon EC2 instances and multiple availability zones, and to ensure that only healthy (determined by configurable response time based health checks) instances receive traffic. Using auto scaling option with load balancing can also enable customers to maintain a predetermined minimum number of available instances. The load balancers support HTTPS and SSL listener protocols for secure communication. [29]

Microsoft's Windows Azure provides load balancing endpoints for the same purpose they are used by Amazon. The endpoints round-robin load balance traffic between two or more customer virtual machines. Just like Amazon, Microsoft uses periodic instance response time based health checks to prevent routing of traffic to unavailable virtual machines and the load balancers support HTTPS. [30]

One of the best load balancing security practice that customers can use to evaluate CSPs is making sure their load balancing mechanism is fault tolerant. Failure of load balancers can create failure of DDoS attack mitigation program. In the case of Amazon, Route 53 (which is a highly available and scalable Domain Name System [DNS]) will fail away from a load balancer if the load balancer itself is unhealthy or if there are no healthy EC2 instances registered with the load balancer. Using Route 53 DNS failover customers can run applications in multiple availability regions and designate alternate load balancers for failover across regions. In the event that customer application is unresponsive Route 53 will remove the unavailable load balancer endpoint from service and direct traffic to an alternate load balancer in another region. When we see the case of Microsoft, the option of Windows Azure traffic manager (WATM) with failover load balancing method can be used to control the distribution of user traffic to similar hosted services that are running within the same data center or in different data centers across the world. [31,32]

5. Conclusion and Future Work

Amazon and Microsoft are offering 99.9 to 100 percent service availability of their cloud compute and storage, but they are not taking responsibility of service downtime as a result of DDoS and outside malicious attacks. This paper discovered that those CSPs are implementing good security controls which can minimize downtime from those attacks. And those CSPs are assisting customers in doing their part of cloud security as cloud security is a shared responsibility.

This study encourages other studies to be conducted on the feasibility of those CSPs to take responsibility of cloud security from those attacks, this means that if it is feasible customers will get better guaranty on the promised 99.9 to 100 percent service availability.

References

[1] Tech Target. Definition: Service-Level Agreement. [Online]. Available: http://searchitchannel.techtarget.com/definition/service-level-agreement

[2] Microsoft Windows Azure. (2013, June). Service Level Agreements. [Online]. Available: http://www.windowsazure.com/en-us/support/legal/sla/

[3] Amazon Web Services. (2013, June 1). Amazon S3 Service Level Agreement. [Online]. Available :http://aws.amazon.com/s3-sla/

[4] Amazon Web Services. (2013, June 1). Amazon EC2 Service Level Agreement. [Online]. Available : http://aws.amazon.com/ec2-sla/

[5] Cloud Security Alliance. (2013, February). Cloud Security Alliance Warns Providers of ' The Notorious Nine 'Cloud Computing Top Threats in 2013. [Online]. Available: http://cloudsecurityalliance.org/csa-news/ca-warns-providers-of-the-notorious-nine-cloud-computing-top-threats-in-2013/

[6] Cloud Standards Customer Council. (2012, April 10). Practical Guide to Cloud Service Level Agreements Version 1.0 .[Online]. Available: http://www.cloudstandardscustomercouncil.org/2012_Practical_Guide_to_Cloud_SLAs.pdf

[7] National Institute of Standards and Technology. (2012, May). Cloud Computing Synopsis and Recommendations: SP 800-146. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-146/sp800-146.pdf

[8] Amazon Web Services. Amazon Simple Storage Service FAQs: How do I know if I lose an RRS object? [Online]. Available: http://aws.amazon.com/s3/faqs/

[9] Payment Card Industry Data Security Standard. (2013, February). PCI DSS Cloud Computing Guidelines Information Supplement V2.0. [Online]. Available: http://www.pcisecuritystandards.org/pdfs/PCI_DSS_v2_Cloud_Guidelines.pdf

[10] Windows Azure. (2013, March). Window Azure Trust Center-Security-Security Resources for Windows Azure-Windows Azure Security Overview. [Online]. Available: http://www.windowsazure.com/en-us/support/trust-center/security/

[11] Amazon Web Services. (2013, June). Overview of Security Processes. [Online]. Available: http://s3.amazonaws.com/awsmedia/pdf/AWS_Security_Whitepaper.pdf

[12] Symantec. (2011). The Secure Cloud: Best Practices for Cloud Adoption. [Online]. Available: http://www4.symantec.com/mktginfo/whitepaper/TheSecureCloudBestPracticesforCloudAdoption_cta52644.pdf

[13] Microsoft. 10 Things to Know about Azure Security. [Online]. Available: http://technet.microsoft.com/en-us/cloud/gg663906.aspx. Accessed on (2013, July 20).

[14] Microsoft. (2011, August 24). Store and View Diagnostic Data in Windows Azure Storage. [Online]. Available: http://msdn.microsoft.com/en-us/library/windowsazure/hh411534.aspx. Accessed on (2013, July 20).

[15] Microsoft. (2013, February 21). Using Multi-factor Authentication with Windows Azure AD. [Online]. Available: http://technet.microsoft.com/en-us/library/jj713614.aspx

[16] Amazon Web Services. (2013, June 15). Authorizing Network Access to your Instances. [Online]. Available: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/authorizing-access-to-an-instance.html

[17] Cloud Security Alliance. (2011, November 14). Security Guidance for Critical Areas of Focus in The Cloud Computing Version 3.0. [Online]. Available: http://cloudsecurityalliance.org/research/security-guidance/. Accessed on (2013,July 20).

[18] Amazon Web Services. (2011, October 4). New-Amazon S3 Server Side Encryption for Data at Rest. [Online]. Available: http://aws.typepad.com/aws/2011/10/new-amazon-s3-server-side-encryption.html. Accessed on (2013, July 20).

[19] Amazon Web Services. Penetration Testing. [Online]. Available: http://aws.amazon.com/security/penetration-testing/

[20] Windows Azure. (2013, March). Security-Penetration Testing. [Online]. Available: http://www.windowsazure.com/en-us/support/trust-center/security/. Accessed on (2013, July 20).

[21] Chris Brenton. (2012, July 05). Pen Testing in the Cloud. [Online]. Available: http://pen-testing.sans.org/blog/2012/07/05/pen-testing-in-the-cloud. Accessed on (2013, July 20).

[22] Amazon Web Services. AWS FedRAMP ATO: Difficult to Achieve, Easily Misunderstood, Valuable to All AWS Customers. [Online]. Available: http://aws.typepad.com/aws/security/

[23] Angus Kidman. (2013, April 12). Top 10 Lesser-known Facts About Windows Azure Security. [Online]. Available: http://www.lifehacker.com.au/2013/04/top-ten-lesser-known-facts-about-windows-azure-security/

[24] Amazon Web Services. (2013, June 16) . Instance Families and Types. [Online]. Available; http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-types.html

[25] Windows Azure. How To Manage Storage Accounts. [Online]. Available: http://www.windowsazure.com/en-us/manage/services/storage/how-to-manage-a-storage-account/

[26] Microsoft. (2011, September 15). Introducing Geo-replication for Windows Azure Storage. [Online]. Available: http://blogs.msdn.com/b/windowsazurestorage/archive/2011/09/15/introducing-geo-replication-for-windows-azure-storage.aspx. Accessed on (2013, July 20).

[27] Microsoft. (2011, October 06). Blobs, Queues, and Tables. [Online]. Available: http://msdn.microsoft.com/en-us/library/windowsazure/gg433040.aspx. Accessed on (2013, July 20).

[28] Wikipedia. Load balancing (Computing). [Online]. Available: http://en.wikipedia.org/wiki/Load_balancing_%28computing%29

[29] Amazon Web Services. Elastic Load Balancing. [Online]. Available: http://aws.amazon.com/elasticloadbalancing/

[30] Windows Azure. Load Balancing Virtual Machines. [Online]. Available: http://www.windowsazure.com/en-us/manage/windows/common-tasks/how-to-load-balance-virtual-machines/

[31] Amazon Web Services. Amazon Route 53. [Online]. Available: https://aws.amazon.com/route53/

[32] Windows Azure. Windows Azure Traffic Manager Overview. [Online]. Available: http://msdn.microsoft.com/en-us/library/windowsazure/hh744833.aspx#BKMK_RRLB