Toward Risk Assessment as a Service in Cloud Environments

Burton S. Kaliski Jr. and Wayne Pauley
EMC Corporation, Hopkinton, MA, USA

Abstract

Security and privacy assessments are considered a best practice for evaluating a system or application for potential risks and exposures. Cloud computing introduces several characteristics that challenge the effectiveness of current assessment approaches. In particular, the on-demand, automated, multi-tenant nature of cloud computing is at odds with the static, human process-oriented nature of the systems for which typical assessments were designed. This paper describes these challenges and recommends addressing them by introducing *risk assessment as a service*.

1. Introduction

Risk assessments have become a standard practice by which organizations determine and demonstrate their privacy, security, and compliance with other policies that protect against loss. Typical risk assessments include five steps: system characterization, threat assessment, vulnerability analysis, impact analysis, and risk determination [4]. Assessment-based approaches to evaluating risk have been the source of continuous research and debate in the literature ranging from theories about the value of quantitative versus qualitative analysis to whether internal staff are more effective for assessing the system than external staff [35][12].

Organizations use the results of a risk assessment in deciding how to apply scarce resources to protect their most important assets [4].

1.1 Security Assessments

Security assessment (SA) methods include international standards such as ISO/IEC 27002:2005 [16], NIST's SP800-53 [27] and "best practice" documents that are developed by security organizations, such as CERT's OCTAVE [28] method. Typical security assessments such as ISO/IEC 27002:2005 are qualitative in nature, relying on a survey instrument that can be used by internal or external security staff. ISO/IEC 27002 also includes a process for certifying the SA auditor.

In addition to a qualitative survey, OCTAVE also has a weighted scoring system that provides a quantitative comparison for threat prioritization and asset valuation. OCTAVE was designed to be executed by internal staff based on the assumption that the context of internal members would offer a better understanding of the value of the assets, and that the staff would gain subject matter knowledge through executing the assessment. In this sense, OCTAVE and other SA's provide the enterprise "security knowledge" by presenting a framework for understanding and evaluating security risk exposures for internal and external infrastructure resources [34].

1.2 External Audits and Assessments

In addition to general standards, assessment methods have also been designed for the needs of specific organizations or communities. These assessments typically address specific regulations for protection of data types such as personally identifiable information (PII), payment account information (e.g. PCI-DSS), personal health records, or corporate financial information. A domain-specific assessment is typically executed either by the body that created it or by a certified authority educated and licensed to deliver and audit.

Relevant examples include service provider assessments such as the SAS 70 Type II audit by the AICPA [1] and the Shared Assessments AUP [30]. Both are designed for assessing financial systems of service providers. The Shared Assessment AUP specifically includes security and privacy controls and was designed by security professionals [33] for the financial services industry.

1.3 Privacy Assessments

Privacy assessments (PA) are a type of risk assessment focusing on unauthorized disclosure. They differ greatly by country and the type information they are designed to assess [5]. A standardized PA is provided by ISO/IEC 22307:2008. Unlike the more general corresponding SA standard, PA 22307 applies only to financial services assets [15].

Comprehensive SA's typically include some PA requirements. For example ISO/IEC 27002 covers restrictions for screening employees, on-line transactions, audit logging, data protection practices, and the requirement for, and definition of, a privacy policy.

1.4 Autonomic Assessments

Although traditional assessments have typically been a manual process, there are more recent efforts to automate some of all of the process. Examples include context-aware privacy protection mechanisms in wireless networks and smart homes [22] [23]. Automated vulnerability assessment has also been studied [13].

2. Why Assess the Cloud?

Since the inception of cloud computing circa 2000, security, information leakage, and loss – or more generally, "trust" -- has been listed as a top concern [24] [11]. Trust in eCommerce is a broadly studied area with research results that demonstrate that improved trust [25] can positively affect eCommerce. In particular, an increased level of trust improves disclosure, reduces the demand for legislation, and reduces perceived risk. An enterprise that uses a self-service cloud provider is effectively consuming an eCommerce based service that provides infrastructure services, so it is reasonable to assume that similar perceptions of trust would apply to self-service cloud providers.

Privacy statements, security policies, and assessments have been established as effective methods for establishing trust in eCommerce service providers. Security and privacy protections have also been shown to engender trust and loyalty to eCommerce sites [10]. Cloud service providers are beginning to adopt elements of assessments such as a published security or privacy policy, or third-party audit pass/fail information. An example is Amazon Web Services (AWS)'s recent proclamation that it had passed a SAS 70 Type II Audit [2].

A cloud assessment may consider one or more of the entities in the cloud environment [21], including the cloud service provider, the cloud consumer, and ultimately, the cloud consumer's end customers.

Several efforts are underway to standardize cloud security, including the Cloud Security Alliance (CSA) [7], European Network and Information Security Agency (ENISA) [9], Cloud Audit (A6) [6], and Open Cloud Computing Interface (OCCI) [29]. These efforts provide requirements against which entities can evaluate security and privacy. However, the CSA and ENISA efforts do not address how such assessments per se would be implemented as an automated service in a cloud environment. They also leave open the question of how a cloud consumer would build a test and development environment that includes security regression testing as well as assessment controls [9]. A6 holds promise as a standard that will automate risk assessment as one of its primary goals.

3. Why Cloud Security and Privacy Are Hard to Assess with Existing Tools

The very characteristics that make cloud computing attractive also tend to make it hard to assess. Indeed, each of the five cloud characteristics articulated in NIST's definitions [21] can be seen to complicate the assessment of security and privacy of a business application deployed into a cloud computing environment. In this sense, cloud security and privacy are "immeasurable" with current assessment approaches.

3.1 On-Demand Self-Service

The avoidance of human interaction in the cloud, while reducing cost and decreasing deployment time, also takes away an important control point. Whereas an organization with a conventional IT environment depends on trained individuals to configure or verify security and privacy controls, a cloud consumer must rely substantially on automated enforcement mechanisms. A traditional assessment, however, may assume the existence of trained individuals in certain roles. To be effective in a cloud environment, it must equally address the increasing presence of their automated equivalents.

One of the promises of cloud computing is that a well governed, automated control mechanism will be more effective than a manual one. From the point of view of an assessment, the "training" of each of the automated entities that monitor and manage security and privacy must be proved. Put another way: Would the control mechanism pass the relevant portion of a CISSP exam? Would a human operator, assisted by such a collection of tools, be more reliable with them than without?

3.2 Broad Network Access

Broad network access affects assessments by changing the attack surface that must be assessed from a relatively static set of approved devices to a dynamic collection of end points of varying security postures and capabilities. To the extent that a traditional assessment inventories and evaluates device types and endpoints, it must also be updated to address an ongoing new set of computing elements (a trend that is affecting security and privacy independent of the cloud).

3.3 Resource Pooling

Resource pooling imposes perhaps the greatest collective set of challenges.

First, the dynamic allocation of resources according to consumer demand means that the specific resources deployed for a given application are not known a priori and therefore cannot be assessed, per se, in advance. An assessment must therefore focus on the correctness of the allocation mechanisms and the qualities of the overall pool. In a conventional IT environment, the parallel might be the inclusion of IT vendors, their inventories, and the company's procurement office in the assessment (which is not unusual in high-assurance evaluations [16] but certainly an additional aspect beyond the assessment of resources already in place).

Virtualization introduces similar concerns due to the separation of the logical entities being assessed from the underlying physical resources. In a conventional IT environment, however, the various resources, physical and virtual, are under the same governance. The difference in the cloud is that the logical entities are subject to consumers' IT requirements, whereas the underlying physical resources are the responsibility of the provider. An assessment must therefore include a translation between the governance of the two domains.

Second, the service of multiple consumers with the same pool of resources means that the impact of the presence of other tenants in the cloud infrastructure must also be taken into account. In a conventional IT environment, the presence of other organizations' actors is generally something to be avoided. In a multi-tenant cloud, the situation is completely reversed. A traditional assessment must therefore be updated to consider the effect of having other tenants present on the same resources before, during, or after the target consumer's use. The cloud provider's multi-tenancy architecture thus becomes a new and critical point of evaluation.

Multi-tenancy of data at a given provider (or even in a conventional IT environment) is of course a well established practice, and assessments of such deployments already address access controls and (external) intrusion protection. In the more general case of compute based multi-tenancy, the threats active (hence the term "actor" above), even to the point in infrastructure-as-a-service [21] deployments of an actor having full control of a (virtualized) IT environment. Accordingly, neighboring "content" is more at risk of contamination, or at least compromise, from the content in nearby "containers" [18][30]. An assessment must therefore contemplate the effectiveness of isolation mechanisms much stronger than those required for resource sharing among applications under common governance. It must also consider compliance mechanisms that verify that the appropriate separations have been achieved.

Finally, location independence of the physical resources introduces the complicating possibility that those resources may be subject to varying local regulations. Privacy is perhaps most affected by this aspect because of the significant diversity of relevant laws [20] [7] [9]. The dynamic nature of resource allocation combined with local variability of external requirements again means that an assessment is not possible based only on an a priori model of the IT environment, but must also consider the policy by which resources are assigned at run time. Indeed, all of the concerns about resource pooling and assignment collectively make the case for policy-based resource management in the cloud, a point more fully explored in [26].

3.4 Rapid Elasticity

The ability "to quickly scale out [and] scale in" at one level simply exacerbates the challenges already described. At another more architectural level, the expression of that ability brings a new set of issues to be bear through the practice of cloud bursting to handle the rapidly increasing workloads, migrating between different clouds to meet demand. In effect, the assessment must not only cover the consumer and a given target provider, but the provider's own sub-providers, and so on recursively.

Recursive implementation of an IT service is again not new; this is the basic model for a service-oriented architecture. The new part in the cloud context is the systematic migration of a consumer's computational workload across multiple providers not specified in advance; again, the movement of actors, not data. An assessment therefore becomes an exercise in evaluating the "transitive closure" of the environments in which those actors may operate not just each one individually.

3.5 Measured Service

Lastly, the "metering capability" by which cloud systems "automatically control and optimize resource use" presents one more challenge for assessments – namely, that the information collected to achieve the capability itself is a potential point of vulnerability.

Although metering information, again, is available in conventional IT environments, the assessment in a cloud environment must consider the much finer level of detail resulting from the focus on cost and dynamic resource sharing. Furthermore, even if the metering information for each tenant is individually well protected, there remains the possibility that an adversarial consumer can infer behavioral patterns of other tenants by analyzing its own usage. The extent of such disclosures, once again, must be factored into the assessment of security and privacy in the cloud.

4. Further Challenges from an Economic Perspective

In a conventional outsourcing model, an IT service provider supports a number of consumers with a dedicated set of resources for each one. (This is effectively the same as a collection of externally run private clouds.) Contractual arrangements can be fairly long term and consumer turnover is relatively low.

The higher churn in a cloud environment combined with the drive for volume creates an incentive for a provider to pay more attention to visible features such as performance than to insurance against eventualities. (Indeed, this is precisely the reason why security and privacy must be more measurable and visible in real-time, not before or after the fact.) Such a provider may therefore trade off immediate reward for eventual risk.

Although the same influences may be present to an extent in conventional IT environments, whether internal or outsourced, cost pressures in the cloud can significantly increase the motivations to make tradeoffs. Furthermore, consumer turnover can significantly decrease the insight of the decisions being made – unless, of course, they can be assessed more automatically.

A more extrinsic concern is the likely ongoing occurrence of mergers and acquisitions within the cloud provider ecosystem as the industry matures. This means that the surviving cloud providers will be in a constant state of forced hybridization. Consistent security and private assessment across such a constantly changing infrastructure may be nearly impossible.

Furthermore, as systems from acquired companies are integrated (many of them start-ups), old problems will be exacerbated such as the presence of rogue system administrators who can become a very real threat on a much larger scale. New problems such as data aggregation across merged companies with trans-border protections will also need to be addressed. Conventional IT, again, has the same issues but the pace of change and the scale of the systems will make due diligence hard to keep up with daily demands of business in the cloud.

5. How Do You Assess the Cloud?

It is clear from the foregoing discussion that the dynamic nature of the cloud makes traditional, more static assessments of resources and their configuration ineffective. Rather, just as the cloud is “on-demand,” increasingly, risk assessments applied to the cloud will need to be “on-demand” as well. Although the underlying *policy infrastructure* by which a cloud service provider (or consumer) applies resources to meet business objectives can and should be assessed – this foundation must be complemented with ongoing *policy compliance* that verifies that the objectives have indeed been met in operation.

We contend that the way that cloud computing should be assessed, is the same as the way cloud computing is delivered: *as a service*. Indeed, the same characteristics of the cloud that makes it hard to assess with existing tools, also make it easy to assess with new ones, especially the metering that is already built in for billing and service-level assurance.

Risk assessment as a service is a new paradigm for measuring risk as an autonomic method [19] that follows the on-demand, automated, multi-tenant architecture of the cloud – a way to get a continuous “risk score” of the cloud environment with respect to a given tenant, a specific application, or more generally, for use by new tenants and applications.

We envision such assessments as being made available in real-time by one or more of the entities in the cloud ecosystem. For instance, a cloud provider could perform continuous self-assessments as a best practice through evaluation of its own run-time environment; a trusted third party could assess the provider on an ongoing basis either through privileged access to certain internal measurement interfaces; or a consumer could assess the provider through non-privileged access. The third avenue is exemplified in approaches such as *proofs of retrievability* [17][3].

In each case, the dynamic assessment service would rest on a foundation periodic, underlying, static assessments. Static assessments should focus on the elements of the provider's underlying IT infrastructure and governance that (a) changes infrequently and (b) drives security and privacy in the dynamic environment. This again points to the importance of assessing security and privacy policies, policy enforcement mechanisms, and policy compliance mechanisms.

Since a provider may itself be a consumer of services from other providers, it is reasonable to expect that a provider would also be assessing the providers it relies on, thus addressing the point above about the recursive nature of cloud computing. Indeed, even if the ultimate business consumers and their customers are not directly assessing providers, the providers themselves will likely be assessing one another.

The addition of real-time assessment capabilities into the cloud environments parallels managed security services whereby an external provider monitors the internal security of a conventional data center. The results of such services are kept confidential to the relevant organization. In the cloud, the comparable results would, like the cloud itself, be open to all consumers.

An assessment service for the cloud involves more than just the automation of traditional surveys and scoring systems. The metrics must also be adapted to the nature of cloud computing, for instance the dynamic allocation of resources and multi-tenancy. Updating a traditional assessment to address cloud characteristics, then applying it manually, although accurate in principle, still may not fit the dynamic nature of the new environment. Rather, the "new wine" of cloud risk assessment should be put into the "new wineskin" of a cloud service.

6. Conclusion and Further Work

Risk assessments provide significant value in increasing trust in a commercial service, and thus appear particularly beneficial to the adoption of cloud computing. However, traditional assessments developed for conventional IT environments do not readily fit the dynamic nature of the cloud. We have proposed a cloud-based *assessment as a service* paradigm as a promising alternative.

We have not implemented such a service but rather offer it as a paradigm to be pursued. For actual implementation, we suggest several research directions briefly here. By definition, autonomic systems have to be reactive and proactive, e.g. an autonomic system must have the ability to measure its environment and then adjust its behavior based on goals and the current context [14]. This type of functionality requires sensors and an autonomic manager that analyzes risks and implements changes. For risk assessment, research is needed on the sensors that would collect relevant data in real time in a cloud environment. The measurements must support the viewpoints of multiple tenants and for service providers, so may be different than previous approaches focusing on a single stakeholder. Automated measurement and analysis is the basis for delivering risk assessment as a service; automated adjustment is a further (and much more complex) extension. (To start with, adjustments could be made in the conventional way based on the risk reports.)

Automated SLAs require a dictionary, an SLA specification language, and a correlation engine [31]. Risk assessment as a service could apply the same principles where the dictionary holds the risk assessment rules and asset valuations based on data entered by the tenant. This would provide the basis for a weighted scoring method such as those that are in OCTAVE [35].

CloudAudit has begun defining a directory/namespace for security audit and assessment that includes PCI DSS, HIPAA, COBIT, ISO 27002, and NIST SP800-53. Such a directory/namespace offers a common language that both tenants and service provider can use to collect information in support of continuous assessment. Further research may be needed on how to express the rules in a cloud setting.

For example, a procedure could be defined that uses dictionary definitions such as those defined by Shared Assessments [32] threats: malicious, natural, accidental, and business changes (scale/volume); assets: information, technology (VM's, storage, network), and pricing model (economic denial of service); vulnerability: threat-asset combination; control determination: preventive, corrective, predictive; determination: accept, avoid, mitigate; response: action, report. Output from the procedure would be fed to the autonomic manager to determine current state, perform pattern analysis against historical state data, correlate to the namespace, and perhaps trigger protective measures.

Another question to explore is the classic "Who guards the guards". Given that much of the proposed measurement relies on provider's own representations of its state, how does one address the possibility that such claims could be unreliable. Finally, from a formal point of view, it would be helpful not only how to implement risk assessment as a service, but whether indeed such a service is more effective than traditional approaches in practice.

References

[1] AICPA. SAS 70. Accessed March 2010. http://infotech.aicpa.org/Resources/Assurance+Services/Standards/SAS+No.+70+Service+Organizations.htm

[2] Amazon Web Services. AWS Completes SAS70 type II Audit. November 2009. http://aws.amazon.com/about-aws/whats-new/2009/11/11/aws-completes-sas70-type-ii-audit/

[3] Cachin, C., Keidar, I., and Shraer , A. Trusting the cloud. *ACM SIGACT News,* 20:4 (2009), pp. 81-86.

[4] Chen, T. Information and Risk Management. In M. Pagani (ed.), *Encyclopedia of Multimedia Technology and Networking, Volume II,* 2009.

[5] Clarke, R. Privacy impact assessment: Its origins and development. *Computer Law & Security Review, 25* (2009), pp.123-135.

[6] Cloud Audit. *The Automated Audit, Assertion, Assessment, and Assurance API.* Accessed March 2010. http://www.cloudaudit.org/

[7] Cloud Security Alliance. *CSA Guide V2.* Accessed March 2010. http://cloudsecurityalliance.org/

[8] Doherty, N. F., and Fulford, H. Aligning the information system policy with the strategic information systems plan. *Computers & Security* (2006), pp. 55-63.

[9] European Network and Information Security Agency. *Cloud Computing Information Assurance Framework.* Accessed March 2010. http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-information-assurance-framework/?searchterm=cloud

[10] Flavián, C. and Guinalíu, M. Consumer trust, perceived security and privacy policy: Three basic elements of loyalty to a web site. *Industrial Management + Data Systems,* 106, 5 (2006), pp. 601-620.

[11] Foster, I., Zhao, Y., Raicu, I. and Lu, S. Cloud computing and grid computing 360-degree compared. In *Proceedings of the IEEE Grid Computing Environments* (2008), pp. 1-10.

[12] Futcher, L., and von Solms, R. Guidelines for secure software development. In *Proceedings of the South African Institute for Computer Scientists and Information Technologists* (2008), pp. 56-65.

[13] F. Guo, Y. Yu, and T. Chiueh, Automated and safe vulnerability assessment. In *Proceedings of the 21st Annual Computer Security Applications Conference* (2005), pp. 24-34.

[14] Huebscher, M. C., and McCann, J. A. A survey of autonomic computing – degrees, models, and applications. *ACM Computer Survey,* (2008), 40(3), pp. 7-28.

[15] ISO/IEC 22307:2008. *Financial Services – Privacy impact assessment.* http://www.iso.org/iso/catalogue_detail.htm?csnumber=40897

[16] ISO/IEC 27002:2005. *Information technology – Security techniques – Code of practice for information security management.* http://www.iso.org/iso/catalogue_detail?csnumber=50297

[17] Juels , A. and Kaliski, B.S., Jr. PORs: Proofs of retrievability for large files. In *Proceedings of the 14th ACM Conference on Computer and Communications Security* (2007), pp. 584-597.

[18] Kaliski, B. Multi-tenant cloud computing: From cruise liners to container ships. In *Third Asia-Pacific Trusted Infrastructure Technologies Conference* (2008), p. 4.

[19] Kephart, J. O., and Chess, D. M. The vision of autonomic computing. *Computer* (2003), pp. 41-50.

[20] Mather, T., Kumaraswamy, S., and Latif, S. *Cloud Security and Privacy* (2009), O'Reilly.

[21] Mell, P., and Grance, T. *The NIST Definition of Cloud Computing. Version 15.* NIST, October 7, 2009. http://csrc.nist.gov/groups/SNS/cloud-computing

[22] Mitseva, A., Imine, M., and Prasad, N. R. 2006. Context-aware privacy protection with profile management. In *Proceedings of the 4th International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots* (2006). pp. 53-62.

[23] Moncrieff, S., Venkatesh, S., and West, G. Dynamic privacy assessment in a smart house environment using multimodal sensing. *ACM Trans. Multimedia Computing. Communications and Applications.* 5:2 (2008), pp. 1-29.

[24] Mowbray, M. and Pearson, S. A client-based privacy manager for cloud computing. In *Proceedings of*

the 4th International ICST Conference on Communication System Software and Middleware (2009), pp. 1-8.

[25] Nemati, H. R., and Van Dyke, T. Do privacy statements really work: The effect of privacy statements and fair information practices on trust and perceived risk in eCommerce. *International Journal of Information Security and Privacy,* 3:1 (2009), pp. 45-65.

[26] Nick, J.M., Cohen, D., and Kaliski, B.S, Jr. Key enabling technologies for virtual private clouds. In B. Furht and A. Escalante (eds.), *Handbook of Cloud Computing*, Springer, to appear.

[27] NIST SP800-53. *Recommended Security Controls for Federal Information Systems and Organizations.* http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/800-53-rev3_final-markup_final-publicdraft-to-final-updt.pdf

[28] *OCTAVE: Operationally Critical Threat, Asset, and Vulnerability Evaluation.* Accessed March 2010. http://www.cert.org/octave/

[29] *Open Cloud Computing Interface.* OCCI Working Group. Accessed March 2010. http://www.occi-wg.org/doku.php

[30] Ristenpart, T., Tromer, E., Shacham, H., and Savage. S. Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM Conference on Computer and Communications Security* (2009), pp. 199-212.

[31] Sahai, A., Machiraju, V., Sayal, M., van Moorsel, A, and Casati, F. Automated SLA monitoring for web services. In M. Feridun et al. (eds.), *Lecture Notes in Computer Science,* 2506, Springer (2002), pp. 28-41.

[32] *Shared Assessments: About.* Accessed March 2010. http://www.sharedassessments.org/about/

[33] *Shared Assessments: Frequently Asked Questions.* Accessed March 2010. http://www.sharedassessments.org/media/pdf-70vsAUPFAQ.pdf

[34] Tsoumas, B., Dritsas, S., and Gritzalis, D. An ontology-based approach to information systems security management. In V. Gorodetsky, I. Kotenko, and V. Skormin (eds.), *Lecture Notes in Computer Science,* 3685, Springer (2005), pp. 151-164.

[35] Vorster, A., and Labuschagne, L. A framework for comparing different information security risk analysis methodologies. In *Proceedings of the South African Institute for Computer Scientists and Information Technologists,* (2005), pp. 95-103.