



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Proposal for standard Cloud Computing Security SLAs – Key Metrics for Safeguarding Confidential Data in the Cloud

GIAC (GLEG) Gold Certification

Author: Michael Hoehl, mmhoehl@gmail.com
Advisor: Manuel Humberto Santander Pelaez

Accepted: TBD

Abstract

Many organizations are adopting cloud solutions as part of their portfolio of business applications. There are many documented advantages to using cloud computing including lower cost of entry, lower maintenance, improved availability, and global access. The good news is that many of the service organizations offer credible Service Level Agreements (SLA) for cloud availability. The bad news is that the same level of SLA maturity and standardization is not as pervasive for confidentiality and integrity. This is a serious concern when confidential and sensitive data is being hosted in the cloud. This paper proposes service level standards for security controls that customers can ask of service organizations offering cloud solutions.

Promises may fit the friends, but non-performance will turn them into enemies. ~Benjamin Franklin

1. Introduction

Cloud computing services provide many technology and business opportunities that were simply unavailable a few years ago. Small and medium businesses now have access to services that had previously been available exclusively to large cap or Fortune 1000-size companies. Large organizations have access to business and specialty applications that in the past had to be developed in-house. Typically, cloud technology frameworks are designed with redundant components and locations to assure a single failure does not interrupt service. The multi-tenant design of cloud computing services brings economy of mechanism and ultimately a lower total cost of ownership to provider and customer. Clearly, cloud computing has opened many doors for organizations to grow revenue, reach new markets, improve productivity, and expand product offering. These same doorways open to a large untrusted network when cloud services are connected to the Internet. “The chief concern firms have with the cloud is the overall security of their data. Putting information into services that are accessible over the public Internet means that criminals have a potential gold mine of targets” (V3, 2014). Unfortunately, customers might discover these security risks only after a data breach.

Cloud customer data protection requirements vary broadly. In some cases, the customer has a protection requirement of only basic logical access controls. Examples include email and research papers. In other cases, the data might be intellectual property, regulated data, or classified as confidential (e.g., Customer Relationship Management, Enterprise Resource Planning, credit card processing, etc.). These require advanced security controls including encryption, data masking, logging, auditing, jurisdiction, redundancy, integrity checking, and others. The cloud provider must strike a balance. A vendor does not want to implement security controls appropriate for the Department of Defense if the business data safeguard requirements do not demand this level of investment. If a disproportionate cost of infrastructure and services are put in place, the vendor might find themselves priced well above competitors and lose business. If sufficient safeguards are not in place, the vendor’s direct liability grows, and customers might not be willing to accept this level of risk to their data. When cloud computing service organizations are storing, processing, or transporting critical customer data, then security goals must be contractually and specifically described in a manner that aligns with the customer Information Security Management System (ISMS). Otherwise, the ISMS is undermined when cloud computing becomes part of the customer Information Technology service catalog. Alignment provides the customer the ability to continue to apply necessary data governance and risk decision-making practices even though the operations duties have been assigned to a vendor. The most common approach to communicate security service requirements is with the use of legal instruments and Service Level Agreements (SLA).

The Service Level Agreement provides a collection of quantifiable and measurable standards that the customer feels is necessary for the effective use of vendor services. These standards are a key component of the service contract to assure that business objectives are met in a manner that aligns with

Author: Michael Hoehl, mmhoehl@gmail.com

customer quality standards and risk appetite. SLAs are a critical element in any successful commercial relationship and define the metrics used to evaluate vendor performance, how the metrics are to be determined, metric targets, and frequency of metric review. “The use of Service Level Agreements for security services has the potential to provide some very tangible benefits to an enterprise. The largest benefits are associated with improved security administration and management practices. The definition of Service Level Agreements forces an organization to think about security” (Henning 1999). The value of an SLA is ultimately risk management alignment between vendor and customer. The remedies (i.e., liabilities and financial penalties) distribute risk between the contracted parties.

Unfortunately, contractual commitments to specific security controls with quantitative performance objectives are not common with cloud computing. Vendor service guarantees are typically focused on availability and lack clear definition for integrity and confidentiality commitments. Furthermore, standards of security performance are not well-established and benchmark data is limited. As cloud computing expands, greater security control visibility and accountability will be demanded by customers. This document explores Security SLA standards and proposes key metrics for customers to consider when investigating cloud solutions for business applications.

NOTE: This document is not intended to provide legal advice. No implied endorsement is intended for vendors mentioned in this document. The purpose of this document is general public education; it is not a substitute for legal or other professional advice. Do not rely exclusively on this document for guidance on contract terms or service agreements. Consult appropriate legal counsel for questions regarding business obligations and risk management for your organization.

2. SLA Primer

2.1. Service Charters, Service Guarantees, and Service Level Agreements

Information Technology Infrastructure Library (ITIL) Service Design defines a service as a means of delivering value to customers by facilitating outcomes customers want to achieve without the ownership of specific costs or risks (ITIL v3 Service Design 2012). According to ITIL, a service is a coherent, ready-to-use deliverable that is of value to the customer. Services allow customers to do business without worrying about underlying technology or IT infrastructure. Service definition and performance metrics are key to Service Management. Service Charters, Service Guarantees, and Service Level Agreements are all instruments that provide vendor service definition clarity and quality measurement. A Service Charter specifies the standards that the vendor intends to sustain as well as indicating the quality of services the customer is entitled to expect. Vendors are not required to compensate customer(s) if service objectives are not met. The Service Charter simply sets common understanding of good professional practice. A Service Guarantee is distinct from Service Charters because of the vendor payment to customer if obligations are not sustained. However, Service Guarantees are unilateral from the vendor (one size fits all) and does not allow for customer input. Service Level Agreements are bilateral and negotiated between customer and vendor. SLAs also have contractual implications if not met. These include financial penalties, regulatory

Author: Michael Hoehl, mmhoehl@gmail.com

compliance violations, additional audit obligations, and early severance of services. When the cloud solution is uniformly provided to all customers, then Service Guarantees are common. However, when a customer wants to closely align the service with unique business requirements, a Service Level Agreement is then appropriate (CIPS, 2009).

2.2. SLA Metric Best Practices Overview

According to Dimitri McKay, “Best practices for security don't change just because you're moving from the old data center model to a cloud-based environment. And it doesn't change who's responsible (ultimately, you), but it does affect who's in control and how close you are to the processes and technologies that are put in place. This is, of course, where the SLA comes in” (McKay, 2010). There are multiple sources for guidance while developing Service Level Agreements (e.g., Gartner, ITIL, etc.). This paper is not intended to provide all necessary guidance for developing and negotiating SLAs in general. This guidance is already well-documented and published. For example, the Cloud Standards Customer Council has published a 10-step approach with “Practical Guide to Cloud Service Level Agreements” that provides a prescriptive guide with requirements and best practice recommendations for each step of the Cloud SLA lifecycle (CSCC, 2011). The CSCC document includes valuable guidance when comparing cloud providers and negotiating terms. The following steps are proposed:

- Understand roles and responsibilities
- Evaluate business-level policies
- Understand services and deployment model differences
- Identify critical performance objectives
- Evaluate security and privacy requirements
- Identify service management requirements
- Prepare for service failure management
- Understand the disaster recovery plan
- Define an effective management processes
- Understand the exit process

The purpose of the aforementioned CSCC paper is to provide broad guidance with Cloud SLAs. This paper advances the CSCC recommendation to use Cloud SLAs to meet critical performance objectives relating to security.

An SLA metric is constructed from a few common elements. These include name of metric, metric source, duration of sampling, frequency of sampling, scope of testing, target range, weight, reporting process, and penalty/incentive calculation. These elements should be found with every SLA metric used to demonstrate services have been sustained to mutually agreed obligations. If some of these elements are not clearly defined, then the service provider operations team might not be clear on what is to be accomplished and how to report service delivery performance. Further, false reading might occur. For example, if the frequency of sampling is every 5 minutes, then a recurring outage lasting only 2-3 minutes might not be

identified and reported as an SLA failure. Not all metrics are the same as far as importance to customer (and vendor). A weight assigned to each metric provides a clear message of the priority and relative importance of a metric to customer. Weight is therefore related to customer risk and influences calculation of vendor penalties or incentives (transfer of risk).

When choosing performance metrics, there is a broadly accepted approach abbreviated as SMART (SANS, 2007). Each letter of SMART has a specific meaning (“Specific”, “Measurable”, “Actionable”, “Relevant”, and “Timely”). “Specific” requires that metrics be of a certain scope and targeted. “Measurable” requires that the data be of a quantitative nature that is easily verified and complete. “Actionable” requires that metrics clearly reveal the corrective action that needs to occur. “Relevant” requires that the metrics all have contextual value—that they are meaningful to the audience. “Timely” metrics require data be collected in a credible and repeatable manner that can be trended. This guidance applies to Service Level Agreement metrics, too. Brian Monahan and Mike Yearworth from Hewlett-Packard Laboratories offer further guidance with selecting SLA metrics (Monahan, 2008):

- Express SLAs in terms of observable, measurable quantities, so that there is a clear statement of what kind of data has to be gathered on a regular basis by each participant.
- State SLAs in terms of outputs of relevant processes – for example, “the number of transactions of this or that performed over the reporting period” - and not the inputs – for example, “the number of servers of this type that are available for transaction processing”. Otherwise, the SLA can all too easily end up saying that some “particular state” was intended and some effort was made to achieve this “particular state”. This is still a valid statement, of course, but probably not the one that was intended. This illustrates an interesting point – merely making a valid statement is not enough to make a useful and meaningful SLA, since perfectly valid statements can also be completely irrelevant and ineffective.
- Avoid “implicit functions” that have been left hanging and undefined. This is a very dangerous source of ambiguity, since a lot of complexity could arise hidden away inside them. In other words, it is important to clearly define the lexicon of general terms available to express the SLA and, where necessary, to define specialized terms as they are introduced.

For SLAs to be successful, metrics must be defined for effective decision-making and with economy of mechanism. Metrics that are unclear, inconsistently derived, irrelevant to vendor service scope of work, expensive to produce repeatedly, and lacking clear call to action can have a negative impact on quality control. More is not necessarily better for metrics. Complexity is the enemy of effective SLA management. Too many metrics bundled indiscriminately together can result in significant cost and false readings (positive and negative). Analysis paralysis might occur when clear patterns among all the unrelated data points are not revealed. Generation of every metric has a cost. This cost includes the vendor computational effort and the customer analysis effort. For SLAs, a balance between level of metric granularity and level of effort to repeatedly create the metric is necessary.

Author: Michael Hoehl, mmhoehl@gmail.com

Service Level Agreement metrics must reflect factors within the vendor's control and scope of service. If the service output has a dependency on customer input, there is no clear accountability. An example of a poor metric would be password resets required within 4 hours when the customer requires management approval. In this case, the service provider could be delayed because of customer lack of performance (delayed approval by manager). A better metric would be password resets required within 4 hours after management approval. For this example, the customer might elect to track the length of time for management approval and combine this with the vendor metric to understand actual the customer experience. After observing the typical length of time to fulfill the password reset (including manager authorization time delay and vendor execution time delay), the customer might elect to change their approach to password resets (e.g., eliminate manager approval and incorporate challenge questions so vendor can identify individual and confirm request is authentic).

In some cases, the SLA metrics are reported in a consolidated manner reflecting delivery performance for all tenants of the service provider. This is a common approach for Service Guarantees. Metric consolidation creates a risk for an individual tenant whose actual experience does not reflect that of the entire customer base. Therefore, consolidation of metrics across multiple customers should be carefully considered before agreeing to this SLA metric reporting approach. Metric averages and summarization might also create decision conflict. Customers should carefully consider how the sample data is being related to avoid a high impact, low frequency event that occurs at a critical business time (e.g., On-line Retailer suffering outage during Cyber Monday) with little or no penalty.

3. Business Considerations

3.1. Sky is the limit

According to Skyhigh Networks, “The cloud has created a new wave of enterprise software that is not only faster to develop, easier to deploy, and more cost effective, but also offers innovative features not found elsewhere. That’s because much of the innovation today is happening in software delivered via the cloud, and for many customers, the cloud is mainstream” (Skyhigh Networks, 2014). An interesting finding was that companies are selecting cloud based solutions because their evaluation has identified the cloud application as the best-in-class enterprise application—not just best-in-class cloud-based application. Customer Relationship Management (CRM), Expense Reporting, Office Automation, Productivity Management, and even Enterprise Resource Planning (ERP) are now pervasive in the Cloud. This creates a condition in which Security or IT departments broadly prohibiting cloud computing might also be preventing their business from using the application with best business functionality and value. This could result in competitive and operational disadvantages for the business. Further, the business application within the cloud computing framework might offer significant security advantages including automated security management, system life cycle management, yearly security control auditing, and DR/BCP. It is becoming increasingly difficult to make compelling arguments to avoid cloud computing as a rule of

thumb. Risk and quality remain appropriate concerns. However, quality and security controls in the Cloud can be effectively addressed using Service Level Agreements.

3.2. Clouds bring risk

The European Network and Information Security Agency (ENISA) performed a Cloud Risk Assessment study (ENISA, 2012). In this study, top security benefits as well as risks were identified. The most important classes of cloud-specific risks were:

1. Loss of governance – gap in security controls because of customer abstraction and poor SLAs;
2. Lock-in – data portability and vendor migration prohibitively difficult;
3. Isolation Failure – failure of mechanisms to enforce separation of tenant resources;
4. Management interface compromise – unauthorized privileged access gained over Internet;
5. Data Protection – Inadequate safeguards (e.g., access control, encryption, etc.) to safeguard confidentiality, integrity, and availability of data;
6. Insecure or incomplete data deletion – adequate or timely destruction of data in an irrecoverable manner not possible;
7. Malicious insider – vendor (and vendor agent) unauthorized administrator access to customer data;
8. Customer's security expectations – contracted and actual security controls do not match or cost cutting results in control elimination;
9. Availability Chain – Poor Internet access reliability at customer location creates single point of failure.

The Cloud Risk Assessment conducted in 2009 and 2012 by ENISA reveal substantially the same risks. Therefore, it can be reasonably concluded that these risks were not temporary and will remain into the future. The baseline for security measurement associated with Service Level Agreements must address these risks. Regulatory and compliance obligations as they apply to each customer (and service provider) are also a significant risk management consideration.

3.3. Asset Valuation

Which security controls are relevant and necessary for cloud computing? This is generally determined by asset value and Cloud model. In this case, asset value is a business measurement—not simply technology bill of materials. Valuation is more than acquisition cost of computer equipment. Asset value is also influenced by business process importance, operational dependency, and data confidentiality. Examples of key business process include customer service, supply chain, health service delivery, and financial management. Manual execution of business processes might not be plausible for an extended period of time. Automation might be required creating an operational dependency on technology. If cloud computing services are used to manage customer data, patient records, trade secrets, vendor contracts, or supply-chain then confidentiality safeguards are much stricter than that of public data. In addition to mandated regulatory and compliance obligations (e.g., EU Protection Directive 95/46/CE, HIPAA, PCI, etc.), many organizations elect to safeguard data as if regulated because of privacy policy commitments to

consumers and harm to reputation if data is accessed unauthorized. When a formal asset classification and valuation standard is not in place, the Cloud Security Alliance proposes the following helpful questions (CSA, 2011):

1. How would we be harmed if the asset became widely public and widely distributed?
2. How would we be harmed if an employee of our cloud provider accessed the asset?
3. How would we be harmed if the process or function were manipulated by an outsider?
4. How would we be harmed if the process or function failed to provide expected results?
5. How would we be harmed if the information/data were unexpectedly changed?
6. How would we be harmed if the asset were unavailable for a period of time?

3.4. Cloud Service and Deployment Models

Once asset valuation is complete, customers must consider what Cloud Service and Deployment Models are available. Understanding the cloud model is critical to understanding the complete risk picture. There is significant different trade-offs to each model in terms of features, complexity, and security. NIST SP800-145 proposes the cloud is essentially composed of three Cloud Service Models and four Cloud Deployment Models (Mell & Grance, 2011). The three Cloud Service Models are Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). An overview of each service model is provided below in Table 1.

Table 1: NIST SP800-145 definition of Cloud Service Models

Cloud Service Model	Description
Infrastructure-as-a-Service (IaaS)	The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).
Platform-as-a-Service (PaaS)	The capability provided to the consumer is to deploy onto the cloud infrastructure consumer - created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.
Software-as-a-Service (SaaS)	The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-

Author: Michael Hoehl, mmhoehl@gmail.com

specific application configuration settings.
--

Essentially, IaaS is intended to provide basic computer infrastructure in a virtual environment so that the consumer does not have to purchase assets. PaaS sits on top of IaaS and is intended to provide developers an environment to host their applications without having to manage the underlying technology complexity (e.g., databases, messaging, queuing, operating system, etc.). SaaS is built on PaaS and intended to provide “on-demand” business software that is available over the Internet. Each Cloud service model builds upon the underlying service model.

For IaaS, the customer typically assumes substantially all data and application security risks. When progressing from IaaS to PaaS to SaaS, more technology abstraction is introduced reducing the customer direct visibility into and control over the environment. Therefore responsibility associated with Data and Application Security controls must be transferred to the vendor since the customer does not operate nor directly manage the entire environment. Security controls, data governance, compliance, liability, incident handling, software life cycle management, and performance level expectations must be contractually stipulated, managed, and enforced. Though responsibility can be transferred to vendor using risk management instruments like MSAs and SLAs, accountability typically remains with the Cloud customer since they provide the data. This is an important operational risk management consideration when negotiating Cloud models. Ultimately, organizations can assign security responsibility, but they cannot assign security accountability. The Cloud Security Alliance visually depicts the relationship between Cloud Service Models, security controls, and compliance in Figure 1 below.

Figure 1: Cloud Security Alliance – Mapping the Cloud Model to the Security Control & Compliance

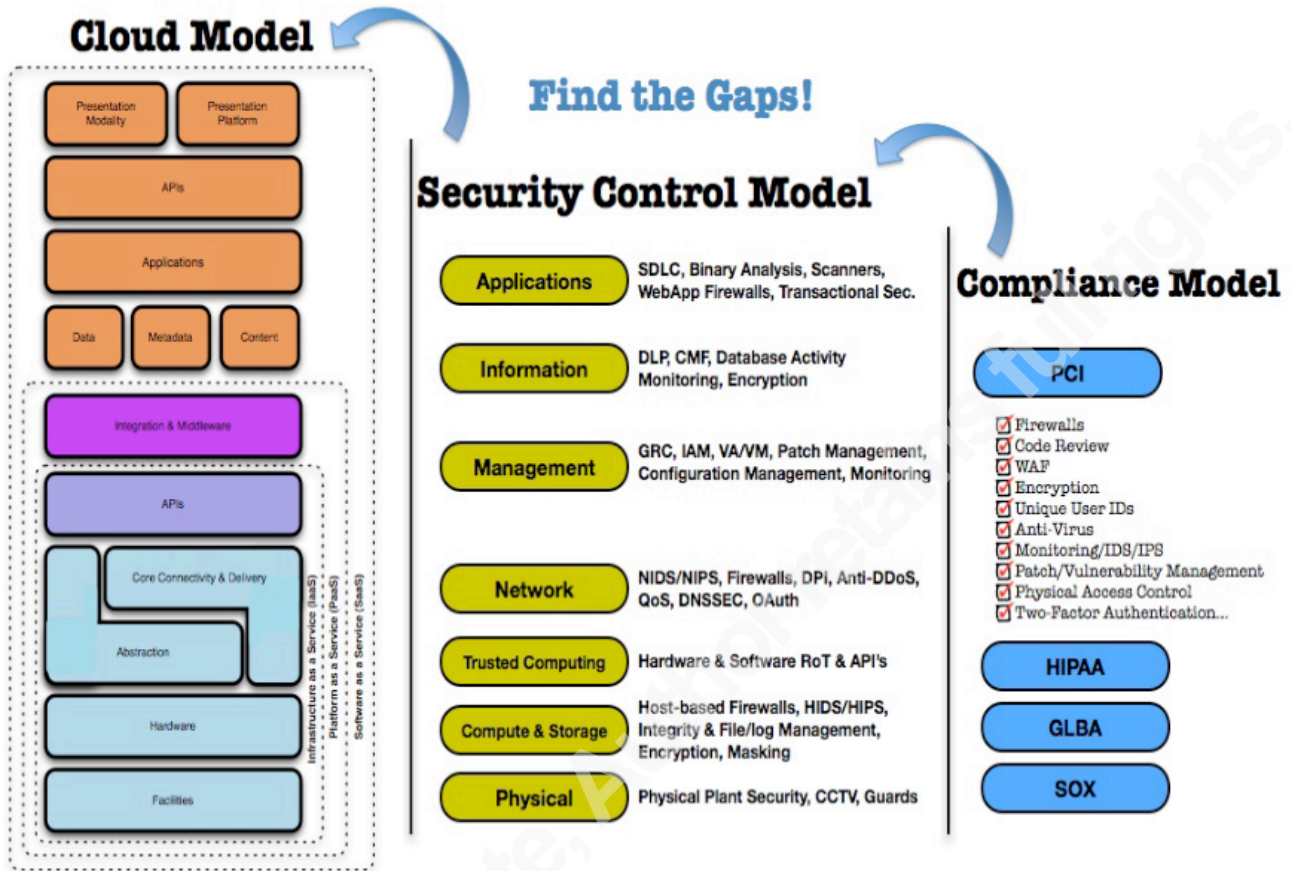


Figure from CSA Security Guidance for Critical Areas of Focus in Cloud Computing

In addition to Cloud Service Model, NIST SP800-145 proposes four different Cloud Deployment Models.

Table 2: NIST SP800-145 definition of Cloud Deployment Models

Cloud Deployment Model	Description
Private	The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.
Public	The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
Community	The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g. mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in

	the community, a third-party, or some combination of them, and it may exist on or off premises.
Hybrid	The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds)

Essentially, Private Cloud is traditional outsourcing because the assets are dedicated to a single customer tenant. This deployment model offers the isolation that a customer might desire when handling highly confidential data. Public Cloud is most familiar to commercial customers and retail consumers. With Public Clouds, the environment is revealed to an untrusted network (e.g., Internet) creating a large surface of attack for systems and applications. With Community Clouds, a smaller target consumer population is in scope as compared to Public. Hybrid Cloud bridges together one or more Public, Private, or Community Clouds.

Once the Asset Value and Cloud Models are well understood, customers can map the two together and make business decisions about which approach to take. These decisions typically result in a selection of a Cloud Service and Deployment that are expected to save time and money. This same mapping is used to determine who is responsible for each security control and associated SLAs necessary to reduce risk. With IaaS, customers are still substantially on the hook to build and operate security controls. However, when progressing to PaaS and eventually SaaS, security controls must be contracted with vendor using SLAs because the customer abstraction from the operations of the technology.

4. Legal Considerations for Security SLAs

Typically, the SLA is part of a Master Service Agreement (MSA). The MSA has many parts to it including Service Description, Confidentiality Requirements, Indemnification, Insurance Coverage, Business Continuity Commitments, Acceptable Use, Security and Privacy Policy and SLA. Metrics and objectives might be found within these other documents. These too should be carefully reviewed to determine if they provide sufficient risk mitigation and ensure there are not conflicting commitments (e.g., Security Policy commitments regarding patching conflicting with SLA commitments). In addition to providing measurable targets for service delivery, remedies are typically identified when SLA obligations are not sustained. As mentioned earlier, not all metrics are equal. Therefore, weights are assigned to metrics as an incentive to assure vendor services are sustained to predetermined commitment levels. In addition to weight, some metrics might have a clause that recurring failures within a finite timeline have multipliers. For example, a service level objective failure in month 1 is 5%, subsequent failure in month 2 is 10% and subsequent failure in month 3 is 20%. Some Cloud services represent complex business processes and require significant investments to integrate before going live (e.g., ERP, CRM, etc.).

Author: Michael Hoehl, mmhoehl@gmail.com

Recovering data and redeploying to another service provider (or internally) might require months or even years to execute. Therefore weights and multipliers for SLA metrics are very important for customers that are not prepared to uproot from a vendor with short notice.

Vendor services evolve and improve over time. This is especially true for Cloud Computing related services. Because of this condition, new service objectives (e.g., recurring web application penetration testing every 90 days) and improved service quality commitments (e.g., account creation in less than 4 hours as compared to next day) will be offered by the vendor. This should be considered when negotiating the MSA. If the vendor offers these new commitments subsequent to the signing of the MSA, then the MSA should be written in a way that the customer automatically inherits these new service commitments without renegotiation of contract and pricing.

Vendors frequently use business partners to help provide the total service offering. Rarely does a vendor use their own resources exclusively to provide services “top-to-bottom”. This pool of resources is commonly referred to as Cloud Federation. For example, the vendor might provide the business application (e.g., Software-as-a-Service Model), however the foundational infrastructure might be provided and maintained by a third-party business partner (e.g., Infrastructure-as-a-Service Model). Cloud Federations work by distributing services and operational risk to providers that specialize in specific services (e.g., data center hosting, database administrator services, virtualization management, application integration, etc.). This approach can be beneficial for the vendor to maximize efficiency, increase scalability, and reduce cost. These are substantially the same benefits endpoint Cloud customers enjoy. With Cloud Federations, key customer security controls might be under the responsibility of a third-party that the customer has no direct, legal relationship with. Risk propagation (and liability assignment) therefore must be clearly understood by Cloud customer. “Weakness assessment and vulnerability analysis must be abstracted, i.e., not based on specific system details, and made relevant to the external black boxed cloud domain...to prevent violation scenarios and thus ensure security control compliance.” (Hale and Gamble, 2012). Further, security certifications (e.g., ISO 27001, SSAE 16, PCI DSS, etc.) might be held by the third-party and not the vendor with which the customer has contracted. This creates a potential compliance issue as well as unintended liability. When negotiating Master Service Agreements and Service Level Objectives, the customer should have a clear understanding of the third-party involvement and risk propagation. There should be mutual understanding that SLAs apply to the vendor and vendor's agent(s).

Timing for negotiating the SLA is important. The recommendation to codify security controls in a formal contract was a good idea in the past when negotiating traditional data center outsourcing and still good advice today with modern Cloud services. Waiting until after the contract is signed to establish a Service Level Agreement severely disadvantages the customer and presents an opportunity for unplanned risk. Further, contracts that reference vendor internal documents and marketing material might result in a moving target. Commitments that existed in marketing materials or vendor standards at the time of negotiation (e.g., encryption and data destruction) might not be sustained. The SLA is a critical part of the

Author: Michael Hoehl, mmhoehl@gmail.com

customer-vendor relationship and should be formally established early within the contract—not informally as a website URL that might change without notice.

Cloud vendors provide a variety of ways to report on service level performance. In some cases, the reporting is similar to that of an outsourcing engagement in which the vendor presents reports monthly (pro-actively or upon customer request). This approach is beginning to fade away as the cost to produce these metrics can be significant. A growing trend is to present the evidence on-line for customer review. Evidence of commitment to security controls including Service Guarantees and SLAs that are common across multiple tenants are also being presented to Cloud security authorities that serve as informal certification and accreditation organizations. For example, the Cloud Security Alliance established Security, Trust, and Assurance Registry (CSA, 2014) in 2011. CSA STAR is a free, publicly accessible registry that documents the security controls provided by various cloud computing offerings. Cloud service providers have the option to complete the CSA STAR control assertion questionnaire to demonstrate due care and satisfy customer evidentiary requirements. The program is based on an open certification framework that begins with vendor assertion, progressing through third-party certification (like that of SSAE 16 or PCI), to finally continuous monitoring-based certification.

Lastly, formal audits of controls can be performed using the same approach as those used for data centers. For example, ISO 27001 registration, AICPA/ISAE SOC1/SOC2 report, and Payment Card Industry Data Security Standard (PCI DSS) Service Provider assessment are common considerations for Cloud vendors. In some cases these audit standards might be required to do business and will be a condition of the Master Services Agreement. Depending exclusively on an audit to confirm the effectiveness of the security controls once a year might be acceptable for some Cloud service offerings. Once a year assurance by an auditor might not be adequate for those customers intending to transport, store, or process confidential data. For these conditions, Service Level Agreements provide continuity of visibility into the operational effectiveness of controls and better risk response at or near the time a control fails.

5. Proposed Standards for Cloud Computing Security SLAs

5.1. Available Cloud-focused frameworks

Today there is no single cloud computing information assurance framework that meets the needs of every customer. There are essentially 2 categories of guidance available today for establishing Cloud SLA standards for security. The first source is existing standards for security that in the past were substantially applied to the more conventional data center model (e.g., ISO 2700x, ITIL, COBIT, PCI DSS, HIPAA, etc.). However, some of these standards might not be versatile enough because they are intended for specific communities (FedRAMP based on NIST SP800-53 for U.S. Government agencies), specific data types (FERPA for student education records or PCI DSS for credit card data), and specific locations (EU Data Protection Directive). The second source of guidance include standards intended specifically for cloud computing (e.g., Cloud Security Alliance Cloud Control Matrix, British Standards Institution

Author: Michael Hoehl, mmhoehl@gmail.com

Security Requirements for Cloud Computing Environments, European Network and Information Security Agency Procure Secure, Jericho Forum Self-Assessment, etc.). This paper considers both sources of guidance to establish a minimum baseline for Cloud SLA security standards.

5.2. Selection criteria for Security SLAs

The SLA standards proposed in this paper are focused on security controls for cloud computing. NIST SP800-53 proposes there are 18 families of security controls organized into 3 classes. Families are assigned to their respective classes based on the dominant characteristics of the controls in that family. The organization of these security controls can be helpful when determining how the Service Level Agreement conditions are to be presented in the Master Services Agreement. The table provided below lists the NIST security controls into technical, operational, and management classes.

Table 3: NIST SP800-53 Control 18 Families and 3 Classes

Technical	Operational	Management
(AC) Access Control	(AT) Awareness and Training	(CA) Certification, Accreditation and Security Assessment
(AU) Audit and Accountability	(CM) Configuration Management	(PL) Planning
(IA) Identification and Authentication	(CP) Contingency Planning	(RA) Risk Assessment
(SC) System and Communication Protection	(IR) Incident Response	(SA) System and Services Acquisition
	(MA) Maintenance	(PM) Program Management
	(MP) Media Protection	
	(PE) Physical and Environmental Protection (PS) Personnel Security	
	(SI) System and Information Integrity	

For the technical class, security controls are generally architectural or policy based. Generally, these controls would be defined as part of an Information Security Policy. An approach might be to have the Master Services Agreement state that the vendor must have an Information Security Policy that requires this class of controls including automatically disabling inactive accounts after a finite time (AC-02), encryption to protect the confidentiality of remote access sessions (AC-17), regularly review/analyze audit records for indications of inappropriate or unusual activity (AU-06), time stamps for use in audit record generation (AU-08), multifactor authentication (AI-02), function isolation (SC-03), and protects the confidentiality of transmitted information (SC-09). The examination of these controls is typically performed by a third-party security practitioner engaged by the cloud vendor yearly, and the audit approach

Author: Michael Hoehl, mmhoehl@gmail.com

can be based on a credible standard (e.g., ISO27001, SOC 2, etc.). The absence of these technical controls might be sufficient cause for contract termination. For this reason, annexing an Information Security Policy might be more effective than creating a unique service level metric for each of the security controls in this family.

The management class of security controls is similar to the technical control family in that the obligations might be stated in the Master Service Agreement without complex metrics. Management security controls can be examined by a third-party auditor to confirm the controls are in-place and practiced. For example, the Master Service Agreement might state that a Risk Assessment must be performed yearly (RA-03), vulnerability scanning must be performed every month and immediately after every significant change (RA-05), and Acceptable Use Policy must be presented for review and signature yearly (PL-04). The absence of these management controls might be sufficient cause for cloud service provider contract termination. For this reason, annexing these requirements in the Master Service Agreement with a simple in-place obligation might be more effective than creating a unique service level metric for each of the security controls in this family.

The operational class of security controls generally require frequent and recurring monitoring to demonstrate commercially reasonable due care. Relying on management attestation or audit once a year might not sufficiently demonstrate operational effectiveness nor manage risk. Therefore these controls lend themselves to Service Level Objectives (SLO). Based on asset valuation and cloud model, the customer might have specific requirements including configuration change control (CM-03), DR/BCP (CP-02), incident handling (IR-04), monitoring physical access intrusion (PE-06), patching (SI-02), malware prevention (SI-03), intrusion detection (SI-04), and error handling (SI-11). Therefore, the NIST operational class of security controls provides a useful reference for determining the SLA metrics.

Since all industries and business requirements cannot be met with a single, universal collection of security control standards, service level obligations share the same limitations. This section proposes a minimum baseline that is intended to be broadly adoptable by all cloud service providers. Standards are categorized by Cloud Service Model (IaaS, PaaS, and SaaS). Since each cloud service model builds upon the underlying service model, the security SLA standards follow the same approach. PaaS proposed security SLA standards are intended to include IaaS. SaaS security SLA standards are intended to include IaaS and PaaS. Additional regulatory (e.g., Health Insurance Portability and Accountability Act), contracted (e.g., Payment Card Industry), and business requirements can supplement this baseline.

5.3. Key Metrics for IaaS SLAs

As stated earlier, Infrastructure-as-a-Service (IaaS) is one of three cloud service delivery models. IaaS is intended to provide basic computer infrastructure in a virtual environment so that the consumer does not have to purchase assets. The customer typically assumes substantially all data and application security risks. When progressing from IaaS to PaaS to SaaS, more technology abstraction is introduced reducing the customer direct visibility into and control over the environment. There are some infrastructure components such as networking that the customer will not have access to, but are critical for the security

program. These require service level obligations for the vendor since the customer has no “hands-on” access to configure or examine the associated security controls. This section proposes key security SLA standards for these components common to IaaS as they serve an important role in the information security program. The authoritative sources that are identified and align with the proposed metrics are NIST SP800-53r4 (NIST), Cloud Security Alliance Cloud Control Matrix v3.01 (CSA), and ISO 27001-2013 (ISO). References to the specific section of each authoritative source are provided with each SLA recommendation for additional guidance.

Table 4: Key Security Service Level Agreement Metrics for IaaS

#	Key Security SLAs	NIST	CSA	ISO
1	Change Control and Configuration Management	CM	CCC	A12.1.2
2	Data Center Asset Management	CM	DCS	A8.1.1
3	Disaster Recovery and Business Continuity Planning	CP	BCR	A.17.1.3
4	Secure Configuration and Server Hardening	CM	IVS	A.12.5.1
5	Malware and Intrusion Prevention	SI	TVM	A.12.2.1
6	Network Vulnerability and Penetration Testing	RA	IVS	A.14.2.3
7	Software Lifecycle and Patch Management	SA	TVM	A.12.6.1
8	Security Incident Handling	IR	SEF	A.16
9	Secure Network Protocols and Data Transport	SC	IPY	A.13
10	Security Event Logging	AU	IVS	A.12.4

5.4. Key Metrics for PaaS SLAs

This section proposes key security SLA standards for PaaS. These are in addition to the aforementioned IaaS SLA standards. Several new SLA metrics for PaaS are proposed (e.g., secure application and program interfaces). Some SLA standards are listed again reflecting the scope change between IaaS and PaaS. For example, Change Control and Configuration Management not only applies to infrastructure, but also middleware, databases, and messaging components introduced as part of PaaS.

There are more “moving parts” to PaaS as compared to IaaS. To deliver the PaaS, multiple vendors might be collaborating creating a cloud federation. For example, separate vendors may be engaged for data center, network, systems, database and middleware services. Therefore, some clarification might be required during the SLA negotiation to understand actual metric source and reporting accountability.

Table 5: Key Security Service Level Agreement Metrics for PaaS

#	Key Security SLAs	NIST	CSA	ISO
---	-------------------	------	-----	-----

1	Change Control and Configuration Management	CM	CCC	A.12.1.2
2	Secure Application and Program Interfaces	SC	AIS	A.14.1.3
3	Disaster Recovery and Business Continuity Planning	CP	BCR	A.17.1.3
4	Secure Configuration	CM	IVS	A.12.5.1
5	Intrusion Prevention	SI	TVM	A.14.1.2
6	Vulnerability and Penetration Testing	RA	IVS	A.14.2.3
7	Software Lifecycle and Patch Management	SA	TVM	A.12.6.1
8	Data Protection/Portability/Retention/Destruction	MP	DSI	A.8
9	Encryption and Key Management	SC	EKM	A.10.1.2
10	Application and Database Logging	AU	IVS	A.12.4

5.5. Key Metrics for SaaS SLAs

This section proposes key security SLA standards for SaaS. The IaaS and PaaS key Security SLAs are cumulative and would apply to a SaaS environment. In addition to this cumulative approach, it is worth noting that some of the same key security metrics are listed in all three cloud service models (e.g., Disaster Recovery, Intrusion Prevention, Software Lifecycle and Patch Management, etc.). These metrics remain relevant because the security requirements and operations duties to fulfill these requirements are substantially different for each cloud service model. Intrusion kill chain analysis demonstrates that cyber-attacks have pivoted using infrastructure, platform, and software to gain unauthorized access to confidential data (USSCCST, 2014). Each cloud service level introduces new surfaces of attack. Intrusion detection and prevention mechanisms are different for each cloud service level because the threats grow. Therefore, the security SLA metrics remain in all three service model recommendations.

Table 6: Key Security Service Level Agreement Metrics for SaaS

#	Key Security SLAs	NIST	CSA	ISO
1	Change and Release Management	CM	CCC	A.12.1.2
2	Secure Application and Program Interfaces	SC	AIS	A.14.1.2
3	Disaster Recovery and Business Continuity Planning	CP	BCR	A.17.1.3
4	Secure Configuration	CM	IVS	A.12.5.1
5	Intrusion Prevention	SI	TVM	A.14.1.2
6	Vulnerability and Penetration Testing	RA	IVS	A.14.2.3
7	Software Lifecycle and Patch Management	SI	TVM	A.12.6.1

8	Secure Coding Practices	AT	HRS	A.14.2
9	Identity Access Management	AC	IAM	A.9.2

6. Conclusion

Prospect and existing customers of service providers are demanding confidentiality, integrity, and availability when contracting with vendors for cloud computing. Contracts and service level obligations that commit to only availability objectives are not adequate to manage business risk. Service Level Agreements provide a vehicle of communication between the vendor and customer regarding performance and quality expectations. SLAs are the de facto standard for managing IT Outsourcing, and remain relevant for cloud computing. Security has become increasingly important and one of the major reasons for delayed adoption of cloud computing. SLAs that include security metrics are needed to reduce risk and effectively transfer responsibility between parties. As customer abstraction from technology operations grows with cloud computing, SLAs become vital to assure security controls are properly sustained. Unfortunately, there is not a single, authoritative standard for cloud computing SLAs that applies to all customer security management needs. There are some promising standards in early development to improve Service Level Agreements including Cloud Security Alliance STAR (Security, Trust, and Assurance Registry) and European Commission SPECS (Secure Provisioning of Cloud Services) and ENISA (European Network and Information Security Agency) Procure Secure. Until an authoritative standard is widely adopted and becomes the benchmark, the approach proposed with this paper can be useful for advancing initial conversations between customer and cloud service provider regarding security SLAs.

7. References

Butler, Brandon. (2012). *Nine security controls to look for in cloud contracts*. NetworkWorld. Retrieved from <http://www.networkworld.com/article/2161443/cloud-computing/nine-security-controls-to-look-for-in-cloud-contracts.html>

Brodkin, J. (2008). *Gartner: Seven cloud-computing security risks*. Infoworld. Retrieved from <http://www.infoworld.com/article/2652198/security/gartner--seven-cloud-computing-security-risks.html>

Chartered Institute of Purchasing and Supply. (2009). *How to prepare Service Level Agreements*. Retrieved from http://www.cips.org/Documents/Resources/Knowledge%20How%20To/How%20to%20prepare%20Service%20Level%20Agreements.pdf?bcsi_scan_3F31264ACB0CFD71=hHhIS/

Author: Michael Hoehl, mmhoehl@gmail.com

- Cloud Security Alliance. (2011). *Security Guidance for Critical Areas of Focus in Cloud Computing Version 3.0*. Retrieved from <https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf>
- Cloud Security Alliance. (2014). *CSA Security, Trust and Assurance Registry*. Retrieved from https://cloudsecurityalliance.org/star/#_overview
- Cloud Standards Customer Council. (2012). *Practical Guide to Cloud Service Level Agreements Version 1.0*. Retrieved from http://www.cloud-council.org/2012_Practical_Guide_to_Cloud_SLAs.pdf
- U.S. Senate Committee on Commerce, Science, and Transportation. (2014). *A Kill Chain Analysis of the 2013 Target Data Breach*. Retrieved from http://www.commerce.senate.gov/public/?a=Files.Serve&File_id=24d3c229-4f2f-405d-b8db-a3a67f183883
- ENISA. (2012). *Benefits, risks, and recommendations for information security*. Retrieved from <http://www.enisa.europa.eu/events/speak/cloud.jpg/view>
- European Commission. (2014). *Cloud Service Level Agreement Standardisation Guidelines*. Retrieved from <http://ec.europa.eu/digital-agenda/en/news/cloud-service-level-agreement-standardisation-guidelines>
- Grance, Timothy and Jansen, Wayne J. (2011). *NIST SP800-144: Guidelines on Security and Privacy in Public Cloud Computing*. Retrieved from http://www.nist.gov/customcf/get_pdf.cfm?pub_id=909494
- Hale, Matthew & Gamble, Rose. (2012). *Risk Propagation of Security SLAs in the Cloud*. Retrieved from <http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=6477665&url=http%3A%2F%2Fieeexplore.ieee.org%2Fiel7%2F6470041%2F6477486%2F06477665.pdf%3Farnumber%3D6477665>
- Henning, Ronda. (1999). *Security Service Level Agreements: Quantifiable Security for the Enterprise?* Retrieved from <http://www.nspw.org/papers/1999/nspw1999-henning.pdf>
- Hogben, G. & Dekker, M. (2012). *Procure Secure: A guide to monitoring of security service levels in cloud contracts. Technical report, European Network and Information Security Agency (ENISA)*. Retrieved from <http://www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/procure-secure-a-guide-to-monitoring-of-security-service-levels-in-cloud-contracts>
- ISO. (2013). *Information Technology – Security techniques – Information security management systems – Requirements*. Retrieved from http://www.iso.org/iso/catalogue_detail?csnumber=54534
- Krutz, Ronald L. & Vines, Russell Dean. (2010). *Cloud Security: A Comprehensive Guide to Secure Cloud Computing*. Wiley Publishing:Indianapolis, IN.

- Kuyoro, S. O. et al. (2011). *Cloud Computing Security Issues and Challenges*. Retrieved from <http://www.cscjournals.org/manuscript/Journals/IJCN/volume3/Issue5/IJCN-176.pdf>
- Lambo, Taiye. (2012). *Why You Need a Cloud Rating Score*. Retrieved from https://cloudsecurityalliance.org/wp-content/uploads/2012/02/Taiye_Lambo_CloudScore.pdf
- Mather, Tim. Et al. (2009). *Cloud Security and Privacy - An Enterprise Perspective on Risks and Compliance*. O'Reilly Media.
- Mell, P., & Grance T. (2011). *NIST Special Publishing 800-145: The NIST Definition of Cloud Computing: Recommendations of the National Institute*. Retrieved from <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>
- McKay, Dimitri. (2010). *Cloud Service SLA Security Tips – What Should You Should Be Asking Your Provider?* Retrieved from <http://www.securityweek.com/cloud-service-sla-security-tips-what-should-you-be-asking-your-provider>
- Monahan, Brian & Yearworth, Mike. (2008). *Meaningful Security SLAs*. HP Labs. Retrieved from <http://www.hpl.hp.com/techreports/2005/HPL-2005-218R1.pdf?q=meaningful>
- Myerson, Judith M. (2013). *Best practices to develop SLAs for cloud computing*. Retrieved from <http://www.ibm.com/developerworks/cloud/library/cl-sla-standards/>
- NIST. (2013). *NIST Special Publication 800-53 Revision 4: Security and Privacy Controls for Federal Information Systems and Organizations*. Retrieved from <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf>
- PCI Security Standards Council. (2013). *Information Supplement: PCI DSS Cloud Computing Guidelines*. Retrieved from https://www.pcisecuritystandards.org/pdfs/PCI_DSS_v2_Cloud_Guidelines.pdf
- Peterson, Brad. (2014). *Ten Key Questions for Developing an Effective Service-Level Agreement*. Retrieved from <http://www.outsourcing-center.com/2014-07-ten-key-questions-for-developing-an-effective-service-level-agreement-63376.html>
- Popa, Reluca Ada. Et al. (2011) *Enabling Security in Cloud Storage SLAs with CloudProof*. Retrieved from <http://web.mit.edu/ralucap/www/cloudproof.pdf>
- SANS. (2007). *MGT 512: SANS Security Leadership Essentials for Managers with Knowledge Compression*.
- Skyhigh Networks. (2013). *Cloud Adoption and Risk Report*. Retrieved from <http://info.skyhighnetworks.com/rs/skyhighnetworks/images/2013%20Cloud%20Adoption%20%26%20Risk%20Report.pdf>

Author: Michael Hoehl, mmhoehl@gmail.com

Tarig, Muhammad Imran. (2012). *Towards Information Security Metrics Framework for Cloud Computing*. International Journal of Cloud Computing and Services Sciences. (IJ-CLOSER) pp 209-217.

V3. (2014). *Top 10 cloud computing risks and concerns*. Retrieved from <http://www.v3.co.uk/v3-uk/news/2343547/top-10-cloud-computing-risks-and-concerns/page/5>

Wrinkler, Vic. (2011). *Securing the Cloud: Cloud Computer Security Techniques and Tactics*. Syngress. Waltham, MA.