

Fall 2011

Continuous Monitoring in the Cloud Environment

Victoria Nyffeler
Regis University

Follow this and additional works at: <http://epublications.regis.edu/theses>



Part of the [Computer Sciences Commons](#)

Recommended Citation

Nyffeler, Victoria, "Continuous Monitoring in the Cloud Environment" (2011). *All Regis University Theses*. Paper 630.

This Thesis - Open Access is brought to you for free and open access by ePublications at Regis University. It has been accepted for inclusion in All Regis University Theses by an authorized administrator of ePublications at Regis University. For more information, please contact repository@regis.edu.

Regis University
College for Professional Studies Graduate Programs
Final Project/Thesis

Disclaimer

Use of the materials available in the Regis University Thesis Collection ("Collection") is limited and restricted to those users who agree to comply with the following terms of use. Regis University reserves the right to deny access to the Collection to any person who violates these terms of use or who seeks to or does alter, avoid or supersede the functional conditions, restrictions and limitations of the Collection.

The site may be used only for lawful purposes. The user is solely responsible for knowing and adhering to any and all applicable laws, rules, and regulations relating or pertaining to use of the Collection.

All content in this Collection is owned by and subject to the exclusive control of Regis University and the authors of the materials. It is available only for research purposes and may not be used in violation of copyright laws or for unlawful purposes. The materials may not be downloaded in whole or in part without permission of the copyright holder or as otherwise authorized in the "fair use" standards of the U.S. copyright laws and regulations.

CONTINUOUS MONITORING IN THE CLOUD ENVIRONMENT

A PROJECT

SUBMITTED ON 9 OF OCTOBER, 2011

TO THE DEPARTMENT OF INFORMATION TECHNOLOGY
OF THE SCHOOL OF COMPUTER & INFORMATION SCIENCES
OF REGIS UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS OF MASTER OF
SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

BY



Victoria Nyffeler

APPROVALS



Robert Bowles, Thesis Advisor



Shari Plantz-Masters



Ranked Faculty Name

Abstract

The National Institute of Standards and Technology introduced a risk management framework that concludes with a process for continuous monitoring. Continuous monitoring is a way to gain near real-time insight into the security health of an information technology environment. The cloud environment is unique from other environments in the way that resources are virtualized and shared among many cloud tenants. This type of computing has been gaining popularity as a solution for organizations to purchase resources as an on-demand service in the same way that an organization purchases utilities today. In order to experience the benefits promised by the emergence of cloud computing the inherent security challenges in utilizing shared resources must be addressed. The proposed continuous monitoring program, based on recommendations from the National Institute of Standards and Technology Draft Special Publication 800-137 (Dempsey et al., 2010), is intended to address these security concerns. The program specifically addresses continuous monitoring activities for cloud providers to implement related to configuration management, patch and vulnerability management, antivirus/malicious software management, firewall management, and access management. This proposal does not address the shared responsibilities between the cloud tenant and cloud provider which is recommended as the next step in this research. The tenant and provider should have complementary controls and continuous monitoring programs to ensure the security of a cloud solution.

Table of Contents

Abstract ii

Table of Contents iii

List of Tables iv

Chapter 1 – Introduction 5

Chapter 2 – Review of Literature and Research 2

 Cost Savings to the Organization..... 4

 Increased Productivity 4

 Simplified IT Management..... 5

 Refocus on Core Competencies 5

 The Security Risks 6

Chapter 3 – Methodology 8

Chapter 4 –Results 10

 Configuration Management 10

 Patch and Vulnerability Management..... 11

 Antivirus / Malicious Software..... 16

 Firewall Management 18

 Access Management 21

Chapter 5 – Conclusions 25

References..... 27

Appendix A..... 31

 Continuous Monitoring Grid..... 31

Glossary 34

List of Tables

Table 1: External Vulnerability Resources 13

Chapter 1 – Introduction

Cloud computing has been gaining momentum over the past decade as a viable option for enterprise applications. The National Institute of Standards and Technology (NIST) describes cloud computing as a solution for organizations to purchase computing resources on-demand that can be rapidly implemented with minimal management or service provider interaction (Mell & Grance, 2009). This definition actually covers many different forms of computing that many are utilizing today and may not be aware that the foundation is based on cloud models, such as Google Docs or Facebook.

The cloud model is expected to make computing resources broadly available to small, mid-sized, and large organizations in the same way that utilities are purchased today. Smith (2009) suggests that cloud computing will evolve in the same way that electricity generation has migrated from an internal model a century ago to the service that exists today. The technology community has already witnessed a similar evolution of services in the telecommunications industry in the 1990's (Smith, 2009). Historically, capacity was hard wired between destinations; but, this design evolved into capacity being managed through Virtual Private Networks (VPNs) that provided a secure path between destinations that was built using many segments (Smith, 2009). This design coined the term “telecom cloud,” which was the first instance of the modern day term for the cloud model (Smith, 2009). Now, we are experiencing the next evolution in technology which utilizes a cloud environment for hosting services that are available on an as needed basis.

Chapter 2 – Review of Literature and Research

The cloud model can be broken into three major offers: (1) Software as a Service, (2) Platform as a Service, and (3) Infrastructure as a Service (Mell & Grance, 2009). Software as a Service (SaaS) provides software applications that are housed in the cloud provider's environment that users can access via a thin client, most commonly a web browser (Mell & Grance, 2009). Platform as a Service (PaaS) provides the ability for consumers to house consumer-created or acquired applications and deploy in the cloud infrastructure (Mell & Grance, 2009). In this case the underlying systems and hardware are maintained by the cloud provider, but the consumer has control over the deployed applications and can also maintain the environment configurations (Mell & Grance, 2009). PaaS can be used for the entire software development lifecycle by hosting the development tools in the cloud infrastructure and utilizing the cloud environment to deploy to users. Infrastructure as a Service (IaaS) provides fundamental computing resources to the consumer (Mell & Grance, 2009). This service offers consumers processing, storage, and network resources that are hosted in the cloud provider's environment. Consumers can still maintain access to configure operating systems, deployed applications, and some networking components (Mell & Grance, 2009).

In addition to the three major cloud environments there are also five key characteristics that the NIST attribute to a service being considered cloud computing (Mell & Grance, 2009). These five critical elements define a service as part of the cloud model: (1) on-demand self-service, (2) broad network access, (3) resource pooling, (4) rapid elasticity, and (5) measured service. On-demand self-service suggests that resources can be automatically provisioned by the consumer without requiring manual intervention (Mell & Grance, 2009). Broad network access is required for consumers to access the platform using a multitude of devices (Mell & Grance,

2009). Resource pooling refers to the basic concept of cloud customers sharing resources that grow and shrink automatically, in some cases, based on demand, known as rapid elasticity (Mell & Grance, 2009). The last characteristic and arguable the most important element of this shared resource space is ensuring that the service being provided can be monitored, controlled, and reported (Mell & Grance, 2009). Measured services indicate that metrics must accompany the solution that is provided to customers; but metrics are also useful for the provider to adequately operate the cloud environment.

Lastly, the cloud model is described by four types of deployment methods: (1) private, (2) public, (3) community, and (4) hybrid. A private cloud can be hosted by a third party provider or managed internally and is intended for use by a single organization (Ryan & Loeffler, 2010). A public cloud deployment is hosted by a third party provider and can experience the greatest financial benefits since it is shared by many customers (Ryan & Loeffler, 2010; Mell & Grance, 2009). The community cloud can be shared by multiple organizations that are connected in some way; the NIST provides potential threads of connection as “mission, security requirements, policy, and compliance considerations” (Mell & Grance, 2009). Lastly, the hybrid model is a combination of two or more deployment methods described above (Mell & Grance, 2009).

Cloud computing is gaining momentum as a viable option for many types of enterprise applications. It provides a solution for organizations to purchase computing resources as an on-demand service in the same way that an organization purchases utilities today (Talbot, 2010; Anthes, 2010; Goodburn & Hill, 2010). Organizations have been driven to take a closer look at incorporating the cloud model into the overall business strategy due to the many substantial benefits that could be gleaned by making the move. These benefits have been described in many

different ways, but seem to fall into the following categories which will be described in more detail below.

- (1) Cost savings to the organization
- (2) Increased productivity
- (3) Simplified Information Technology (IT) management
- (4) Refocus on core competencies

Cost Savings to the Organization

First and foremost, an organization can see cost savings simply by not hosting these technology resources on-site. Smith (2009) indicates that by using a third party cloud provider an organization will avoid the added cost of providing an adequate environment for technology equipment. Traditionally, the facility hosting business critical resources would require increases to the electrical system, isolated floor space, modifications to regulate the air conditioning, as well as staff to manage the technology program (Smith, 2009). The core function of cloud computing, as the NIST's definition confirms (Mell & Grance, 2009), is that resources are elastic (Owens, 2010). This characteristic of the cloud provides a major savings for organizations as they only need to purchase what they need when they need it. Rather than purchasing resources to support peak times that only last a portion of the year, the organization only pays for increased computing resources during those peak times when needed. Smaller organizations that may just be getting off the ground can forecast expected computing needs, but if those forecasts are too high or too low this elasticity feature will allow their customers to remain unimpacted.

Increased Productivity

The cloud model allows organizations to respond to changes and be more agile when making decisions that support their business. This increased flexibility enables organizations to

respond to market changes faster and easier than traditional technology environments (Goodburn & Hill, 2010). The ability to implement change rapidly within an organization provides the opportunity to capitalize on market shifts immediately and as a result experience an increase in productivity (Goodburn & Hill, 2010).

Simplified IT Management

Two views can be explored relative to the simplification of IT management: first, from the perspective of a small to mid-sized business; and second, from the perspective of a large organization that does not specialize in technology management. The cloud model provides a small to mid-sized business with access to “the same technology infrastructure and support as a Fortune 500 company” creating a huge advantage today over conventional IT models for these companies (Goodburn & Hill, 2010). In many IT departments today technology associates are spread so thin that there are few opportunities to become an expert in a specific technology. Large organizations can also see the benefits from this simplified IT management structure by capitalizing on the knowledge and experience of the cloud community (Goodburn & Hill, 2010). Both small and large organizations will be able to take advantage of the growing competition between cloud providers and have the opportunity to move to a provider if needs are not being met (Smith, 2009).

Refocus on Core Competencies

Organizations should be focused on fulfilling their primary business function, not on becoming experts in managing IT resources. Goodburn and Hill (2010) emphasize that the cloud allows organizations to reconnect with their core competencies and redirect resources from managing internal technology to focusing attention on their primary business objectives. This

ability to “outsource” these technical responsibilities frees employees to attend to “other aspects of their work that could otherwise have been neglected” (Cloud Computing, 2009).

The Security Risks

In order to experience the benefits promised by the emergence of cloud computing the inherent security challenges in utilizing shared resources must be addressed. The cloud infrastructure is based on the virtualization of processors, networks, and disk drives which allow multiple users to run concurrently on a single physical server (Talbot, 2010; “Hypervisor,” 2011). However, it has been demonstrated that services running on a single piece of hardware has an increased potential for the system to be compromised which is explored further in the following two examples.

The first example is presented by Owens (2010) as a vulnerability that was identified in November 2009 that allows a user to traverse from one virtual machine client environment to another client environment managed by the same hypervisor. Owens (2010) further emphasizes this security vulnerability specifically in relation to elasticity; one of the basic functions of cloud computing that allows users to grow and shrink resources on-demand. This highlights this vulnerability given a public cloud model and the lack of user control over where data may be physically stored (Owens, 2010).

The second example presented by Talbot (2010) is regarding sharing hardware between multiple cloud customers. Researchers have demonstrated that an attacker can successfully steal data using an eavesdropping program when two programs are running in parallel on the same operating system (Talbot, 2010). Talbot (2010) then suggests that this same kind of attack could penetrate a cloud environment when virtual machines run on a single server. Anthes (2010) confirms Talbot’s claim by referencing a successful side-channel attack using virtual machines

located on the same hardware that was conducted by computer scientists at University of California and MIT.

These examples highlight the fact that security risks are currently the barrier for widespread adoption of cloud computing for mission critical applications and data (Chen, Paxson, & Katz, 2010; Kontzer, 2010). In addition, data is not the only security concern for organizations in the cloud. Chen et al. (2010) also identifies activity patterns as a vital asset that must be protected; activity patterns could be visible to other users sharing the same resources. The recommendation of this study is for cloud providers to address security concerns via a continuous monitoring program that is consistent with the recommendations included in NIST Special Publication 800-37, Revision 1 (National Institute of Standards and Technology [NIST], 2010).

Chapter 3 – Methodology

The Risk Management Framework (RMF) described in NIST Special Publication 800-37, Revision 1, is a framework that is intended to improve information security and strengthen the risk management process within federal agencies (NIST, 2010). The RMF is a six step process that defines risk related tasks that are to be executed during the system development life cycle, or against legacy systems if applied as a gap analysis (NIST, 2010). It provides guidance on how to maintain effective security controls despite constant changes in the internal and external environment while still allowing a high degree of flexibility to be exercised in implementing the process (NIST, 2010). This flexibility is what contributes to the effective application of this process within non-government organizations.

The last step, step six, in the RMF describes the security control monitoring process coined continuous monitoring (NIST, 2010). “Continuous monitoring is a proven technique to address the security impacts on an information system resulting from changes to the hardware, software, firmware, or operational environment” (NIST, 2010). Using the concept of step six in addition to the controls identified in Special Publication 800-53 Recommended Security Controls, continuous monitoring guidelines have been provided for a cloud environment to provide adequate security for itself and its tenants (Joint Task Force Transformation Initiative, 2009).

The goal of the continuous monitoring program within the cloud environment is to provide a clear picture of security on a near real-time basis. This program delivers consistent monitoring via automated tools with built-in steps for external review and testing to ensure that the controls are working as expected. This also provides tenants, or potential customers, of the cloud environment to remain confident in the security framework that the cloud provider offers.

A grid of the continuous monitoring recommendations is captured in Appendix A. Ultimately, this program is intended to help all parties manage risk within the environment and maintain the highest level of availability (Dempsey et al., 2010).

Chapter 4 –Results

Configuration Management

Configuration Management tends to be an area that experiences significant volatility and is a foundation for the need for continuous monitoring activities (Dempsey et al., 2010). The main components of managing configuration settings are to (1) document the baseline configurations, (2) identify/control necessary changes to the baseline, (3) implement configuration changes, and (4) monitor the configuration settings against the baseline (Johnson, Dempsey, Ross, Gupta, & Bailey, 2011). Each of these components are critical to maintaining security requirement baselines within the cloud environment.

Baseline configurations are documented when the security-focused configuration management (SecCM) program is first introduced or a new system is being introduced into the existing environment (Johnson et al., 2011). Automation tools like ServiceNow can help manage the baseline creation and storage (ServiceNow, 2011). The baseline configurations will include all systems (hypervisor, workstation, server, firewall, router, database, etc.) within the cloud architecture (Joint Task Force Transformation Initiative, 2009). Each system baseline will contain security requirements, at a minimum, for the operating system, applications, current versions, patch versions/service packs, model, hardware specifications, and location within the architecture (Joint Task Force Transformation Initiative, 2009). Using an application like ServiceNow provides a streamlined tracking system when issues are identified (incident or problem) to the correction activity (change request) (ServiceNow, 2011). This type of system distributes the information gathering activities, rather than managing it all within the Information Technology team, and creates a well-rounded inventory of changes.

As updates are made to systems, after testing and approvals are secured, the baselines are simultaneously updated using the change request information. The Change Management team is responsible for managing this continuous updating process. The Change Management team will work with the Patch and Vulnerability Group when they acquire, test, and distribute patches to the organization. Those systems that are customer impacting will include an additional communication step to update the customer of such changes. The Information System Security Officer is responsible for this communication.

Real-time monitoring of baseline security requirements is accomplished using an automated tool that can monitor the system infrastructure against baseline configurations to confirm compliance. The Network Administration team is responsible for this monitoring. Alerts are built in to notify the team of non-compliance indicating that remediation is necessary. Baseline security requirements will also be reviewed semiannually to ensure that the configuration is appropriate. Reports are pulled from the automated tools and reviewed by the Network Administration team, Information System Security Officer, and Security Control Assessor.

Patch and Vulnerability Management

The Patch and Vulnerability Group (PVG) manages the patch and vulnerability program (Mell, Bergeron, & Henning, 2005). With an enterprise patching solution in place it allows the PVG to “automatically push patches out to many computers quickly” (Mell et al., 2005). Having standardized system configurations provides an environment that is consistent from a maintenance standpoint but also makes testing patches much more streamlined and ultimately more successful (Mell et al., 2005). It is recommended to have as few system images as possible

to better maintain each variation and reduce the amount of testing and potential issues that can arise with non-standard system configurations.

The risk assessment policy is an input for this PVG team to operate effectively. In accordance with RA-1 the risk assessment policy and procedures must be documented and include roles and responsibilities, coordination between organizational entities, required compliance, and above all contain the appropriate level of leadership support (Joint Task Force Transformation Initiative, 2009). Having this policy in place provides the framework for how to address issues as they arise. It instructs the PVG and its leadership on what types of risks fall outside of the organizations risk tolerance. This policy need not be a separate document but must be addressed in order for this team to be effective.

This patch management program is applicable to all cloud models - SaaS, PaaS, and IaaS. Patches for the following are considered in scope for this program: Operating Systems, Client Applications, Server Applications, Enterprise Firewalls, Enterprise Network Intrusion Prevention Systems, Enterprise Antivirus and Antispyware Software, and Security Applications. In order to stay apprised of the constantly shifting security environment an enterprise patch management tool is utilized to obtain all available patches from supported vendors (Mell et al., 2005). Non-supported vendor patches are managed individually and are tracked using a separate system inventory.

A patch management tool, like IBM Tivoli Endpoint Manager, is utilized to capture updates automatically for supported vendors (IBM Corporation, 2011). The following are vulnerability management resources that are utilized as sources of timely information on security threats and for non-supported vendor patch information. In addition, vendor websites are

manually reviewed to ensure all applicable patches in the environment are reviewed on a monthly basis.

Table 1: External Vulnerability Resources

Source Name	Source Location	Description of Use
US-CERT National Cyber Alert System	http://www.us-cert.gov/cas/	A shared PVG email address is used to manage updates from the National Cyber Alert System. The PVG group will manage this mailbox daily and compile any updates that have not been captured via the patch management tool.
US-CERT Vulnerability Notes Database	http://www.kb.cert.org/vuls/	To help determine the priority of patches based on the cloud environment.
Open Source Vulnerability Database	http://www.osvdb.org/	Review manually daily. The PVG group will compile any updates that have not been captured via the patch management tool.
SecurityFocus Vulnerability Database	http://www.securityfocus.com/vulnerabilities	A shared PVG email address is used to manage updates from SecurityFocus. The PVG group will manage this mailbox daily and compile any updates that have not been captured via the patch management tool.
System	www.sans.org/sac	Review manually daily. The PVG group

Administration, will compile any updates that have not
Networking, and been captured via the patch management
Security Institute tool.
(SANS Institute)

All applicable patches are reviewed by the PVG on a daily basis. This review will consist of reviewing the patch management tool updates, the PVG mailbox updates, and the additional external vulnerability resources listed in Table 1. Once the vulnerabilities have been captured that apply to the cloud environment the PVG will then assess whether there are redundant patches and remove any duplication until they are left with a complete list of new vulnerabilities that apply to the environment. Lastly, the PVG team will determine the risk level of the vulnerability using a calculation of how many systems / users are impacted, how the vulnerability can be exploited, and the potential result if the vulnerability is exploited (Brykczynski & Small, 2003). This provides the PVG with the risk level of the vulnerability that has been identified in order to evaluate it against the stated organizational risk tolerance. These steps help the PVG determine the priority of each patch and whether the emergency patch process should be initiated.

The PVG will determine the testing schedule for the recommended changes based on the prioritization determined in the previous steps. Testing within a non-production environment is required for all patches prior to being deployed to the production environment to reduce the impact to cloud tenants. Once testing is completed a phased approach is used to apply to production starting with the least impacted areas first (Network World Staff, 2008). This provides the PVG an opportunity to evaluate and measure the impact of the update prior to

releasing to the entire impacted environment. Non-critical patches are released into production weekly using scheduled downtime that clients can anticipate and plan around (Thurman, 2006). This release schedule is published and released to cloud tenants on an annual basis. Any adjustments to this schedule are communicated to the tenants by the Information System Security Officer.

All system images that are used within the environment will also require updates to ensure that newly deployed equipment is up-to-date. All images are updated on a quarterly basis and all equipment released during that time will have updates pushed prior to deployment (Thurman, 2006). In addition, any vulnerability that is identified that does not have a patch developed is assessed for immediate configuration changes based on the risk assessment. This gap should follow the same emergency process as high risk vulnerabilities that have patches available (Mell et al., 2005).

An emergency process is in place to address and deploy critical patches in the environment immediately. The PVG reviews the patches that have been released and identifies any patches that require immediate action due to the vulnerability putting the organization outside of its risk threshold. This is determined using the risk assessment process defined in the risk assessment policy. The patch, or patches, that fall within this recommendation are fully documented by the PVG and presented to the Configuration Management Board which consists of the Information System Owner, Authorizing Official, and Senior Information Security Officer who will evaluate for immediate mitigation. These high risk vulnerability patches must be installed as soon as possible, but no later than 24 hours after the initial notification of the vulnerability.

The patch management tool will also be used to continuously monitor system compliance to ensure that all patches are implemented successfully and remain up-to-date. In addition, a semiannual review of the patch and vulnerability management program is conducted by the Configuration Management Board in coordination with the PVG. This certifies that the approach within the cloud environment still meets the needs of its customers. During this review, metrics are used to indicate the effectiveness of the program. The data points reviewed will cover susceptibility to attack, mitigation response time, and cost (Mell et al., 2005).

Patch management is a critical component to ensuring that systems and applications don't have exposed vulnerabilities. An example of how to confirm that patches are installed and vulnerabilities mitigated within the environment can be found in Appendix B. However, patch management is often a time consuming endeavor that requires appropriate testing prior to implementation. This lag time provides an opportunity for that vulnerability to be exploited prior to the patch getting implemented. Mell et al. (2005) emphasize that the time between a vulnerability being published and the release of malware developed to exploit the vulnerability has significantly reduced to weeks or even days. Given that our program is to capture new patches daily but implement non-critical, or lower risk, patches weekly creates risk in the environment. In this case, we will rely on additional controls to detect and stop malware prior to patches being deployed.

Antivirus / Malicious Software

Antivirus software is the key mitigating control that exists to catch known threats or infections prior to patches being installed (Mell et al., 2005). Least privileges are implemented to ensure that only a limited group of users have administrative access on servers, network devices, and desktop / laptops (Mell et al., 2005). Implementation of this control limits the effectiveness

of malware to exploit vulnerabilities since it typically requires administrator access to deploy (Mell et al., 2005).

Antivirus software typically uses two main techniques for identifying malicious software (malware), signature dictionary and suspicious behavior (Weaver, 2007). As malware is detected antivirus developers update the dictionary of known virus signatures. According to Mell et al. (2005) “major antivirus vendors usually release signatures for a significant new threat within several hours.” In this way clients receive signature updates frequently to ensure that known malware is identified and either deleted or quarantined (Weaver, 2007). The second approach is to monitor system activities and identify when certain behaviors appear suspicious (Weaver, 2007). This is done via heuristics techniques that assess files for suspicious code sequences or by looking for irregular activities when running the file in a virtual machine (Mell et al., 2005).

Antivirus software is installed on all systems that support the cloud services. The application settings are owned by the PVG and individual operators will not be able to make updates to the configuration or disable the service. Similar to the patch and vulnerability management program, centrally managed antivirus software, like Symantec Protection Center and Symantec Endpoint Protection, is utilized to acquire, review, test, and deploy signature updates (Symantec Endpoint Protection, 2011). Antivirus software should have various modules to address modern threats. The Symantec Endpoint Protection family covers the standard elements that should be addressed in a cloud environment by combining antivirus, antispayware, desktop firewall, intrusion prevention, device and application control, and network access control into a single agent (Symantec Endpoint Protection, 2011). The following are some specific features that should be included, whether the Symantec product or another application is implemented:

- (1) Protection against viruses, worms, Trojan horses, spyware, bots, zero-day threats and root kits (Symantec Endpoint Protection, 2011; Mell et al., 2005);
- (2) Rules-based firewall engine, browser protection, Generic Exploit Blocking shields (Symantec Endpoint Protection, 2011);
- (3) Real-time activity monitoring;
- (4) Scan storage (local and removable) weekly (Mell et al., 2005).

The cloud environment will house a minimum of two antivirus servers that are used for managing client software and distributing updates (Mell et al., 2005). These servers will have unrelated operating systems to reduce the impact of an attack against these servers. If one is taken down due to a targeted attack all servers won't be impacted (Mell et al., 2005). Staying consistent with the patch and vulnerability management program, all images are updated with current antivirus software versions and signatures on a quarterly basis and all equipment released during that time will have updates pushed prior to deployment. In addition, the centralized antivirus software management console will provide reports to determine any systems that are out of compliance and require updating or manual intervention. The PVG will monitor this report daily and address issues within 24 hours. Appendix B provides an example of how to confirm that antivirus signatures are installed and up-to-date within the environment. Global issues are escalated to the Senior Information Security Officer and the Risk Executive for mitigation plan.

Firewall Management

Firewalls are implemented to separate system and network environments within the cloud. This allows the cloud provider to offer more secure environments potentially for a private, community, or hybrid cloud deployment model that requires additional segregation of systems. Firewalls are configured to work with a set of rules to determine what network traffic is

acceptable and will be permitted (Mell et al., 2005). A network-based Intrusion Prevention Systems (IPS) is incorporated into the firewall to provide layers of protection at the perimeter and between network environments. An IPS is used in place of an Intrusion Detection System (IDS) due to its ability to not only perform up-front identification of an attack but it will also attempt to stop or block the attack that is detected. IDS technology will detect the attack but will not take any counter measures against it. The network-based IPS monitors network activity and detects irregular deviations from the baseline activity (Mell et al., 2005). These systems work together to determine what to allow or disallow into the network (Mell et al., 2005).

The necessary firewall and IPS patches are managed via the patch management process. However, there are other continuous monitoring activities that must be maintained to effectively manage the firewall and IPS solution (Scarfone & Hoffman, 2009). Performance of the firewall is critical and is monitored every hour. Alerts are configured to send an email to the Network Administration team when the performance strays from baseline. This alert is configured to send an email to the shared Network Administration mailbox in addition to mobile devices to ensure that the alert is responded to immediately. Monthly and year-to-date performance logs and reports are reviewed on a monthly basis by the Network Administration team, Common Control Provider, Information System Owner, and Senior Information Security Officer.

The firewall policy guides the set of rules that the firewall uses to direct incoming and outgoing network communications. Over time these policies will require adjustments to accommodate environmental changes or as a result of a new threat(s) (Scarfone & Hoffman, 2009). Changes can be introduced as needed to accommodate new products or services required within the organization. When changes of this kind are requested a complete impact assessment is conducted to ensure that changes don't impact other policies or rulesets already in place. This

assessment is driven by the Common Control Provider who will work with the Network Administration team, Information System Owner, and Senior Information Security Officer.

An emergency process is in place to address and deploy critical firewall updates to the environment immediately. The Network Administration team reviews the changes that are needed that require immediate action due to the vulnerability putting the organization at risk. This is determined using the risk assessment process defined in the risk assessment policy. The adjustments that are needed that apply for this emergency process are documented by the Network Administration team and presented to the Common Control Provider, Information System Owner, and Senior Information Security Officer. This team will evaluate the required change and determine the impact to the existing ruleset and configuration in place. This team indicates if there are any dependencies for the change and provides approval to implement. These changes will be implemented in production as soon as possible, but no later than 24 hours after the initial change notification.

Review of the firewall policies and rulesets is assessed quarterly, independent of changes, to ensure that all are still necessary and none are inadvertently missing. The review is conducted by the Common Control Provider, Network Administration team, Information System Owner, and Senior Information Security Officer. It will cover the complete assessment of the current state, any changes that occurred since the last review, who made the changes, who approved the changes, and what triggered the change (Scarfone & Hoffman, 2009). A complete policy and ruleset review is conducted by the Authorizing Official on an annual basis to confirm that the rules are appropriate and align with the organizations goals (Scarfone & Hoffman, 2009).

Since it is expected that firewall policies and rulesets change over time it is critical to keep frequent backups (Scarfone & Hoffman, 2009). Backups are scheduled monthly and are

completed and stored by the Network Administration team. Backups will also be taken prior to any changes being implemented to ensure that if a rollback of changes is necessary that it's accessible for immediate return to normal. The Common Control Provider will log all policy and ruleset decisions that are implemented.

In addition to maintaining the firewall configuration and assessing annually to validate that they support the goals of the organization, it is also important to confirm that the rules are complete and perform as expected. Therefore, penetration testing is scheduled semiannually to evaluate the overall security of the network. This is performed by a minimum of two associates on the Network Administration team with the oversight of the Information System Owner and Senior Information Security Officer. The schedule for this test is kept confidential to get a true simulation of network security. Appendix B also provides an example of how to confirm that firewall configuration working as expected within the environment.

Access Management

Managing operator access within the cloud environment is an essential control that limits access to only authorized users and distinguishes between functional responsibilities. An application is used for access requests that has a built in workflow component to automate the approval process prior to access being fulfilled. Aveksa is one example of a system that provides self-service to the user that requires access as well as workflow that facilitates human resource (HR) and organization management approval and then sends the access request to the Information Security Administration team for provisioning (Aveksa Inc., 2011). When HR provides approval for an employee to become active in the system the request will then route to the manager to ensure that the access is appropriate for the role of that individual.

Separation of duties must be maintained in order to thwart insider threats. The concept of separation of duties is a security principle with the primary objective of preventing fraud and errors by implementing a two-person integrity control (Humphreys, 2008). In the cloud environment it is crucial “that no single employee is in a position to introduce fraudulent, malicious code or data without being detected” (Humphreys, 2008). Role based access (RBAC) will be used in order to maintain this separation. RBAC essentially means that a role, or a collection of access entitlements, is assigned to a user based on their function within the organization (Joint Task Force Transformation Initiative, 2009). In the cloud environment the administrative functions will be segregated using roles to reduce the likelihood of an individual having significant access that could individually compromise the system or tenant data.

Consistent with AC-5, the Separation of Duties control, the following support functions will be segregated using roles: systems management, systems programming, configuration management, quality assurance and testing, network security, and database security (Joint Task Force Transformation Initiative, 2009). These roles will be enforced to ensure that no individual will perform the other’s responsibilities or have access to system layers that is outside of the stated function.

Roles are configured with a focus on separation of duties. They are defined by the Information System Security Officer and Information System Owner to ensure that they do not violate the separation of duties rules that have been defined within the cloud environment. The Information System Security Officer and Information System Owner will not have access to make updates to the role privileges. Roles are defined by the Information System Security Officer and Information System Owner and will require approval from both individuals prior to any changes being introduced to these roles. Approval will be submitted in the form of a change

request via a change management tool, like ServiceNow (ServiceNow, 2011). The Information Security Administration team is responsible for creating and updating the entitlements within each role. This is a manual layer of control that will limit any updates to the roles that may contradict the separation of duties rules defined.

Reporting and logs are used to monitor activities within the system to ensure that no user is performing functions outside of their role access. This automated control is continuously monitored on a daily basis to ensure that these barriers remain intact. If the separation of duties rules are violated an alert will be sent to the Information Security Administration team and the Information System Security Officer for immediate review, root cause assessment, and remediation. Additional teams, like the Network Administration team, will be a part of the review and root cause analysis as needed. In addition, semiannual reviews of the roles will be conducted to recertify that the roles in use are appropriate. This review also assesses any changes to the environment and determines if the roles are still sufficient or if adjustments are necessary to retain complete separation of duties. Reviews will be conducted by the Information System Security Officer, Information System Owner, Senior Information Security Officer, and Risk Executive.

If an employee changes job function, HR notifies the Information Security Administration team. The Information Security Administration team then requests a recertification of access to the employee's new manager. If approval to retain access is not received the Information Security Administration team will revoke access. If the Human Resources system is not integrated with the access provisioning system, termination notices are submitted daily to the Information Security Administration team in order to completely revoke access. Terminations are processed within the same day. An entitlement review is initiated using

an automated tool, like Aveksa, each quarter to validate the active users in the system as well as individual user entitlements (Aveksa Inc., 2011).

Using an automated tool for managing operator accounts and operator entitlements provides an audit trail to support future audit requests to ensure that all access provisioned was appropriate. Organization Managers will review the active account report quarterly to ensure that the accounts are appropriate. Operator entitlements will also be reviewed by the Organization Manager to validate that the user is serving in the function that the role and access provisioned is appropriate for. Testing a sample of user access requests to validate that the process is working effectively is done on an annual basis by the Information System Security Officer and Security Control Assessor. Reports will also be produced to track the speed that system access has been revoked when no longer needed. These reports are reviewed on a quarterly basis by the Information System Security Officer, Security Control Assessor, and Information Security Administration team to confirm that the process is working effectively and in sync with the entitlement review process.

Chapter 5 – Conclusions

The cloud environment is unique in the way that there is shared responsibility to maintain effective security controls and ensure that there is limited security exposure. Tenants will be responsible for a piece of the security puzzle and must be accountable for a portion of the security controls and continuous monitoring. “Cloud Providers and Cloud Consumers collaboratively design, build, deploy, and operate cloud-based systems. The split of control means both parties now share the responsibilities in providing adequate protections to the cloud-based systems” (Liu et al., 2011). The next step in this research is to determine what a continuous monitoring program looks like for the cloud tenant to complement the cloud provider’s continuous monitoring program (Liu et al., 2011).

Cloud providers, like Amazon, have internally developed tools that are used for monitoring and security related activities (Amazon Web Services, 2011). This proprietary development has some pros and cons for the cloud model and for its customers. One of the benefits is that there can be a custom integration process for managing the security within the unique cloud infrastructure. This can lead to a higher level of service for the tenant and allows the cloud provider to provide a high level of security and availability. A drawback to this internal development is the lack of public awareness for vulnerabilities that may exist. It can also mean that the systems may not hold to industry parity over time. Internal development is known to be more costly and given increased competition in the cloud sector it may become challenging for the cloud provider to maintain the highest quality of security in the long term.

This program should be implemented in a cloud environment to understand the full benefits of the proposed continuous monitoring program. This program is intended to provide a mature security infrastructure that the cloud provider can maintain and communicate to

customers. A high level of transparency is required between the cloud provider and its tenants in order for the program to be successful given the unified nature of the cloud environment.

References

- (2011). Hypervisor. Retrieved from <http://en.wikipedia.org/wiki/Hypervisor>.
- Amazon Web Services. (2011). Amazon Web Services: Overview of Security Processes. Retrieved from http://d36cz9buwru1tt.cloudfront.net/pdf/AWS_Security_Whitepaper.pdf.
- Anthes, G. (2010). Security in the Cloud. *Communications of the ACM*, 53(11), 16-18. doi:10.1145/1839676.1839683.
- Aveksa Inc. (2011). Automating Access Governance: Overview. Retrieved from <http://www.aveksa.com/resources/upload/Aveksa-Overview-DS-web.pdf>.
- Brykczynski, B., & Small, R. (2003). Reducing Internet-Based Intrusions: Effective Security Patch Management. Retrieved from <http://www.computer.org.dml.regis.edu/plugins/dl/pdf/mags/so/2003/01/s1050.pdf?template=1&loginState=2&userData=Regis%2BUniversity%253ARegis%2BUniversity%253AAddress%253A%2B207.93.211.102%252C%2B%255B140.98.196.191%252C%2B%2B64.215.172.240%252C%2B207.93.211.102%255D>.
- Chen, Y., Paxson, V., & Katz, R. (2010). What's New About Cloud Computing Security?. Retrieved from <http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-5.pdf>.
- Cloud Computing. (2009). *Library Technology Reports*, 45(4), 10-12. Retrieved from EBSCOhost.
- Dempsey, K., Johnson, A., Jones, A.C., Orebaugh, A., Scholl, M., & Stine, K. (2010). Information Security Continuous Monitoring for Federal Information Systems and

- Organizations. Retrieved from <http://csrc.nist.gov/publications/drafts/800-137/draft-SP-800-137-IPD.pdf>.
- Goodburn, M. A., & Hill, S. (2010). The Cloud Transforms Business. *Financial Executive*, 26(10), 34-39. Retrieved from EBSCOhost.
- Humphreys, E. (2008). Information security management standards: Compliance, governance and risk management. Retrieved from http://www.sciencedirect.com.dml.regis.edu/science?_ob=MiamiImageURL&_cid=271961&_user=1922016&_pii=S1363412708000514&_check=y&_origin=&_coverDate=30-Nov-2008&_view=c&_wchp=dGLbVlt-zSkWz&_md5=d2da22b7dd759128de2b386128f1bdcf/1-s2.0-S1363412708000514-main.pdf.
- IBM Corporation. (2011). IBM Tivoli Endpoint Manager for Security and Compliance. Retrieved from <http://public.dhe.ibm.com/common/ssi/ecm/en/tid14075usen/TID14075USEN.PDF>.
- Johnson, A., Dempsey, K., Ross, R., Gupta, S., & Bailey, D. (2011). Guide for Security-Focused Configuration Management of Information Systems. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-128/sp800-128.pdf>.
- Joint Task Force Transformation Initiative. (2009). Recommended Security Controls for Federal Information Systems and Organizations: NIST Special Publication 800-53 Revision 3. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final.pdf>.
- Kontzer, T. (2010). Cloud Forecast 2015. *CIO Insight*, (114), 8-10. Retrieved from EBSCOhost.

- Liu, F., Tong, J., Mao, J., Bohn, R., Messina, J., Badger, L., & Leaf, D. (2011). NIST Cloud Computing Reference Architecture: Recommendations of the National Institute of Standards and Technology, Special Publication 500-292. Retrieved from http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/ReferenceArchitectureTaxonomy/NIST_SP_500-292_-_090611.pdf.
- Mell, P., Bergeron, T., & Henning, D. (2005). Creating a Patch and Vulnerability Management Program. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-40-Ver2/SP800-40v2.pdf>.
- Mell, P., & Grance, T. (2009). The NIST Definition of Cloud Computing: Version 15. Retrieved from <http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc>.
- Network World Staff. (2008). Guide to Patch and Vulnerability Management: Patch management best practices. Retrieved from http://www.pcworld.com/businesscenter/article/144636/guide_to_patch_and_vulnerability_management.html.
- National Institute of Standards and Technology. (2010). Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach, NIST Special Publication 800-37 Revision 1. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf>.
- Owens, D. (2010). Securing Elasticity in the Cloud. *Communications of the ACM*, 53(6), 46-51. doi:10.1145/1743546.1743565.
- Ryan, W., & Loeffler, C. M. (2010). Insights into Cloud Computing. *Intellectual Property & Technology Law Journal*, 22(11), 22-28. Retrieved from EBSCOhost.

Scarfone, K., Souppaya, M., Cody, A., Orebaugh, A. (2008). Technical Guide to Information Security Testing and Assessment: Recommendations of the National Institute of Standards and Technology, Special Publication 800-115. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-115/SP800-115.pdf>.

Scarfone, K., Souppaya, M., & Hoffman, P. (2011). Guide to Security for Full Virtualization Technologies: Recommendations of the National Institute of Standards and Technology, Special Publication 800-125. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf>.

ServiceNow. (2011). Solutions. Retrieved from <http://www.service-now.com/solutions.do>.

Smith, R. (2009). Computing in the Cloud. *Research Technology Management*, 52(5), 65-68. Retrieved from EBSCOhost.

Symantec Endpoint Protection. (2011). Data Sheet: Endpoint Security. Retrieved from http://www.symantec.com/content/en/us/enterprise/fact_sheets/b-sep_DS_21194634.en-us.pdf.

Talbot, D. (2010). Security in the Ether. *Technology Review*, 113(1), 36-42. Retrieved from EBSCOhost.

Thurman, M. (2006). WMF Vulnerability Sparks Patch Program. Retrieved from <http://search.ebscohost.com.dml.regis.edu/login.aspx?direct=true&db=bth&AN=19753190&site=bsi-live>.

Weaver, R. (2007). Guide to Network Defense and Countermeasures. Boston: Cengage Learning.

Appendix A

Continuous Monitoring Grid

Category	Control	Continuous Monitoring Frequency	Roles Responsible
Configuration Management	Review Baseline Configuration for Real Time Compliance	Daily	Network Administration team
	Review Baseline Security Requirements	Semiannually	Network Administration team, Information System Security Officer, Security Control Assessor
Patch and Vulnerability Management	Capture New Patches	Daily	Patch and Vulnerability Group
	Review New Non-Critical Patches	Daily	Patch and Vulnerability Group
	Implement Patches that require Server Reboot	Weekly	Patch and Vulnerability Group
	System Compliance with Patch status	Daily	Patch and Vulnerability Group
	Emergency Patches	As Needed / within 24 hours of notification	Patch and Vulnerability Group, Configuration Management Board
	Maintain Patches on Images	Quarterly	Patch and Vulnerability Group
	Review Overall Patch and Vulnerability Program	Semiannually	Patch and Vulnerability Group, Configuration Management Board
Antivirus / Malicious Software	Maintain Signatures on Images	Quarterly	Patch and Vulnerability Group
	Manage Compliance with systems	Daily	Patch and Vulnerability Group
	Real-time activity monitoring	24/7/365	End Point Application

	Full System Scanning	Weekly	End Point Application
	Corporate Wide Signature Updates	Daily	Patch and Vulnerability Group, Antivirus Central Management Console
Firewall Management	Firewall Performance Real-Time Monitoring	Hourly with Alerts configured	Network Administration
	Firewall Performance Management - Reporting	Monthly	Network Administration, Common Control Provider, Information System Owner, Senior Information Security Officer
	Firewall Policy and Ruleset Maintenance	As Needed	Common Control Provider, Network Administration, Information System owner, Senior Information Security Officer
	Emergency Firewall Policy and Ruleset Changes	As Needed / within 24 hours of notification	Network Administration , Common Control Provider, Information System Owner, Senior Information Security Officer
	Log Policy and Ruleset Changes	As Needed	Common Control Provider
	Review Firewall Policy and Ruleset	Quarterly	Common Control Provider, Network Administration, Information System Owner, Senior Information Security Officer
	External Review Firewall Policy and Ruleset	Annually	Authorizing Official

	Policy and Ruleset Backup	Monthly / As Needed	Network Administration team
	Penetration Testing	Semiannually	Subset of Network Administration team, Information System Owner
Access Management	Access Provisioning	As Needed	Information Security Administration team
	Role Changes	As Needed	Information System Security Officer, Information System Owner, Information Security Administration team
	Separation of Duties Monitoring	Daily	Information Security Administration team, Information System Security Officer
	Role Review for Separation of Duties	Semiannually	Information System Security Officer, Information System Owner, Senior Information Security Officer, Risk Executive
	Termination Fulfillment	Daily	Human Resources, Information Security Organization Manager
	Active Account Review	Quarterly	Information Security Organization Manager
	Operator Entitlements	Quarterly	Information Security Organization Manager
	Testing Access Process	Annually	Information System Security Officer, Security Control Assessor
	Report Metrics for Access Revocation	Quarterly	Information System Security Officer, Security Control Assessor, Information Security Administration team

Appendix B

Example Control Enforcement

Patch and vulnerability management example.

Confirming that a remediation has been successful and the vulnerability has been mitigated is an important step in the patch and vulnerability management process. One way to provide assurance that patches have been installed as planned is to perform a network scan with a vulnerability scanner (Mell et al., 2005). “A vulnerability scanner identifies not only hosts and open ports on those hosts, but also associated vulnerabilities” (Mell et al., 2005). These systems use databases of vulnerabilities which must be updated frequently so that it can identify the newest vulnerabilities (Mell et al., 2005). This program recommends weekly patch implementations. Vulnerability scans are conducted post patch implementation to ensure the remediation was successful. Logs will be reviewed from the vulnerability scanner to identify any false positive results. Notations must be made in the change management tool to indicate that the patch implementation was successful and attach the log files as evidence. The change request cannot be completed until this confirmation step has been completed.

Antivirus / malicious software example.

Confirm that antivirus signatures are installed and up-to-date within the environment in addition to the managed enterprise application monitoring real time. Reviewing antivirus logs can provide details on update failures and other indications of outdated signatures and software (Scarfone, Souppaya, Cody, & Orebaugh, 2008). Logs are reviewed weekly by PVG and Information System Security Officer. Using an automated audit tool to review the logs and provide a specific view on the information needed to confirm that the environment is in compliance with the appropriate antivirus definition files with less effort than reviewing

manually (Scarfone et al., 2008). This summary view produced by the audit tool can also be shared with cloud tenants and auditors as evidence that antivirus program is effective.

Firewall management example.

Confirm that firewall policy and rulesets are configured within the environment and are working as expected. Reviewing firewall and IPS logs can provide details on the traffic that is being allowed into the network (Scarfone et al., 2008). Reviewing these logs can identify issues with the current firewall configuration if traffic is coming through the firewall that should be disallowed based on the policy (Scarfone et al., 2008). Using an automated audit tool to review the logs and provide a specific view on the information needed to confirm that the environment is in compliance with the appropriate policy and rulesets with less effort than reviewing manually (Scarfone et al., 2008). This summary view produced by the audit tool can also be added to the inventory to confirm that the baseline configuration is in place and working as expected.

Glossary

Term	Definition
AWS	Amazon Web Services
Community	<p>An environment shared by many organizations in particular industries, by geography, along similar supply chains or otherwise connected. Establish cooperation between suppliers, providers, customers.</p> <p>Cloud infrastructure shared by several organizations that support a specific community that has shared concerns. The shared concerns could be the mission, security, privacy, policy, or regulatory compliance</p>
DoS	Denial of service
DDoS	Distributed denial of service
HR	Human Resources
Hybrid	Involves a composition of two or more of the three preceding models.
Hypervisor	In computing, a hypervisor, also called virtual machine monitor (VMM), is one of many virtualization techniques which allow multiple operating systems, termed guests, to run concurrently on a host computer, a feature called hardware virtualization.
IaaS	Cloud Infrastructure as a Service. Storage, processing, and network services.
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
IT	Information Technology
PaaS	Cloud Platform as a Service. Development, testing, deployment, hosting, and maintenance services.
NIST	National Institute of Standards and Technology

Private	<p>Cloud deployment model: a “closed” environment for a single organization hosted by a third party.</p> <p>Maintain all the technology components, servers, and software for a single organization. The solution may be managed by the user or a third party but is provided for the benefit of only one organization. Private clouds are increasingly being deployed within larger enterprises.</p>
Public	<p>A shared environment used by many organizations.</p> <p>Available to anyone or to large industry groups and is owned by the provider of the service. Offers the greatest potential flexibility and savings but also involves granting the service provider the greatest control over the enterprise’s technology capabilities. Large enterprises are using this deployment for discrete services and are evaluating ways to further use the model.</p>
PVG	Patch and Vulnerability Group
RBAC	Role Based Access
Risk	<p>Risk is a measure of the extent to which an entity is threatened by a potential circumstance or event, and a function of:</p> <p>(i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence.</p>
SaaS	Cloud Software as a Service. Web application usage services.
SecCM	Security-focused configuration management
SPI Models	IaaS, PaaS, SaaS
ToS	Terms of Service

Virtualization	The simulation of the software and/or hardware upon which other software runs.
----------------	--
