

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) OCTOBER 2013		2. REPORT TYPE Conference Paper		3. DATES COVERED (From - To) APR 2011 – JUN 2013	
4. TITLE AND SUBTITLE Near-Real-Time Cloud Auditing for Rapid Response (POST PRINT)				5a. CONTRACT NUMBER IN-HOUSE GGIHZORR	
				5b. GRANT NUMBER N/A	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) <i>Syracuse University: Joon S. Park</i> <i>State University of New York, Institute of Technology: Edward Spetka</i> <i>Air Force Research Laboratory: Keesook J. Han, Hassan Rasheed, Paul Ratazzi</i>				5d. PROJECT NUMBER GGIH	
				5e. TASK NUMBER ZO	
				5f. WORK UNIT NUMBER RR	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Syracuse University, 900 South Crouse Ave, Syracuse NY 13244 State University of New York, Institute of Technology, 100 Seymour Rd, Utica NY 13502 Air Force Research Laboratory/RIGA 525 Brooks Rd, Rome NY 13440				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/Information Directorate Rome Research Site/RIGA 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TP-2013-060	
12. DISTRIBUTION AVAILABILITY STATEMENT Distribution Approved For Public Release; Distribution Unlimited. PA Case number: 88ABW-2012-4497, dated 18 Aug 2011					
13. SUPPLEMENTARY NOTES © IEEE 2012. This paper was published in the Proceedings of the 26 th International Conference on Advanced Information Networking and Applications Workshops (WAINA 2012) Fukuoka, Japan, 26-29 Mar 2012. This work is copyrighted. One or more of the authors is a U.S. Government employee working within the scope of their Government job; therefore, the U.S. Government is joint owner of the work and has the right to copy, distribute, and use the work. All other rights are reserved by the copyright owner.					
14. ABSTRACT Due to the rapid emergence of Information Technology, cloud computing provides assorted advantages to service providers, developers, organizations, and customers with respect to scalability, flexibility, cost-effectiveness, and availability. However, it also introduces new challenges and concerns, especially in terms of security and privacy. One of the major security obstacles to widespread adoption of cloud computing is the lack of near-real-time auditability. In particular, near-real-time cloud auditing, which provides timely evaluation results and rapid response, is the key to assuring the cloud. In this paper, we discuss security and privacy concerns in cloud computing and the current status of cloud auditing efforts. Next, we address the strategies for reliable cloud auditing and analyze the deficiencies of current approaches. We then discuss the summary of our case study with Amazon CloudWatch, which is one of the most developed cloud-monitoring APIs.					
15. SUBJECT TERMS Cloud Auditing, Cloud Computing, Network Monitoring, Rapid Response, Security					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES ï	19a. NAME OF RESPONSIBLE PERSON KEESOOK J. HAN
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) N/A

Near-Real-Time Cloud Auditing for Rapid Response

Joon S. Park
School of Information Studies
(iSchool)
Syracuse University
Syracuse, NY 13078, USA
jspark@syr.edu

Edward Spetka
State University of New York
Institute of Technology (SUNYIT)
Utica, NY 13502, USA
ed@spetka.net

Hassan Rasheed, Paul Ratazzi,
Keesook J. Han
Information Directorate
Air Force Research Laboratory
(AFRL)
Rome, NY 13440, USA
HassanRasheed@acm.org, {Paul.
Ratazzi, Keesook.Han}@rl.af.mil

Abstract— Due to the rapid emergence of Information Technology, cloud computing provides assorted advantages to service providers, developers, organizations, and customers with respect to scalability, flexibility, cost-effectiveness, and availability. However, it also introduces new challenges and concerns, especially in terms of security and privacy. One of the major security obstacles to widespread adoption of cloud computing is the lack of near-real-time auditability. In particular, near-real-time cloud auditing, which provides timely evaluation results and rapid response, is the key to assuring the cloud. In this paper, we discuss security and privacy concerns in cloud computing and the current status of cloud auditing efforts. Next, we address the strategies for reliable cloud auditing and analyze the deficiencies of current approaches. We then discuss the summary of our case study with Amazon CloudWatch, which is one of the most developed cloud-monitoring APIs.

Keywords *cloud computing; monitoring; auditing; rapid response; security*

I. INTRODUCTION

Cloud computing has become a compelling business model for companies that own large data centers to essentially rent out different layers of their computing resources. This phenomenon has emerged in large part because many companies now internally rely both on technologies that must be able to scale dynamically and on large sets of computing hardware for their processing. These organizations have developed a high level of proficiency, deploying scalable applications over virtualized architectures [1, 2] with commodity hardware to the extent that they can easily provide these resources as a service to others. When economies of scale are achieved, companies achieve the ability to provide a service to consumers more cheaply than it would cost for the consumer to make that same resource internally within their own organization. Therefore, organizations have found that they can reduce IT costs by merely outsourcing one or more types of their internal IT infrastructure to a cloud service provider.

The National Institute of Standards and Technology (NIST) has formalized the cloud services being offered to consumers into three different types [3]: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS).

At the most basic level, IaaS is the model in which a service provider makes a set of virtualized computer components (e.g., virtual machines or networked storage) available to consumers as a fully outsourced service that can be used to build and run applications without purchasing the actual computer components. Each of these hardware components is visible to the consumer. An example of an IaaS service is Amazon EC2 [4], which allows customers to dynamically provision new virtual machines on the Amazon cloud using configuration templates with their own custom settings.

Moving to the next highest level, PaaS is the model in which a service provider or vendor offers to the client an entire runtime environment for an application (i.e., design, deployment, and test). This would include the basic hardware in addition to critical libraries that the consumer's application will rely on. The underlying hardware, however, is not visible to users, and they do not have the same level of control over those resources as IaaS customers do. An example of a PaaS service is Google AppEngine [5], which allows developers to make use of some of the same tools and libraries used in Google Web applications to build their own custom applications by building on a common application platform.

At the highest level, SaaS provides consumers with the use of on-demand software (e.g., business applications) running on the provider's servers and the associated data over the Internet. A consumer only views the specific application features that the service provider surfaces for him or her. An example of an SaaS is Google Apps, which provides consumers with personal, education, and business Web applications such as email, calendar, video, and documents for teams.

In a cloud computing business model, many aspects of the infrastructure are abstracted away from the end-user. Costs for operation, sustainment, technology update, etc., are directly borne by the cloud service provider (CSP) and amortized across all customers by way of their service level agreement (SLA). CSPs also typically provide basic security features such as physical security, customer isolation, authentication, customer-configured firewalls, and APIs (Application Programming Interfaces) for accessing certain security-related statistics and parameters. While none of these issues is new in the world of computing, compared

with traditional infrastructures, cloud computing architectures exhibit a different partitioning with respect to security and privacy issues [6, 7, 8, 9, 10].

In this paper, we discuss security and privacy concerns in cloud computing and the current status of cloud auditing efforts. Next, we address the strategies for reliable cloud auditing and analyze the deficiencies of current approaches. We then discuss the summary of our case study with Amazon CloudWatch, which is one of the most developed cloud-monitoring APIs.

II. SECURITY AND PRIVACY ISSUES IN CLOUD COMPUTING

With the movement toward widespread adoption of cloud computing, security and privacy issues have become major concerns [11, 12, 13]. For instance, third-party possession of personal documents raises questions about control, service payment, and ownership of the contents. Currently, there are no standards governing the cloud-provider industry with regard to security and privacy. As a result, each provider has full discretion over how it will manage these important issues.

Threats against cloud architectures include most, if not all, of the traditional threats to information security, such as race conditions, buffer overflows, SQL injection, etc. Cloud architectures also introduce new classes of attacks. These include poisoned virtual machines, attacks against the CSP (Cloud Service Provider) management/administration console, attacks based on knowledge of default security settings, abuse of billing systems, attacks that abuse the trust associated with the CSP's namespace, data leakage via uniform resource locators (URLs), and others. Dealing with some of these is the responsibility of the CSP, while others lie within the customer's purview. Understanding this division of responsibility is critical when implementing security measures designed to counteract these threats. Some risks can be substantially reduced by the customer, while others are entirely dependent on the CSP's approach to security.

Laws and regulations such as the USA Patriot Act require cloud providers to release personal and other user data to government authorities in response to a search warrant or subpoena. Data residing in a public cloud is not secured or maintained by the owner, which means that he or she has no recourse when governmental or law enforcement agencies (whether in the owner's country or somewhere else) demand delivery of the data. Additionally, the Patriot Act deters some governments outside of the U.S. from allowing data generated within their borders to reside in the U.S. This places a data portability constraint on cloud computing. If a provider does not have a data center in a particular country where a potential customer has operations, it may not make sense for the potential customer to utilize that provider's services.

One of the major security obstacles to widespread adoption of cloud computing is the lack of near-real-time auditability. While the general purpose of monitoring is to observe the current status of the target system, we define *auditing* as evaluation of the target system based on pre-

defined criteria and monitoring results. In particular, near-real-time cloud auditing, which provides monitoring results and rapid response corresponding to the outcome of evaluation, is the key to assuring the cloud. For instance, when the monitoring results report an ongoing cyber attack in the cloud, the security evaluation determines a threat level higher than that of a normal condition in runtime, which initiates a rapid response to protect the cloud (e.g., deployment of additional defense mechanisms in the cloud).

The majority of cloud services are public and multitenant, which means multiple users access a single application instance via the same infrastructure. This is unlike multi-instance architecture, which provides each user with one application instance. Because of its resource-sharing capability, multitenant architecture requires the prevention of data exposure among the users. Even if some clouds have strong security measures in place, the ultimate security level of the entire service is only as strong as the weakest point in the clouds. For instance, one user's security vulnerability can become the entire cloud's weakness. A service provider could open the door to security threats in resource provisioning or during distributed application execution. Furthermore, the vulnerabilities in cloud computing environments can be exploited to carry out Internet crimes.

Currently, CSPs do not have robust technical solutions that can protect their clouds from harmful malware, virus infection, botnets, distributed denial of service attacks, or other types of cyber attacks. Furthermore, there is no effective mechanism to help cloud users evaluate the security measures of their service providers and ensure the protection of their data while taking into consideration industry standards or personal preferences. Therefore, the near-real-time auditing capability is an immediate requirement for reliable cloud computing services.

III. CURRENT STATUS OF RELATED EFFORTS

A. Working Groups

Recognizing the demand for cloud computing, several working groups have recently launched to meet the business and assurance challenges in the cloud environment. For instance, in January 2010, the CloudAudit/A6 working group [14] was launched. Its ultimate goal is to help cloud service providers with possible automation of the audit, assertion, assessment, and assurance of their services. The Cloud Security Alliance (CSA [15]) was formed to promote education about and use of practices for cloud computing security. It has announced the general security guidelines for cloud computing and top threats to cloud computing. NIST (National Institute of Standards and Technology) launched a cloud-computing group in its computer security division to promote the secure use of the technology within government and industry by providing technical guidelines and standards [16]. The U.S. Department of Energy, the Defense Information Systems Agency (DISA), and other cross-industry groups have also presented their visions of cloud computing from the perspectives of the federal government, the military, and industry, respectively [17, 18, 19, 20, 21].

B. Standards

PCI DSS [22, 23] is one of the standards most frequently referred to in the online payment industry, as it specifies the requirements for handling cardholder information for the major transaction cards, such as credit, debit, and ATM cards. The standard contains 12 core requirements in six main areas: building and maintaining a secure network, protecting cardholder data, maintaining a vulnerability management program, implementing strong access control measures, regularly monitoring and testing networks, and maintaining an information security policy. Organizations wishing to gain certification against the requirements of this standard must get an assessment from a security specialist that is approved by PCI DSS. The most recent version of the PCI DSS added language which explicitly clarifies that virtual components are also included under the heading of system components to which the standard applies. It also makes allowances for multitenant virtual architecture that enables a single hardware server to effectively host multiple virtual machines with different functions, as long as each of the virtual machines has only one primary function.

SAS70 (Statement on Auditing Standard) is a widely recognized auditing standard developed by the American Institute of Certified Public Accountants (AICPA). SAS70 is an auditing methodology rather than an actual auditing rule. Therefore, the organization being audited should specify a series of control objectives and corresponding activities. SAS70 is the source of criteria and guidelines for CPAs to conduct auditing on the control capability of service providers. There are two different audit types according to the standard: Type I (quick, usually first time) and Type II (long, thorough).

C. Industry Recommendation

Recently, a group of industry organizations—including HyTrust, Cisco, Coalfire, Savvis, and VMware—recommended an architecture that would theoretically allow organizations using private cloud architectures to achieve PCI DSS compliance [24]. This group proposed technical solutions for various risks and controls that were needed to meet the PCI DSS requirements, including firewall maintenance, default setting change, protection of stored data, encryption of data transmission, usage of antivirus software, development/maintenance of secure systems, access control to data, accountability, physical access control, tracking/monitoring, security test of systems/processes, and maintenance of security policies. Although many of these risks are abstract and lacking in specificity, they provide insight into the unique security challenges for virtualized cloud environments.

IV. STRATEGIES FOR RELIABLE CLOUD AUDITING

Although cloud computing provides assorted advantages to service providers, developers, organizations, and users with respect to mobility, flexibility, cost-effectiveness, availability, and maintenance, it also introduces new challenges and concerns, which are described in Section II.

Many researchers, government agencies, and vendors have exerted significant efforts to make the clouds more reliable. There are many ways to achieve this goal, but we believe that situation awareness via reliable cloud auditing is the key element of the ultimate solution. Therefore, we have surveyed the currently available monitoring tools, research literature, standards, and other resources related to IA (Information Assurance) metrics and IT auditing. In the following subsections, we identify and discuss the deficiencies of the current approaches to reliable cloud auditing based on the outcomes from our research.

A. Scope of Cloud Auditing

In general, the main purpose of monitoring is to observe the target system. In our work, we define *auditing* as evaluation of the target system based on the monitoring results and pre-defined criteria. Therefore, monitoring is still part of cloud computing, but the monitoring results should be evaluated in a timely manner for the purpose of auditing. Apparently, most current monitoring services need to be improved with respect to their evaluation capability. In particular, applying the auditing concept to the clouds, a reliable cloud environment requires near-real-time auditing services that can provide the current evaluation results of the performance and IA status in the clouds. The evaluation results may initiate a corresponding rapid response in order to maintain the quality of service in the cloud environment. For instance, when the monitoring service reports that the user's location has been changed (e.g., from a secure office to an insecure public area), the result of the trust-level evaluation is changed (e.g., from High to Low). The evaluation result may initiate rapid response by providing the result to the access control module, which will degrade the user's privilege accordingly while the user stays in the insecure area.

B. IA Metrics for Cloud Auditing

The practical application of cloud computing is relatively new. Many previous and current efforts mainly focus on the promotion of cloud computing usage and addressing its potential benefits. Fewer efforts have been directed in the area of cloud auditing. Furthermore, little effort has been exerted on IA auditing, while monitoring services primarily depend on performance-related metrics such as CPU/memory/disk usage, latency, bandwidth, I/O rate, network traffic status, and so on. However, in order to support more robust cloud auditing services for assured clouds, both performance and IA metrics should be measured and evaluated in a timely manner.

C. Near-Real-Time Evaluation

There are comprehensive security and IT standards that can provide the basis for cloud auditing metrics, including ISO27002, CoBIT, SAS70, and so on. However, they were not originally developed for multitenant cloud environments. Furthermore, their auditing scopes are too broad, including organizations, systems, users, policies, and operations for the purpose of near-real-time cloud auditing. Many of these metrics require human interaction or synchronized

communications via interviews, questionnaires, surveys, etc. Although we still need these comprehensive measurements to enhance the overall quality of cloud computing, we cannot measure all of them in runtime. Therefore, we should define the IA metrics that can be evaluated in near-real-time. Some examples of metrics are the numbers of outstanding vulnerabilities, attacks detected, risk acceptances, weak passwords, and anomalies.

D. Commonality of Audit Trails and Metrics

While conducting the survey on monitoring tools and IA evaluation metrics, we found that there is a lack of commonality in audit trails and metrics. For instance, different tools adopt various technologies for data management and use different formats for their audit trails, which are typically large in size and accumulated every second in runtime. In reality, we most likely need multiple monitoring tools and sensors for cloud auditing in a distributed environment. A single tool cannot cover all the target areas, including infrastructures, platforms, applications, and user behaviors. In fact, the audit trails and metrics from different sources should be converted to a consistent format before the evaluation. Technically, retrieving only related audit trails from a large amount of data set delays the evaluation process. Therefore, we need to improve the overall performance of audit trail analysis by developing the commonality of audit trails and metrics.

E. Cloud Audits for Different Actors

One of the main advantages of cloud computing is that the technology can support multitenant services which consist of multiple actors, including service and infrastructure providers and customers. Currently, several working groups and research efforts have focused on the CSPs' ability to monitor the status of the clouds and the customers. However, there has been no significant effort to determine the customer's ability to evaluate the CSPs' service quality in terms of performance and security measures. In order to satisfy the customer's requirements, RSA, Intel, and VMware recently announced the proof of concept, the Hardware Root of Trust [25], whose end goal is to enable the cloud providers and their customers to measure and monitor security conditions within a private cloud. Therefore, for a more reliable cloud-auditing service, we believe that both the cloud providers and the customers should be able to audit the current status of the clouds with customized interests.

V. A CASE STUDY ON AMAZON CLOUDWATCH

Since Amazon CloudWatch [26] is one of the most developed cloud monitoring APIs (Application Programming Interfaces) that provides real-time monitoring on the Amazon Web Service (AWS), in this section we analyze its primary functionalities and discuss the capabilities to be improved for reliable cloud auditing based on the strategies we propose in Section IV.

A. Analysis of CloudWatch Capabilities

Amazon CloudWatch monitors cloud resources such as Amazon EC2 (Elastic Compute Cloud), EBS (Elastic Block Store) volumes, RDS (Relational Database Service), DB instances, and ELB (Elastic Load Balancing). The updated monitoring information is periodically reported to the subscribers according to their configuration. Depending on the level of service to which a customer is subscribed, services may send updated information to CloudWatch every five (free basic monitoring) or one (premium detailed monitoring) minutes. The measuring metrics mainly focus on resource usage and performance monitoring, including CPU utilization, latency, memory usage, transaction volumes, and error rates. Different services provide different metrics based on the nature of the service. There are also four aggregate operations that can be used to tailor calculation of the metrics—minimum, maximum, sum, and average—although not all operations are available with each metric. One of the limitations of CloudWatch is that the information it provides is region-specific, and customers must perform additional processing if they want to aggregate data across deployments in multiple regions.

We analyze Amazon CloudWatch's measuring metrics on key AWS components as follows. EC2 (Elastic Compute Cloud) is a Web service that allows customers to dynamically provide virtual machines with varying configurations on the Amazon cloud infrastructure. New instances can be created using templates or by creating a machine image and specifying its application, libraries, and settings. CloudWatch can measure the following metrics on EC2.

- CPUUtilization - The percentage of allocated EC2 compute units that are currently in use on the instance. This metric identifies the processing power required to run an application upon a selected instance. Units: Percent
- NetworkIn - The number of bytes received by the instance on all network interfaces. This metric identifies the volume of incoming network traffic to an application on a single instance. Units: Bytes
- NetworkOut - The number of bytes sent out by the instance on all network interfaces. This metric identifies the volume of outgoing network traffic to an application on a single instance. Units: Bytes
- DiskWriteOps - Completed write operations to all hard disks available to the instance. This metric identifies the rate at which an application writes to a hard disk. This can be used to determine the speed at which an application saves data to a hard disk. Units: Count
- DiskReadBytes - Bytes read from all disks available to the instance. This metric is used to determine the volume of data the application reads from the hard disk of the instance. This can be used to determine the speed of the application for the customer. Units: Bytes
- DiskReadOps - Completed read operations from all disks available to the instances. This metric identifies the rate at which an application reads a disk. This can be

used to determine the speed at which an application reads data from a hard disk. Units: Count

- DiskWriteBytes - Bytes written to all disks available to the instance. This metric is used to determine the volume of data the application writes onto the hard disk of the instance. This can be used to determine the speed of the application for the customer. Units: Bytes

Amazon ELB (Elastic Load Balancing) is an additional service that runs on top of EC2 and allows the customer to dynamically route traffic to various compatible EC2 instances to distribute the traffic and processing load. CloudWatch can measure the following metrics on ELB.

- Latency: The amount of time between a request and the corresponding response as seen by the load balancer. Units: Seconds. Valid Statistics: Minimum, Maximum, Average, and Count
- RequestCount: The number of requests processed by the LoadBalancer. Units: Count/Second. Valid Statistics: Minimum, Maximum, and Sum
- HealthyHostCount: The number of healthy instances. Both Load Balancing dimensions, LoadBalancerName and AvailabilityZone, should be specified when retrieving HealthyHostCount. Units: Count. Valid Statistics: the Average for a LoadBalancer within an Availability Zone
- UnHealthyHostCount: The number of unhealthy instances. Both Load Balancing dimensions, LoadBalancerName and AvailabilityZone, should be specified when retrieving UnHealthyHostCount. Units: Count. Valid Statistics: the Average for a LoadBalancer within an Availability Zone

Amazon RDS (Relational Database Service) is another Web service which allows the customer to dynamically create and scale database instances. CloudWatch can measure the following metrics on RDS.

- CPUUtilization - The percentage of CPU utilization
- FreeStorageSpace - The amount of available storage space
- DatabaseConnections - The number of database connections in use
- ReadIOPS and WriteIOPS - Measure the average number of read or write disk I/O operations per second, respectively
- ReadLatency and WriteLatency - Measure the average amount of time taken per read or write disk I/O operation, respectively
- ReadThroughput and WriteThroughput - Measure the average number of bytes read from or written to disk per second, respectively

B. Discussion

Amazon CloudWatch provides a significant set of data on resource usage and server performance for consumption by cloud service consumers. In this case, the responsibility boundary has shifted in favor of the customer. The metrics

and monitoring framework provided by the CloudWatch API are useful and available to the customer either for free or for a nominal cost. Many of these metrics would likely be collected by Amazon anyway for billing purposes. In fact, CloudWatch simply provides another service to customers and another opportunity for Amazon to monetize their existing monitoring infrastructure by having clients pay to subscribe to certain types of data. In addition, it mitigates the fact that the customers do not have complete unencumbered monitoring control over their systems by providing them with essential information to help them monitor and adjust deployments.

In May 2011, Amazon started providing custom metrics, a service that enables AWS users to select the application metrics for monitoring and retrieve statistics for reporting. This configuration can provide users with more customized services, but the selection is limited to the pre-existing performance metrics. Users are not yet allowed to define new metrics.

Therefore, CloudWatch cannot completely ameliorate concerns regarding the level of audit or security that may be required or simply prudent to implement in order to satisfy the requirements for cloud auditing. For instance, if we were to try and answer the questions from a PCI DSS audit with current CloudWatch data, few, if any, of those questions would be answered by the data provided in the CloudWatch service. The primary use of this type of data is to help customers evaluate the resource usage and performance of their applications and systems so that they can possibly troubleshoot the performance issues that are showing up on the system's front end. An example of such a use would be an organization with an application running on several systems hosted by the Amazon EC2 service. If they begin to get reports of slow application response time during a specific period, they could compare the metrics gathered through the CloudWatch service during that period with the previous baseline data and use this as a starting point in looking for performance anomalies.

In 2009, Amazon announced that they had undergone the Type II SAS70 audit standard. Although this was a significant milestone for their business, the security community is looking for more details about the controls behind the auditing results for the actual assessment because Amazon was reluctant to provide these details. As we described in Section III, SAS70 is an auditing methodology rather than an actual auditing rule, which means the organization being audited should specify a series of control objectives and corresponding activities. Therefore, Amazon's reluctance to disclose the details about the controls may create a false sense of control capability. Today, most cloud providers in the online transaction industry undergo the SAS70 audit standard. Other standards, such as ISO 27002, CoBIT, and PCI DSS, would be the next adoptions for cloud providers.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have discussed the security and privacy concerns in cloud computing and the current status of cloud auditing efforts. Based on the findings from our research, we

have addressed strategies for reliable cloud auditing and analyzed the deficiencies of current approaches. Furthermore, we have summarized the findings from our case study of Amazon CloudWatch, one of the most developed cloud-monitoring APIs. Finally, as an example of rapid response, we have discussed how to use the cloud auditing results for an advanced access control service that provides context-aware and dynamic access control decisions in the clouds. There can be many other ways to use the results of near-real-time cloud auditing as the inputs to rapid response in order to improve the overall security capability of clouds [27, 28, 29].

We conclude that current auditing efforts by cloud providers still fail to address some of the most pressing concerns of their customers due to multiple issues. We believe that our research outcomes will offer an array of robust cloud auditing and access control for assured clouds. We expect that the trusted, large-scale, resource-sharing services in the assured clouds will increase the organization's productivity and ability to accomplish its mission while allowing servers and providers to manage their resources securely and efficiently.

ACKNOWLEDGMENT

This research was performed in part while Dr. Joon S. Park held a National Research Council (NRC) Research Associateship (RAP) Award at the U.S. Air Force Research Laboratory (AFRL), Rome, New York, USA. This material is based upon work supported by the Air Force Office of Scientific Research under LRIR 11RI01COR and the U.S. Air Force Research Laboratory (AFRL) in-house Job Order Number GGIHZORR and GAIHCYBR.

REFERENCES

- [1] Sören Bleikertz, Matthias Schunter, Christian W. Probst, Dimitrios Pendarakis, and Konrad Eriksson. 2010. Security audits of multi-tier virtual infrastructures in public infrastructure clouds. In *Proceedings of the 2010 ACM workshop on Cloud computing security workshop (CCSW '10)*. ACM, New York, NY, USA, 93-102.
- [2] Kretzschmar, M.; Hanigk, S.; "Security management interoperability challenges for Collaborative Clouds," *Systems and Virtualization Management (SVM), 2010 4th International DMTF Academic Alliance Workshop on*, vol., no., pp.43-49, 25-29 Oct. 2010
- [3] Cloud Computing. NIST (National Institute of Standards and Technology) Computer Security Division. <http://csrc.nist.gov/groups/SNS/cloud-computing/> retrieved in 2010.
- [4] Amazon Elastic Computing Cloud (Amazon EC2). <http://aws.amazon.com/ec2/>
- [5] Google App Engine. <http://code.google.com/appengine/>
- [6] Jerry Robinson and Joon S. Park. Security mechanisms for trusted cloud computing. *Cloud Computing & Virtualization*, Singapore, May 17 – 18, 2010.
- [7] Jerry Robinson and Joon S. Park. Towards trusted cloud computing. *iConference*, University of Illinois at Urbana-Champaign, Illinois, February 3 – 6, 2010.
- [8] Grobauer, B.; Walloschek, T.; Stocker, E., "Understanding Cloud Computing Vulnerabilities," *Security & Privacy, IEEE*, vol.9, no.2, pp.50-57, March-April 2011
- [9] Chen et al. What's new about cloud computing security. Report No. UCB/EECS-2010-5, 2010.
- [10] Hassan Takabi, James B. D. Joshi, and Gail-Joon Ahn. 2010. Security and Privacy Challenges in Cloud Computing Environments. *IEEE Security and Privacy* 8, 6 (November 2010), 24-31.
- [11] Kaufman. Data security in the world of cloud computing. *IEEE Security & Privacy*, vol. 7 (4) pp. 61 - 64, 2009.
- [12] Siani Pearson and Azzedine Benameur. 2010. Privacy, Security and Trust Issues Arising from Cloud Computing. In *Proceedings of the 2010 IEEE Second International Conference on Cloud Computing Technology and Science (CLOUDCOM '10)*. IEEE Computer Society, Washington, DC, USA, 693-702.
- [13] Wayne A. Jansen, "Cloud Hooks: Security and Privacy Issues in Cloud Computing," Hawaii International Conference on System Sciences, pp. 1-10, 2011 44th Hawaii International Conference on System Sciences, 2011.
- [14] CloudAudit. A6 - The Automated Audit, Assertion, Assessment, and Assurance API. <http://www.cloudaudit.org/> retrieved in 2011.
- [15] Cloud Security Alliance (CSA). <http://cloudsecurityalliance.org/> retrieved in 2011.
- [16] Cloud Computing. NIST (National Institute of Standards and Technology) Computer Security Division. <http://csrc.nist.gov/groups/SNS/cloud-computing/> retrieved in 2011.
- [17] Cloud Computing Information Assurance Framework. ENISA (The European Network and Information Security Agency), November 2009.
- [18] Peter Tseronis. Cloud Computing Overview: A Federal Government and Agency Perspective. US Department of Energy, August 2009.
- [19] Shahid N. Shah. Cloud Computing by Government Agencies: Meeting the Business and Security Challenges in the Cloud. IBM developerWorks, August 2010.
- [20] Tom Greenfield. Cloud Computing in a Military Context - Beyond the Hype. DISA (Defense Information Systems Agency), 2009.
- [21] Wang et al. Toward publicly auditable secure cloud data storage services. Network, 2010.
- [22] Payment Card Industry. Payment card industry data security standards: Requirements and security assessment procedures, version 2.0. Accessed Electronically December 2010, URL: https://www.pcisecuritystandards.org/documents/pci_dss_v2.pdf.
- [23] Hugo Harber. Comment: Pci dss compliance in the cloud. Accessed Electronically, December 2010. URL: <http://www.infosecurity-magazine.com/view/9381/comment-pci-dss-compliance-in-the-cloud/>, May 2010.
- [24] Hemma Prafullchandra, Ken Owens, Tom McAndrew, Charu Chaubal, Davi Ottenheimer, Cuong Tran, and Han Yang. Cloud-based reference architecture for the payment card industry data security standard (PCI-DSS) 2.0. Technical report, HyTrust, 2010.
- [25] Securing Private/Public Clouds. http://virtualgeek.typepad.com/virtual_geek/2010/08/rs-a-vmware-and-intel-securing-privatepublic-clouds.html retrieved in 2011.
- [26] Amazon. "amazon cloudwatch". Accessed Electronically December 2010, URL: <http://aws.amazon.com/cloudwatch/>.
- [27] Park, J. S., An, G. and Liu, I. Y. Active access control (AAC) with fine-granularity and scalability. *Security and Communication Networks*, 4(10), 2011.
- [28] Joon S. Park, Ravi Sandhu, and Gail-Joon Ahn. Role-based access control on the Web. *ACM Transactions on Information and System Security (TISSEC)*, 4(1):37–71, 2001.
- [29] Joon S. Park, Gaeil An, and Deepak Chandra. Trusted P2P computing environments with role-based access control (RBAC). *IET (The Institution of Engineering and Technology, formerly IEE) Information Security*, 1(1):27-35, March 2007.