Public Cloud Service Agreements: What to Expect and What to Negotiate Version 2.0.1

August, 2016

Contents

Acknowledgements

The major contributors to Version 1.0 were **Claude Baudoin** (cébé IT & Knowledge Management), **Jordan Flynn** (eFortresses), **John McDonald** (CloudOne), **John Meegan** (IBM), **Michael Salsburg** (Unisys), and **Steven Woodward** (Cloud Perspectives).

Out of those, Claude Baudoin, John Meegan and Steven Woodward also participated in the writing of Version 2.0. They were joined by **Dr. Rizwan Ahmad** (Cianaa Technologies), **John Bruylant** (The Cloud Turbo), **Marcus Busby** (cébé IT & Knowledge Management), **Stephen Cushing** (Bendigo Adelaide Bank), **Mike Edwards** (IBM), **Rajesh Jaluka** (IBM), **Roberta Mazzoli** (Schlumberger), **Sanjay Mundergi** (Albertsons), **Arvind Radhakrishnen** (TATA Consultancy Services), **Karolyn Schalk** (IBM), **Prasad Siddabathuni** (Edifecs), **Rampal Singh** (HCL Technologies), and **Long Wang** (IBM).

What is New in Version 2.0

Version 1.0 of this white paper was published in 2013. In the interval, some cloud service providers have appeared, disappeared or merged; the language of the agreements has occasionally changed, perhaps even because of discussions with customers whose understanding of the issues had been heightened by our work; and our own knowledge of what customers need has been sharpened by our experience and by the addition of new co-authors.

Version 2.0 takes this maturation of the topic of service agreements into account. For example, Step 5 includes new considerations about data residency, the References section links to many more examples of service agreements than the earlier version, and several other updates were made throughout the document.

Version 2.0.1 contains a few minor editorial changes made after publishing version 2.0.

Executive Summary

As CIOs and CFOs search for efficient, agile and cost-effective ways to deliver business services to the enterprise, they naturally consider public cloud solutions. Cloud technology supports all types of IT capabilities, from basic computing and storage to platforms and applications. These cloud services can be orchestrated to deliver what is consumed by the enterprise – business services. If any portion of this orchestration does not meet service level objectives, the business can be impacted, from slow response time to debilitating outages and damage to the enterprise's reputation. Moreover, the broader adoption of hybrid cloud solutions requires management visibility across both in-house systems and public cloud services to ensure the availability and performance of critical services. Therefore, service agreements from cloud service providers need to be understood and balanced against the needs of the business.

CIOs who have already outsourced parts of their infrastructure understand the value of Service Level Agreements (SLAs), and will readily accept the need for formal Cloud Service Agreements (CSAs) and their associated SLAs. For organizations that are using a cloud service for the first time, CSAs may be totally new. IT managers who rely on computing resources that are located and managed outside their immediate control quickly realize that in order to ensure the level of service required by the business, they must understand their objectives and transform them into formalized service levels, agreed with the cloud service providers.

This paper provides cloud service customers with a pragmatic approach to understand and evaluate public CSAs. The recommendations are based on a thorough assessment of publicly available agreements from leading providers. In addition to this paper, a great deal of research and analysis regarding the landscape of CSAs is available in the CSCC's *Practical Guide to Cloud Service Agreements* [3].

In general, we have found that the current terms proposed by public cloud service providers fall short of the commitments that many businesses require. Of course, these providers have reputations to establish or maintain, therefore they are likely to employ all reasonable efforts to correct problems, restore performance, protect security, and so on. However, neither the specifics of the measures they take, nor the remedies they offer if they fall short, are currently expressed well enough in most of their standard formal agreements. Furthermore, the language about service levels is often distributed among several documents that do not follow a common industry-wide terminology. We hope that one impact of this paper will be to improve this state of affairs.

A development of interest in this area is the work currently underway to create an international standard for CSAs, ISO/IEC 19086 [8]. Once published, it should help provide a common vocabulary for use in CSAs and in their associated SLAs.

When specific examples are used in this paper, they reflect the state of the practice as of the date of this document – they can be neither permanent nor exhaustive. In addition, such examples are NOT intended to compare or recommend specific cloud service providers, but rather to provide illustrations and observations from a vendor-neutral perspective, leading to key considerations for evaluating a public CSA. Similar text will be found across multiple cloud service providers, and customers need to perform their own analysis of relevant agreements and other contractual expectations and obligations.

Current Anatomy of a Cloud Service Agreement

No standard nomenclature is used across the various public cloud service providers to define their CSAs (see references [12] through [64]). The CSA could itself be a part of a Master Service Agreement or called a Service Level Agreement, Business Continuity Policy or simply a service agreement. This section and the artifacts mentioned in it, offers a structure that cloud service customers can use to compare agreements from different public cloud service providers.

Customers are advised to pay great attention to the language used in the agreements. Not all agreements are written or edited with the care they require. Wording errors can radically alter the meaning of a clause, making it much more broadly applicable than intended. The right time to catch and correct these errors is before signing a contract, not when a dispute arises.

In general, the CSA can be decomposed into four major artifacts: "Customer Agreement," "Acceptable Use Policy, "Service Level Agreement," and "Privacy Policy." Bear in mind that these artifacts may change at different times, independently from each other.

Customer Agreement

Since business service management includes the processes and procedures of the cloud service provider, explicit definitions of the roles, responsibilities and execution of processes need to be formally agreed upon. The "Customer Agreement" fulfills this need, using various synonyms such as "Master Agreement," "Terms of Service," or simply "Agreement." In general, all the public cloud Customer Agreements we reviewed contained the following critical sections, each using slightly different terminology.

- *Use of Service Offerings*. This defines how the customer is expected to use the public cloud service. Alternate terminology includes "Terms of Use," "Provision of the Service" and "Services Description."
- *Fee and Payment*. This describes the methods of charging and paying for cloud services. Other terminology includes "Service Charges Schedule," "Purchasing Services," and "Payment Terms."
- *Temporary Suspension*. This describes a process whereby the provider suspends for a time the use of the cloud service by a specific customer, based on an issue such as abnormal use of the cloud service, security risks, or delinquency in payment. Other terminology might include "Suspension and Removals" and "Term, Termination and Suspension."
- *Terms and Termination*. This addresses the terms of the agreement and the process for termination. Other terminology includes "Agreement Termination and Closing the Account." As noted above, the provider may also specify in this section a temporary suspension clause.
- *Indemnification*. This addresses holding the provider harmless against various claims, damages and loss.
- *Disclaimer*. This section describes what is not included in the agreement. It is described under headings such as "Warranties and Disclaimer."

- *Limitation of Liability*. In the event of a problem, this section specifies a limit on the amount of compensation a customer can claim. (See Step 8 for further discussion of the impact of disclaimers and limitations of liability in the context of disaster recovery).

Acceptable Use Policies (AUPs)

By definition, an Acceptable Use Policy (AUP), sometimes called an Acceptable *Usage* Policy or Fair Use Policy, is a set of rules followed by users of a network, website, or service. It serves to stipulate constraints and guidelines that must be followed when using that resource.

All of the public cloud service providers we reviewed included acceptable use terms for both the cloud service provider and the cloud service customer:

- It is typical for the provider to restrict cloud service use for "unlawful, obscene, offensive or fraudulent content or activity," which includes security-related items such as "interfering with or violating the integrity or security of a network or system, evading filters, sending unsolicited, abusive or deceptive messages, viruses or harmful code."
- Conversely, the provider usually agrees not to violate the intellectual property rights of the customer.

In most cases, an AUP is provided as a separate artifact on its own web page. The AUP sometimes overlaps with, or replaces, the Security/Privacy terms of the Customer Agreement.

Penalties for violation of the AUP terms can be severe – including suspension or termination of the customer's use of the cloud service.

Cloud Service Level Agreements

Service Level Agreements (SLAs) are formal documents, agreed on by both parties that define a set of service level objectives. These objectives may concern availability, performance, security and compliance/privacy. However, the analyzed cloud SLAs focused solely on availability and on the remedies offered if the availability target is not met. This primary focus on availability objectives and little else is the norm across the three traditional cloud service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) [9].

Privacy Policies

Most public cloud service providers issue a separate privacy agreement or statement that highlights their commitments to maintaining the privacy of all collected data. However, we found several instances where security and privacy policies are discussed jointly.

The depth and breadth of privacy commitments vary significantly across providers. In general, the privacy policy describes the different types of information collected; how that information is used, disclosed, and shared; and how the provider protects that information. As discussed in Step 5, there is an issue of *whose* data is covered by this document – whether it is limited to the data about the cloud service customer, or extends to the personally identifiable information (PII) of which the customer is the custodian, but which belongs to third parties (e.g., the account holders for a bank, the patients for a hospital, etc.). The latter type of data, for which the cloud service customer is termed a PII Controller, is the subject of regulations and laws and is of significant concern for many cloud service customers.

What You Can Expect and What You Should Negotiate

The CSCC *Practical Guide to Cloud Service Level Agreements* white paper [3] prescribes a series of ten steps that cloud service customers should take to evaluate CSAs in order to compare public cloud service providers or negotiate terms with a provider. The following steps are discussed in detail:

1. Understand roles and responsibilities
2. Evaluate business level policies
3. Understand service and deployment model differences
4. Identify critical performance objectives
5. Evaluate security and privacy requirements
6. Identify service management requirements
7. Prepare for service failure management
8. Understand the disaster recovery plan
9. Develop an effective governance process
10. Understand the exit process

This section uses the same list of ten steps as a straightforward way to complement and extend the original Guide. For each step, the corresponding subsection describes the range of statements found in the CSAs that were reviewed, highlights best-of-breed statements, and provides recommendations for what customers should negotiate with providers. Example language from actual agreements is quoted to highlight key points. Assistance on where to find specific information is also provided for each step (i.e., which service agreement artifact should be examined – Customer Agreement, AUP, Cloud SLA, or Privacy Policy).

Step 1: Understand Roles and Responsibilities

The AUP is the primary artifact that should be thoroughly reviewed by cloud service customers to understand their responsibilities and those of the provider. AUPs are generally not related to technology or financial performance of the cloud service relationship, but rather govern the valid and invalid customer behaviors related to using the service.

There are typically differences in AUPs that can be expected based on the service model (IaaS, PaaS or SaaS). Some AUP terms, especially for SaaS services, tend to be superseded by a specific contract or agreement or are simply presented in such documents rather than in an explicit AUP.

Although the AUPs that were reviewed contained some common points, each was original to a surprising degree. Some providers focus more on the illegal usage of their services, such as inappropriate material or copyright violations, while others are more concerned with abuse of network bandwidth or overloading the service itself.

Hence, customers need to perform due diligence and exercise caution to ensure their proposed usage of the service does not violate the AUP – especially in case of abstract or ambiguous AUPs. Also, some of the providers' AUPs include clauses like "Please note that we may change our Acceptable Use Policy at any time, and pursuant to the Provider Terms, it is your responsibility to keep up-to-date with and adhere to the policies posted here."

Appendix A contains key observations and actual language examples for the most common aspects of public cloud AUPs.

Recommendations

When evaluating the **Acceptable Use Policy** of a public cloud service provider, customers should expect the following, and if needed should request clarification.

- *Clarity*. Since the terms of an AUP apply to the overall use of the services, and it is difficult to foresee every possible situation, it is important for the customer to clearly understand all aspects of the AUP. You should ask the vendor to clarify, in writing, any items for which there is confusion or open interpretation.
- *Brevity*. Most of the AUPs analyzed were succinct and clear. However, a few were filled with legal jargon and seemingly duplicate provisions from one part to another. Such lengthy, wordy provisions were probably never tested in a court of law, and you do not want to be the first customer to defend yourself against them.
- *Completeness*. While many AUPs covered all the provisions mentioned in the above "Anatomy" section (content, security, service integrity, and rights of others), some AUPs were missing certain provisions. For example, one large cloud service provider said absolutely nothing about the content prohibited on the service, instead relying on vague language that allowed them, in theory, to deem unacceptable anything they chose. This open language is not in the customer's best interest, because it places the burden of proof on the customer, and there is no clear language for a judge or jury to consider in deciding a case.
- *Focus*. Some AUPs define a very broad range of actions that the service provider may deem as unacceptable. Absent scope limitations, this might place the user in breach of contract for an action seemingly unrelated to the cloud service. Customers should shy away from such broad commitments, or ask for clarification in writing.

In summary, AUPs have little consistency in wording, although there is a clear pattern to the types of provisions they include. To safely navigate these waters, customers should exercise caution and thoroughly review every provision before agreeing to an AUP. It might be helpful for the customer to elaborate on their expected usage of the service and have that validated by appropriate parties on both ends.

Step 2: Evaluate Business Level Policies

Cloud service customers must consider matters of governance, risk compliance and business policy when reviewing a public CSA since there are interdependencies between the policies expressed in the agreement and the business strategy and policies followed in other aspects of the business. Organizations that have adopted hybrid cloud computing need to consider how to harmonize the policies of the multiple service providers they work with, as well as with the policies that apply to their in-house systems. For example, cooperation between providers when it comes to incident resolution or change notifications should not be taken lightly or assumed. Guidance specific to governance of hybrid cloud computing environments may be found in the CSCC's *Practical Guide to Hybrid Cloud Computing* [5].

Areas that are typically most relevant to business concerns are:

- Data policies – residency, storage, disposal and migration
- Change notification and change management (services, APIs, or agreements)
- Suspension of services
- Limitations of liability
- Intellectual Property

Data Policies

The data policies of a public cloud service provider are perhaps the most critical business-level policies to be evaluated. While these are most often expressed in the overall CSA, there may be provider policies included in the AUP or elsewhere that need to be included in a thorough review.

The obligation that a cloud provider has to its clients and their data is partly governed by the data protection legislation applicable to PII in the user's location, as well as the legislation for those locations in which data may reside or may be made available. Customers should carefully consider these legal requirements and how the CSA deals with issues such as movement of data to offer multisite redundancy across several geographies without violating applicable laws or regulation. For commercial information which is not PII, and therefore not covered by data protection legislation, the Customer Agreement should contain the appropriate language.

In general, all public cloud Customer Agreements reviewed contain the following clauses:

- The customer is solely responsible for the development, content, operation, maintenance, licensing and use of their content.
- The customer retains all rights, title, and interest in their content and data.
- The customer is responsible for its end users' use of their content and of the cloud service, and for their compliance with the terms of the Cloud Services Agreement.
- The customer is responsible for any individual's personal information (or any other confidential information) stored in the cloud. The customer agrees to comply with all applicable privacy and data protection laws, to obtain all necessary consents, and make all necessary disclosures

before including personal information in their content. This is a logical requirement – the provider cannot be held responsible for any potential violations of privacy laws by the customer.

The responsibility for maintaining appropriate security, protection and backup of the customer's data may be shared in a way that needs to be reviewed. In the IaaS model, the customer may be entirely responsible for this, unless an additional service is purchased from the provider at an extra cost. Even in PaaS and SaaS models, the provider may include such a clause in order to minimize their responsibility in case of a catastrophic loss of information. This needs to be carefully reviewed.

Early Customer Agreements did not allow the customer to specify where its content would be stored. As concerns about data residency surfaced, received publicity and got amplified by legal decisions such as the rejection of the "safe harbor" ruling between the European Union and the United States, this situation has changed. Increasingly, providers with an internationally distributed infrastructure allow customers to select where their data should – or should not – be permanently stored. This option is generally offered to government customers, but extends to commercial entities as well. It is a critical provision for customers in certain vertical industries (financial services, health care, oil and gas, etc.) on which authorities often impose stringent data residency obligations. Note that such storage location constraints should include the location of backup data, and may also need to extend to "in transit" data. This is further discussed under Step 5.

A cloud service provider may leverage a third party to store data (for example, a SaaS provider may rent storage from an IaaS provider), to perform data and content migration, or to manage incidents (e.g., call center). There is a need to ensure that the third party is also bound, through appropriate agreements, to protect the customer's data.

Finally, the cloud service provider must commit to notifying the customer in advance of any changes in policies or in systems that affect the way in which customer data and content are protected.

Law Enforcement Access

The Customer Agreement should explicitly state that the provider will not access the customer's content. However, it usually includes an exception in which the provider signifies that it will comply with properly formulated requests by law enforcement agencies. In the event of such valid legal or governmental requests, customers should require prompt notification from their provider, enabling them to file without delay for a restraining order if possible (some countries do not allow this), or at least to know that the data was accessed and notify their own users or owners of the data.

As has been shown in well-publicized lawsuits, *who* can issue a valid order to produce the data can be unclear. Therefore, the provider should state whether it will comply with a request based on the country where it is based, the country where the data is stored, the nationality of the customer, the nationality of the person whose data is being requested, etc.

When evaluating the **data policies contained in the Customer Agreement**, customers should consider the following best practices:

- Ensure that the agreement allows the customer to specify the physical location of their security-

sensitive content, or content subject to data residency requirements (acceptable locations vary across industries and national legislations).

- Ensure that cloud provider personnel will not access the customer's data, except when required by law and duly requested by law enforcement authorities.
- Under such circumstances, ensure that the agreement specifies that the cloud service provider will give prompt notice, allowing the customer an opportunity to file for a stay of the request, where permitted by law.
- Understand what capabilities the provider offers for redundancy, replication and backup of customer data, and what actions the customer needs to perform in order to make use of these capabilities.

Changes to Services, APIs or Agreements

Provisions for changes to services, APIs and agreements are typically included in the Customer Agreement, describing in detail the circumstances under which cloud service providers can make such changes. Customers must fully understand the impact that such changes may have on their data and business services, and should develop a plan to minimize business disruption.

In most cases, the onus is on the cloud service provider to give advance notice (typically 30 days) to their customers for any such material change. For services, providers usually give themselves the right to change, discontinue, or deprecate any service offering, or change or remove features or functionality of the service offering – at any time. For APIs, providers may change, discontinue or deprecate any APIs for the services from time to time, but will typically commit to apply commercially reasonable efforts to continue supporting the previous version of any API for a period of time (typically 12 months) after the change, discontinuation, or deprecation.

When evaluating the **policies concerning changes to services** contained in the Customer Agreement, customers should consider the following best practices:

- Ensure that the agreement specifies that advance notice (minimum of 30 days) will be given for all changes initiated by the cloud service provider.
- Ensure that the agreement commits the provider to use commercially reasonable efforts to maintain backward compatibility, or continue to operate the applicable service/API, for an extended period of time (minimum of 12 months) after the effective date of the change.
- Understand whether a change in services that might "break" a customer application is sufficient cause to terminate the agreement with the cloud service provider.

Suspension of Services

Customers must fully understand the impact that potential suspension of services might have on their data and business services, and on their own clients, and should develop a plan to ensure business continuity in such an event. A suspension of services clause is typically part of the Customer Agreement and describes the circumstances under which the cloud service provider can suspend services to a customer. Reasons for suspension will typically include:

- Breach of contract, including payment delinquency and violation of the AUP

- Behavior posing a security risk to the service or to any third party
- Actions that may subject the cloud service provider to liability
- Usage that represents a direct or indirect threat to the provider's network function or integrity, or to anyone else's use of the service

In most cases, suspension of service is applied to the minimum necessary portion of the service and will only be in effect for as long as reasonably necessary to address the issues giving rise to the suspension. Advance notice is typically given before service is suspended, except in emergency situations. Customers are typically given 30 to 60 days to address the reasons for suspension before termination of service is initiated.

When evaluating the **service suspension policies contained in the Customer Agreement**, customers should consider the following best practices:

- Ensure that the agreement specifies that advance notice will be given for all suspensions initiated by the cloud service provider (minimum of 30 days), with the possible exception of well-defined emergency situations.
- Ensure that the agreement provides sufficient time to address the reasons for suspension (minimum of 60 days).
- Ensure that the agreement specifies that the customer's content will not be deleted during service suspension.
- Ensure that advance notice will be given before termination commences (refer to the "Understanding the Exit Process" section below).
- Ensure that payment will not be due for the suspension period if it is determined that the provider incorrectly decided that the customer was at fault.

Limitations of Liability

Typically, the limitations of liability expressed in a public CSA protect the cloud service provider and greatly limit the compensation provided to the customer in cases of breach of contract. Details of liability limitations are contained in the following sections of the Customer Agreement:

- *Limitations of Liability*. This section contains language stating that the provider will not be liable for any deletion, damage or destruction of the customer's content, loss caused by the inability of the customer to use the service, etc. In addition, the aggregate liability is specified (i.e. the maximum amount the provider is liable for). This amount varies for different providers but is typically capped at the amount the customer has paid the provider for services during the 12 months preceding the claim. The potential issue with this language is that it may run contrary to local laws aimed at preventing unreasonable limitations. Such laws should be in the customer's favor in case of a conflict, but if the customer and the provider are from different states or countries, it is important to know in advance which jurisdiction will prevail. This may found in a "Governing Law" clause of the Customer Agreement.
- *Disclaimers*. This section contains language stating that the service offerings are provided "AS IS" and sometimes states that the provider makes no warranties that the customer's content will be

secure or not otherwise lost or damaged. The language differs across the public cloud service providers that were reviewed, but the general intent is to exonerate the provider in advance, even if it is unrealistic for the customer to make their own backup of the data on a continuing basis, which would negate the advantage of using a public cloud service in the first place.

- *Indemnification*. This section states that the customer and provider will indemnify, defend, and hold each other harmless from all liabilities, damages, and costs arising from a third party claim that technology used to provide the service infringes or misappropriates any patent, copyright, trade secret or trademark of such third party. Although the language differs across the public cloud service providers that were reviewed, the general intent and provisions are consistent, although indemnification is not always reciprocal.

When evaluating the **liability limitations** contained in the Customer Agreement, customers should:

- Carefully review the provider's *aggregate liability,* since this amount differs across providers.
- Ensure that the disclaimers exclude cases where the provider is grossly negligent.
- Compare the indemnification and disclaimer clauses to ensure there are not significant differences between the public cloud service providers being considered.
- Verify that the indemnification clause is reciprocal – it's not just the customer protecting the provider, but the other way around too.
- Understand the legal environment in which the liability limitations apply since some jurisdictions prevent unreasonable limitations of liability.

Intellectual Property

Besides the protection of the cloud service customer's confidential information, which may contain non-public intellectual property, there are additional potential issues to consider.

In delivering its cloud service, the provider must not violate any applicable law, rule or regulation, contracts with third parties, or infringe on patents, trademarks, copyrights, trade secrets, and so on. Doing so might expose the provider to suspension of its right to operate, which would cause harm to the customer. The agreement should include an indemnity clause to ensure that customers are held harmless in case of a third party claim of violation of intellectual property. Indemnity clauses in CSAs are often written to protect the cloud service provider against the consequence of customer actions (and this may be legitimate), but the reverse is not as common.

Customer content stored in the cloud by the cloud service customer is normally protected and remains the customer's property. The provider may claim a license to use the customer content, but purely for the purpose of providing the cloud service itself. Customer content can include the following categories: software, machine images, data and text, audio, video and images.

Where material (data, software, etc.) is supplied by the provider as part of – or in association with – the cloud service, the situation can be more complex. The cloud service customer may own copyright in the supplied materials or may have a license to use the materials, but the cloud service provider can retain rights in the materials (e.g., to use them with other customers or other services).

Cloud service providers who support community education and user support forums for their customers make a distinction between "customer content" (as just described) and "customer submissions," which are considered public material. In some cases, submissions may be subject to public licensing rules such as the Apache Licensing model, making the submissions openly reusable. Companies that have strong internal policies about ownership of intellectual property are advised to educate staff on any limitations applying to submissions to such forums. They should make the regular review and communication of such policies part of their ongoing information security program.

Step 3: Understand Service and Deployment Model Differences

Most services offered by cloud service providers follow one of three major *service models:* Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Service models are described in greater detail in the CSCC's *Practical Guide to Cloud Computing* [1], the *Practical Guide to Cloud Service Agreements* [3] and the NIST Reference Architecture [9], and therefore do not need to be explained here.

What is important is that each model presents significant differences in the types of cloud resource, service level objectives, and key performance indicators that are specified in the SLA. The unique characteristics of each service model are described under Step 4 below.

In addition to the service models, we also have deployment models that are classified as *Private*, *Community*, *Public*, or *Hybrid*. Again, this is described in the CSCC's *Practical Guide to Cloud Computing* [2], which offers considerations on selecting a deployment model. This paper addresses exclusively service agreements for public cloud services, and the other deployment models are out of its scope. However, when evaluating CSAs proposed by public cloud service providers, customers with very stringent requirements should remember than the other deployment models may provide appropriate alternatives.

There are in general significant differences between the CSAs across service models:

- IaaS services typically offer basic IT resources such as computing (virtual servers) and storage. Since most of the capabilities of applications and systems deployed on such cloud services are in the hands of the customer, the CSA is likely to be relatively lightweight. Many capabilities such as encryption of data, both at rest and in motion, may depend on specific actions of the cloud service customer, including the need to install, configure and run specific software components.
- At the opposite end of the spectrum, SaaS services offer complete application capabilities, with the provider usually handling the customer data that the cloud service uses as part of its functioning. Given that the responsibilities of the provider are much larger than in the IaaS case, it is not surprising to find much more substantial CSAs covering a wider range of service capabilities. The provider must be clear about data handling, information security, and the protection of PII within the service.
- PaaS services can be more complex. Much of the responsibility for applications and data placed into the cloud service lies with the customer. However, the provider is responsible for the installation and operation of substantial software stacks, such as database engines, etc. The customer should aim to find specific CSA statements that relate to these software components,

especially where such components are critical to the operation of customer applications deployed on the PaaS. Unfortunately, customers may find that important information about specific software and services is scattered across different documents.

Step 4: Identify Critical Performance Objectives

The cloud SLA is the document that specifies service level objectives by the cloud service provider. All of the public cloud SLAs that were reviewed consisted of four key components: *service level objectives, credits, credit process*, and *exclusions*. Credits and the credit process are often jointly called "remedies" by the legal profession, and this term is adopted in the ISO/IEC 19086 standard under development.

Service level objectives differ across cloud service models; therefore, different types of cloud SLAs were analyzed: IaaS SLAs (with a distinction between Compute and Storage services), PaaS SLAs, and SaaS SLAs. In general, service level objectives varied across service models, but credits, credit process, and exclusions were consistent.

- *Service level objective*. All service level objectives across service models (IaaS, PaaS, and SaaS) focused almost exclusively on uptime/availability. Few other metrics were specified. Uptime/availability is expressed as a percentage that ranges from 99.0% to 99.9%, 99.95% and even 100%, depending on the service model, and is typically measured on a monthly basis (one SLA measured it on an annual basis). The providers use percentages to express the availability SLA; however, the calculations, exclusions and algorithms vary.

  For IaaS services, downtime is measured differently across the various SLAs that were reviewed:

  - Total minutes when the service is unavailable during a billing cycle (e.g., per month)
  - Total number of errors divided by total number of requests during a specific time interval (which ranged from 5 minutes to 1 hour)
  - Elapsed time from when a case is filed until when the service is reinstated
  - For at least one SLA, "Failed Storage Transactions" included transactions not processed within a specified time period (although it is not clear how this is measured or monitored)
  - For at least one SLA, the contiguous downtime must be greater than 5 minutes before the downtime is recognized by the provider

  For PaaS or SaaS services, similar remarks are true with the definition of downtime varying significantly across providers. For example:

  - An application error rate exceeding 10% for at least 5 consecutive minutes
  - All attempts to connect fail or take longer than 30 seconds to succeed during a 5-minute period

- *Credits.* Credits are the sole form of compensation for missed service commitments across all the SLAs that were reviewed, regardless of the service models. The calculation of service credits differs significantly from provider to provider. For example:

- Tiered credit of 10%, 25%, and 50%
- Prorated credit based on unavailability
- 5% of fees for each 30 minutes of downtime

In all cases, maximum credit cannot exceed 100% of the monthly service charge. In some cases, the maximum credit is less than 100% (50% maximum in one instance). This may of course be considerably less than the damage suffered by the customer (on the other hand, when a customer suffers a failure of its own on-premise resources, it does not recover anything).

In most cases, if there is more than one service level objective impacted by an incident, only one service credit can be claimed.

- *Credit Process.* Most of the SLAs that were reviewed required the cloud service customer to take specific action to receive credit. The customer is required to identify and report failures. The timeframe for reporting them varied significantly: 48 hours, 5 days, 7 days, 30 days, 10 business days after service is restored, etc. The onus is on the customer to provide proof of the problem, including dates and times, server request logs, network trace routes, full description of the service interruptions the duration of the incidents, and, in the case of PaaS SLAs, the names of the affected databases, failed operations, and so on. In all cases, the cloud service provider reviews claims and makes a final, unilateral judgment on service credits. In some cases, the provider processes credits automatically, based on the outages calculated by the provider.
- *Exclusions*. For the most part, exclusions are similar across all of the SLAs that were reviewed. The following events are typically excluded:
    - Factors outside of the provider's reasonable control
    - Force majeure conditions
    - Failures resulting from any actions or inactions of the customer or any third party, or from equipment, software or other technology operated by the customer or a third party
    - The customer's refusal to allow the provider to perform maintenance deemed necessary to maintain the service – whether it is scheduled or emergency maintenance
    - Periods of emergency maintenance activities, or a customer-requested maintenance downtime
    - Problems with the customers' connectivity to the Internet, or other factors outside of the providers control
    - Outages that last less than a certain amount of time

When the principal capabilities of the cloud service are particular API calls (alternatively called service operations), service level commitments are typically worded in terms of requests made against that API – and in particular the number or percentage of API calls giving an error. One interesting issue for these cases is that failures can occur not only when the API call returns an error, but also when the response

time is greater than some predefined limit. The latter case can be just as important as the error case: if the API call takes too long, it may adversely impact any customer applications that are using the cloud service API.

It is important for cloud service customers to consider requesting response time service level objectives for any cloud service APIs – they are not common today, but it is clearly unacceptable from the customer perspective if an API call takes a long time to complete.

Appendix B highlights the key observations for each of the four aspects (service level objectives, credits, credit process, exclusions), focusing on the commonalities and differences that were found, and provides example language to illustrate the observations.

Appendix C provides more recommendations about the establishment of metrics definitions and a metrics program.

Recommendations

When evaluating the **service level objectives** of a public cloud service provider, or comparing providers, customers should take the following steps:

- Carefully analyze the service availability guarantees and associated credits.
- Find the observation period over which commitments are measured, and understand the business impact of a single outage corresponding to the maximum downtime occurring once during that time window.
- Analyze service credit calculations and maximum credit limits.
- Compare service credit processes, particularly the timeframe within which incidents must be reported and the type of information required to prove that a failure occurred.
- Examine commitment exclusions.
- Automate the process for detecting and logging service outages, for example by using tools that exercise the cloud service through periodic dummy transactions, recording the response time as well as detecting failures.
- Look for API call response time service level objectives, for any cloud service APIs that are time critical for cloud service customer applications.
- Recognize that the SLA metrics are limited and no standards currently exist, therefore it is ultimately the customers' responsibility to evaluate and understand them such that meaningful comparative analysis and assessments can be performed.[1]

Step 5: Evaluate Security, Privacy and Data Residency Requirements

The three interrelated but distinct concepts of security, privacy and data residency should arguably be discussed as separate steps in this white paper. Since we follow the same steps as the CSCC's *Practical Guide to Customer Service Agreements*, we have chosen to keep these issues together in this white paper since they are all covered in Step 5 of the Practical Guide.

Public cloud service providers often place considerations about security and privacy in a variety of different documents, with inconsistent titles and language. For example, security language was found in documents called "Customer Agreement," "Support Agreement," "Service Level Agreement," "Enterprise Agreement," "Contract," "Technical Overview," "Acceptable User Practices," "Security Practices," "Terms of Service," and "Privacy Statement." That last case indicates not only inconsistent naming across providers, but inconsistent classification of content by the same provider, which includes some security terms inside a privacy statement.

It is also fairly common for one of these documents to refer the reader to another document. Sometimes there is more than one level of indirection. This does not make it easy to compare security statements across providers. It also makes it hard for cloud service customers to understand the total

[1] ISO/IEC 19086 Part 2 will eventually provide a standard for Service Level Objectives.

set of statements contained in the agreement. This can lead customers to "sign with their eyes closed" rather than spending the effort required to fully understand what the agreement says about security and privacy.

Therefore, there is a need to harmonize the names and scopes of documents used across the industry in order to make it easier for customers to locate and review the relevant language. Otherwise, compliance with the clauses of these documents is made more difficult, and disputes will be harder to arbitrate.

Data residency, the set of issues raised by the location and movement of data across geographies and jurisdictions, is not often mentioned explicitly in CSAs, and many customers are unaware of the complexity and implications of this issue.

In a global environment, providers should also indicate with which national and regional security and privacy regulations they comply.

One-Sided Security Obligations

Most agreements impose stringent security obligations on the customer to protect the provider, and there are often serious consequences if these obligations are not met. While it is legitimate for the provider to tell the customer that certain practices that would endanger the security of the provider and of its other customers are not acceptable, there are several problems with such clauses:

- The provider is solely responsible for determining that a security violation occurred – opening the door to subjective interpretation leading to arbitrary actions.
- The actions taken by the provider are typically drastic, namely suspension or termination of the account, without easy recourse or mechanism for complaint submission or dispute resolution.
- Absence of any compensation for the loss of business if the suspension is found to be unwarranted.
- The jurisdiction clause limits the customer's ability to challenge a vague agreement.

On the other hand, the security language often does not impose *any* obligation on the provider to protect the security of the customer. The language in the analyzed agreements falls in the following categories:

- Generic language that says that the provider will protect the customer's data with the same level of care as if it was its own. While not very specific, this is standard language in Non-Disclosure Agreements and we therefore take it that this can be considered sufficient to hold a negligent provider accountable in a court of law.
- Language to the effect that the provider will provide some sort of "help," usually poorly specified, to allow the customer to maintain its security.
- Vague language about the provider maintaining certain security measures, usually accompanied with an obligation on the customer to determine if such measures are adequate or not. There were a couple of exceptions where the provider included a detailed description of their process.
- No mention of the provider's security measures at all.

- "Worse than nothing": in at least one case, not only does the provider fail to make any security commitment, but it explicitly declines responsibility to restore any lost data "under any circumstances" even though such circumstances could include its failure to maintain proper security.
- Finally, and fortunately, some CSAs contain security policy sections that indicate that the provider knows and applies serious measures to secure the service. Cloud service customers should look for the cloud service security measures outlined in the CSCC white paper *Security for Cloud Computing Version 2.0* [4]. The best providers list certifications they have received for their cloud services. Examples include ISO 27001 (with ISO 27017 added in some cases), SOC2, CSA Star. The advantage of this approach is that assurance is given with respect to a long list of security controls without the cloud service provider having to list them in detail (which could itself be a security risk). It also removes the need for customers to perform their own audit.

Transparency of Security Measures

Given the increasing prevalence of cyberthreats, cloud service customers need information from the provider beyond general statements that good security practices are followed.

Customers should inquire about the following points, and ask where there are corresponding commitments by the provider[2]:

- Use of data encryption within the provider's facilities, to protect backup copies, or in transit between data centers.
- Availability of reports following penetration testing or security audits.
- Notification to the customer of security breaches, violations, or suspicious activity.
- Obligation to promptly apply security patches to the operating system, database system and middleware or management tools upon supplier notification, and to keep an auditable log of these updates.
- In case there is no regular external security audit process, can the customer perform their own vulnerability testing of the provider before migrating to the cloud service or when adding a new application?
- If the cloud service provider uses subcontractors for any parts of the service, including system administration personnel, do these third parties provide an equally strong level of security?
- If PKI or symmetric keys are used to secure access to the cloud service, how are the keys managed and protected?

[2] In many cases, these security commitments are maintained in security policies separate from the CSA.

Privacy or Protection of Personally Identifiable Information

Most providers address privacy only to the extent that they tell the customer what data they will collect from the customer in order to provide or support the service, and what rights they give themselves to use that data. This data includes customer contact information, IP addresses, billing information, etc., that is, data collected in order to manage the customer relationship.

This is not what most customers are concerned about when they think of "privacy in the cloud." They're not so much concerned about their own names and addresses, but rather about the personally identifiable information (PII) they hold in the cloud about others (called PII principals in the ISO standards):

- The medical history of patients in a health care system
- Account numbers and balances of the clients of a financial institution
- Personal information about customers in a CRM system
- Accounts payable and receivable information in an ERP system
- Personal information about employees in an HR system

For some cloud services, especially IaaS, the provider typically does not know whether the customer data contains PII. As a result, these cloud services rarely offer terms that relate to the handling of such PII. Some IaaS providers acknowledge that their services can be used to store and process such data, but then they place the onus for its protection onto the cloud service customer. In some cases, particularly when the provider is certified as meeting one of the cloud security or privacy standards (e.g., ISO 27018) the provider may indicate that the cloud service offers the underlying technical means that enable the customer to protect PII if they use those capabilities appropriately.

The providers of cloud services (SaaS) that knowingly deal with PII typically pay more attention to data protection and to the various laws and regulations that apply to it. Examples include Human Resources applications, Customer Relationship Management applications, credit card payment services, and social media hosting services. In such cases, there is often (and there should always be) an extensive Privacy Policy or Data Protection section in the CSA. This is an area where CSAs have made good progress in the last few years.

Cloud service customers need to understand how PII is handled across the many systems that the provider uses in relation to the cloud service. This can include backup services, monitoring and management systems, or incident handling systems. If PII is transferred to those systems, or if PII can be inspected by those systems, then the provider must provide assurance to the customer that appropriate controls are in place to protect the PII and prevent data breaches or misuse of the PII.

Finally, there is the issue of law enforcement requests or warrants for access to customer data, which may contain PII. In some jurisdictions, the cloud service provider may be ordered not to inform the customer that the data has been accessed. However, when not prevented by the authorities, the provider should promptly inform the customer of the request, and in fact many providers indicate that they will do so.

The Need for Data Residency Commitments

Data residency is defined by the Object Management Group as "the issues and practices related to the location of data, movement of data across geographies and jurisdictions, and protection of that data against unintended access" [75]. OMG further explains that this issue is not limited to cloud computing deployments, but can also arise in other contexts; and that it is not solely an issue of personal data protection, but can also concern the right to move "sovereign data" belonging to governments or data sets with specific licensing constraints imposed by the jurisdiction where it resides.

Many organizations define "residency" as a synonym for "location." This is a narrow view that can hide some issues. For example, a person can be a resident of the UK even though they are not currently present in the UK. Their resident status submits them to certain obligations (e.g., to pay taxes on their income) even though they are not always physically in the country. The same subtle distinction can be true of data.

It is legitimate for cloud service customers to want to know:

- Where their data or application resides at a given time
- Whether this location is fixed, or can vary over time at the provider's discretion (for example, for load balancing or cost reduction reasons), including moving data across borders
- What unintended access may result, such as access by a foreign law enforcement or regulatory agency

However, the burden of properly handling this issue should not be entirely moved to the provider. The customer has a responsibility to understand how *sensitive* their data is to its location. For example, does the customer hold personal information about European Union citizens? In that case, does the cloud service provider meet the demands of the European Union in terms of data protection? The provider needs to understand the issues and must be able to comply with such requirements, but it is the customer who knows the data.

A red flag should be raised if the provider stores sensitive data outside of the jurisdiction of the data owner's country *and* is not able to describe competently the data residency regulations of all the countries where the data may end up residing. Similarly, the provider should describe whether they are using partners or subcontractors for some of their capabilities and a list of such partners should be available to the customer on request. For example, even the remote access to customer data by an agent working for an outsourced call center might present a challenge: in the course of fixing an issue, records or files manipulated by the remote technician may reside, even if temporarily, in a different jurisdiction than was initially intended.

Disaster prevention measures (covered in Step 8) may lead to additional risks. A provider may replicate customer data, for backup/recovery or "hot standby" purposes, to another data center they operate in a different country.

Cloud service providers vary in their statements about the locations in which customer data is (or may be) stored. Some say rather little, while others give precise lists of their data centers and their locations. Some providers offer no choice about the location(s) where data is stored and processed, while others

give control to the customer. In the latter case, the customer must choose and manage the locations to be used – or the locations to be excluded.

Recommendations

Customers should request, and providers should consider, the following reasonable practices regarding **security, privacy** and **data residency**:

- Security, privacy and data residency statements should be explicit, separate, and in clearly identified documents.
- The customer should look for – or demand – information about certifications held by the cloud service provider in relation to security and privacy/data protection. The customer needs to understand that it is common for such certifications to be specific to particular cloud services and needs to check the documentation carefully.
- The provider should commit to specific physical and logical security practices aimed at avoiding disruption to the customer's business (not just the other way around).
- When a provider seeks to protect itself by granting itself the right to suspend access to services by a customer when a security breach is suspected, it needs to provide an emergency mechanism to resolve the issue if the customer acted in good faith or was actually not responsible for the breach.
- The provider must investigate any incident with due diligence and inform the customer about the findings. The customer should have a fair opportunity to answer any adverse findings and defend itself. Ideally, this process should be concluded before suspension of services; however, if there is a very serious incident and the provider believes that they have clear evidence of a violation and that there is an immediate risk of further or irreparable damage, expect that they will not consent to that delay.
- If the provider takes such a measure, which is determined later to not be justified, the customer should be entitled to compensation for the business disruption suffered.
- If a security attack on the provider causes the loss of customer data, the provider should be obligated to restore the data from a recent, pre-attack backup.
- The provider should offer or subcontract (at a commercially reasonable cost) a security professional service to help the customer assess and select the appropriate security mechanisms. That service should also be available in an emergency to help diagnose and repair security issues.
- The provider should describe what facilities it offers to implement user authentication. In particular, federated identity management (with the customer's own identity management system, or with a trusted third party) can improve security by avoiding password proliferation and allowing immediate deprovisioning of a terminated employee. This information may be contained in technical documentation of the cloud service rather than in the CSA.

Recommendations (continued):

- The protection of PII contained in customer data (e.g. data about account holders when the cloud service customer is a bank) must be addressed in multiple ways:
  - The provider should disclose the measures it takes to prevent its own personnel's access to confidential information contained in the cloud systems and services rented by the customer[3]; and
  - The provider should provide advice to the customer about the vulnerabilities that exist and the possible remediation, such as the potential need to encrypt data in transit and/or at rest so that confidential information, even if intercepted, cannot be exploited.
- The provider must promptly notify the customer when data is handed over to a third party or to law enforcement, unless such notification is explicitly and lawfully prohibited.
- The provider must provide a contact or method to handle privacy issues in accordance with the data protection laws of the customer's country.
- The provider should specify where the customer's data and applications may be stored, including as a result of backup or redundancy measures. If the provider has infrastructure in multiple countries or jurisdictions, it should offer its clients the ability to specify, in the service agreement they sign, locations in which the data must or must not reside.
- The provider should demonstrate that it has knowledge of the data residency and data protection laws and regulations of each of the countries or regions where it operates.
- The customer must understand the location sensitivity of its data, and select a cloud service that will not result in violating data residency laws and regulations.

Many of the recommendations in the above list are things that many providers do not offer today as a standard part of their customer service agreement, especially for IaaS cloud services. Therefore, customers may not be able to use those considerations as hard selection criteria. Instead, they fall into the "what to negotiate" area: they should be discussed with providers, whose willingness (or not) to make reasonable commitments help determine whether they are a suitable supplier.

Step 6: Identify Service Management Requirements

The findings related to service management and maintenance in public CSAs indicate that customers should perform due diligence to ensure that the level of service is managed appropriately by the provider. Customers should not expect much to be specified within the standard service agreements, as most public cloud services are provided "as is" with the customer having sole responsibility to monitor and manage the consumed services.

Customers should also be aware that they may need to improve their internal service management capabilities and resources, including monitoring, in order to comply with terms in the CSA as well as to

[3] In many cases, cloud security standards *certification* of the cloud service provider addresses this requirement.

validate the level of service from their provider and to obtain a sufficient level of control of their own use of the cloud service.

Service management provisions and language are primarily included in two artifacts, the Customer Agreement and the cloud SLA, across service models (IaaS, PaaS, and SaaS). The service management considerations covered include: provisioning, audit, on-boarding account setup, services enablement, reporting and monitoring, metering, and support and maintenance.

Customers should also consider whether test environment(s) are required. If so, the customer must confirm that the provider can support this, and agree how test data is provisioned. This is not typically included in current public CSAs, so customers are likely to need a separate contract addendum. While there may be nothing to negotiate if this is not part of a provider's services, this fact should definitely influence the choice of provider and/or hosting model.

The use of cloud services continues to evolve into more complex multi-service arrangements involving a mix of public and private cloud resources; the business world is requesting many best-of-breed cloud services and combining them to form the optimal solution. Taking this into account, Cloud Management Platforms (CMP) are fast becoming an important component in allowing customers to successfully leverage and broker multicloud environments [70]. Effective cloud service management can therefore include a CMP, compatible with the range of cloud-based services contracted by the customer, to provide enhanced cost management, redundancy, as well as more visibility of facts about the services contracted from multiple providers.

CMPs allow customers to better benefit from multiple cloud service providers while putting in place a formal portal/dashboard/ticketing/process interface between the customer and its growing number of providers. This is an emerging area, there are only few products in this space, and in all cases work is needed to integrate the various data sources into a CMP [73].

Service Management Practices

The description of service management practices has improved in CSAs for public cloud services. In some cases, the delivery of mature service management practices by providers is implied; in other cases, the provider may state in general term that they adhere to the practices of ITIL v3 (Information Technology Infrastructure Library) [1]. In any case, the customer needs to determine what service management practices the provider employs. This is crucial to an understanding of the working relationship between customer and provider.

Customers may expect certain capabilities to be provided as standard: software maintenance and upgrades, backup, recovery, encryption, etc. In fact, there are three possible situations:

- Some providers include these capabilities automatically, and they form a foundation for their service offering.
- Others require the customer to sign up for higher, more expensive levels of service.
- Some do not offer them at all.

These capabilities may be critical considerations for a cloud computing initiative; therefore, they must be carefully evaluated and clarified.

Some system management agreements are complex and/or involve external partners of the provider (such as a CMP provider). Agreements can be different across different cloud services and geographical areas, adding to the complexity of fully understanding the agreement's obligations and constraints.

Maintenance and Updates

Within a CSA, maintenance is usually mentioned in the context of availability to explicitly state that "planned maintenance time is excluded when calculating availability." Another major provision typically states that the provider may change or remove functionality (including enhancements) at any time, with appropriate notice. Such a change could result in preventing the customer, or its own clients, from operating a business function. In turn, this makes the customer incur additional costs, such as having to fail over to another provider's cloud service. These considerations impact the total cost of ownership (TCO) of a cloud service and hence influence the cost/benefit calculation. Moreover, an immature public cloud service with frequent releases that modify or remove existing functions may force customers to consider changing providers.

It is also important for cloud service customers to understand that certain types of maintenance are highly desirable – for example, the patching or updating of software with security fixes to address known vulnerabilities. Customers should look for statements about such maintenance in the CSA.

Maintenance means different things across service and hosting models. The key is to clarify early what the maintenance services include, such as delivery cycles and assurances of quality. Service and product defects are seldom inferred in any of the service agreement documents.

One-Sided Change Management Constraints

Most agreements impose stringent process constraints on the customers, but seldom outline the services or processes that the provider utilizes to manage the services that are provided. The various agreements are written by the providers to protect the provider's assets rather than protect the customer. In many instances, these agreements state that the agreement itself may be subject to change and termination at the discretion of the provider.

Change management and configuration management are very important cloud considerations as asset licensing and volatility of functionality have significant impact on cloud computing justifications. Most of the responsibility will ultimately fall onto the customers to ensure that they comply with agreement terms and prepare for changes. Good configuration management (CM), based on solid enterprise architecture approaches, is extremely valuable to optimize cloud management and to comply with the agreement's requirements. For example, a CM product may help answer the question: "Which applications use service X, which is not compatible with a planned operating system upgrade?"

Service Metrics Definitions

Clarification of SLA metrics – and how they are monitored, measured and reviewed – remains critical: while different cloud service providers often use the same names for metrics, the detailed definitions and usage are often different.

To take an example, *availability* is the primary metric identified in the SLAs, but as the "Service Commitments" section highlights, availability is calculated and used in many different ways. Thus, a

99.5% commitment may result in a higher guarantee of service than 100%, due to the way a provider calculates and credits outages.

Another issue may present itself when one provider relies upon others to deliver the complete end-to-end service experience. For example, a customer may procure a SaaS solution that in turn relies on IaaS services from a different provider. In such a case, it is important to understand whether the first cloud service provider accepts full responsibility for meeting the service objectives, or attempts to shield itself from that responsibility when one of the supporting IaaS providers is not delivering the expected service. These *cascading SLAs* along the supply chain logically depend on each other, but the customer should not have to deal with parties other than the primary cloud service provider, whose responsibility is to shield the customer from the way it assembles the solution it delivers. Any agreements that exonerates the provider when it can shift the blame to a third party should be viewed with suspicion.

A cloud service customer must understand the provider's proposed service metrics, how they are derived, and how they are used (e.g., to calculate credits or trigger escalation). Customers may want to collect additional measurements that allow analysis aligned with their business objectives. Some providers may agree to supply this information, possibly for an additional charge. Providers who flatly reject such requests open themselves to the suspicion that their systems are not capable of collecting such data. More information about metrics approaches appears in the CSCC *Practical Guide to Cloud Service Agreements* [3].

Service Pricing

The costs of the services need to be discussed, understood and negotiated when and if possible. Services often include both non-recurring charges (NRC) and monthly recurring charges (MRC). The NRC are fixed fees that most often cover installation and configuration of the service. The MRCs are variable costs based on consumption, and are applied in accordance to service variants (e.g., Gold, Silver, Bronze, or Tiny) and selected services.

Pricing needs to be directly attached to the specific service units so that clear billing occurs for both the customer but also to support the internal business chargeback model, if one is in place. The monthly bill may vary according to consumption and to the dynamic provisioning and de-provisioning of services. Billing reviews are an important part of cloud service management. Providers are not immune to billing errors and will usually not detect those in their favor.

If there is an element of variable pricing related to user requests or excess usage, then the customer should challenge the provider to offer tools to monitor requests and usage in order to maintain control and avoid surprises. In particular:

- The pricing structure should be simple and easy to understand.
- The provider should be accountable to provide evidence of the events that resulted in variable costs.
- Contracted usage should be capped in order to prevent accidental overruns, or if consumption of variable-cost resources is uncapped, the provider should offer a facility to monitor usage and alert the customer about a potential overrun.

The itemization of costs can be tedious to negotiate but it is most useful to streamline and automate the service management and its ongoing costs. The list will be quite different according to the service model (IaaS, PaaS, SaaS). For example, itemized IaaS costs may include:

- License charges for the OS, hypervisor, antivirus, and other components of the infrastructure
- Fees for provisioning or deprovisioning of a virtual machine
- RAM or storage
- Database instances
- SSL endpoints
- Customization and configuration tasks performed by a professional services team
- Security monitoring and reporting services

Accreditations and Certification

The most unequivocal assurances often provided in a CSA concern a provider's accreditations or certifications by one or more standard-developing organizations (SDOs) or their certified auditors. The agreements reviewed mentioned the following:

- ISAE 3000 international attestation and/or US AT 101 attestation such as a Service Organization Control (SOC) report – especially SOC 2 and SOC 3 reports, which address security and trust
- FISMA (Federal Information Security Management Act) compliance
- FedRAMP
- Cloud Security Alliance – STAR registry
- Payment Card Industry Data Security Standard (PCI DSS) certification
- ISO 27001, 27002, 27017 and 27108 compliance certification by an "accredited certification body"
- FIPS (Federal Information Processing Standard) 140-2 validation, related to data encryption

Most US-based healthcare-related organizations are concerned about compliance with HIPAA, the Healthcare Insurance Portability and Accountability Act. However, there is no direct HIPAA certification for a cloud service provider. Instead, most providers align themselves with one of the existing certifications and state that this ensures that the cloud service customer can be HIPAA-compliant as a result. NIST supports this approach in SP 800-66, "An Introductory Resource Guide for Implementing the HIPAA Security Rule," which refers to NIS 800-53.

Some accreditations require assessment of critical service management processes. Specific service management requirements are not usually cited directly in the agreement, but many accreditations imply that certain mature service management processes will be utilized.

Most customers should ask for ISO 20000-1 certification, which is more recent but most useful. ISO 20000-1 is the first international standard for IT service management. It was originally developed to

reflect best practice guidance contained within the ITIL framework, although it equally supports other IT service management frameworks and approaches including the Microsoft Operations Framework and components of ISACA's COBIT framework. Some highly regulated sectors, such as banking, may find that ISO 20000-1 falls short of their regulatory authority requirements, in particular because it is a supplier attestation (not a third party's) and it represents a snapshot at a given time. For those customers, a SOC 2 assurance report (for example) may be more appropriate.

Audit

Audits (by customers or independent auditors) are not usually specified in CSAs. The certifications included in many CSAs are usually based on third party audits, intended to infer credibility without customers needing to visit facilities and perform audits. For public cloud services, providing cloud service customers with a right to audit the provider's systems is very challenging and is not provided in most cases.

If the right to audit is an important factor, the customer should attempt to negotiate it as part of the contract, but this will be at the provider's discretion. Multi-tenant cloud solutions are particularly challenging with respect to auditing and penetration testing, since the audit process by client A might impact the delivery of services to client B, or may allow client A's representatives to observe information about client B's use of the services.

Recommendations

When evaluating the **service management policies** contained in the CSA and SLAs of a public cloud service provider, customers should consider the following:

- They have the ultimate responsibility to fully understand the agreements, terms, responsibilities, activities and accountability related to service management.
- They must precisely define their objectives and ensure that the provider offers the level of support necessary to meet these objectives.
- Customizations or supplementary agreements may be needed to address specific service management objectives and concerns, but obtaining them is unlikely or at best difficult. For services requiring such specific provisions, private or hybrid clouds should be considered instead. Integration of cloud-based services from best of breed providers (e.g., Security as a Service, Disaster Recovery as a Service, Compliance as a Service) should be considered to cross-check and complete the infrastructure implementation.

Recommendations (continued):

- Customers should understand the service management capabilities available with the cloud service, whether in the form of applications or in the form of APIs.
- Customers need to consider the provider's commitments to stability of functionality over time, including APIs and Web services, and how changes can create extra costs or impact users.
- Customers must examine the definitions and potential impact of each service metric, and the extent to which the metric represents a serious commitment, which can be partially assessed from the way credits for outages are calculated. Customers may consider contracting an alternative public cloud service provider as a backup solution for the prime provider's degradation or failure of services. This may lead the customer to implement a full hybrid cloud solution.
- Customers should ask questions related to service management maturity in the various topic areas to distinguish actual capabilities from marketing claims. Discussions with other customers will help assess the provider's capabilities, and may lead to an agreement to include additional SLAs or commitments in the CSA. For business critical scenarios, customers should consider obtaining independent assurance to validate service management maturity, including commitment to keep the assurance current through annual renewals. This will ideally include a period of monitoring to ensure that stated practices are really occurring – for example via a SOC 2 Type 2 assurance report.
- Customers should not totally outsource service management; they need to retain in-house the service management expertise required to monitor and improve cloud performance.
- Customers should ask for detailed and regular metrics on contracted services. For critical services and/or large contracts, the customer should seek to establish regular operational performance review meetings, in which performance and cost data gathered by both customer and provider are reviewed and acted upon.

Step 7: Prepare for Service Failure Management

In a traditional data center, organizations are able to manage failures using a centralized service management system. In the increasingly common case where an organization builds systems that use cloud services from multiple cloud service providers, managing these multiple systems becomes a bigger challenge.

The public CSAs reviewed discuss service commitments, credits, and credits process in detail. However, when it comes to service failure management capabilities or expectations, the details are sparse. Although not much mentioned, most cloud service providers follow IT Infrastructure Library (ITIL) or ITIL-compatible practices for managing their cloud services. Customers need to pay attention to three key processes and systems used in failure management: event management, incident management and problem management.

- **Event management** involves the cloud services and their related components, generating different types of events related to the monitored functions, and then distributing,

consolidating, delivering and processing these events. The monitored functions include machine states (up/down), the status of hypervisors, stages of service processing, performance metrics collection, and more. Most cloud service failures are automatically handled by the event management system; however, there are cases when automation is not sufficient. In such cases, the event management system passes control to an incident management system by generating a ticket.

- **Incident management** involves ticket generation, ticket assignment to administrators, tracking of ticket resolution, as well as checking and updating the ticket processing status, and escalation procedures.
- **Problem management** is aimed at preventing problems, in particular by analyzing recurring incidents in order eliminate them, and minimizing the impact of incidents that cannot be totally avoided. This is an area of constant innovation through the use of analytics and predictive maintenance. Customers should find out whether a cloud service provider is employing such preemptive problem identification and resolution techniques.

Cloud service providers offer multiple mechanisms to notify customers of failures from these systems. However, the burden is on the customer to aggregate this information from multiple providers to determine the impact of such failures on their business operations. Further, the financial burden for service failure, too, falls predominately on the cloud service customer, with compensation from the provider capped at one month of service credit in most cases. In addition, the onus may be on the customer to identify any failures and to provide proof of the failure to the provider. There are also numerous exceptions for which a provider does not provide compensation. Refer to "Step 4: Identify critical performance objectives" for details.

Apart from service commitments and credits, customers may want to dig into failure metrics such as:

- Mean Time Between Failures (MTBF) – the arithmetic mean, over a period of time, of the intervals between failures. While this is a well-known concept and customers are legitimately concerned if failures occur often, MTBF is not often incorporated in cloud service SLAs.
- Mean Time to Recover (MTTR) – the arithmetic mean of the time required to repair.
- Mean Time to Failure (MTTF) - the arithmetic mean of elapsed times between a recovery and the next failure. MTTF can also be derived by subtracting MTTR from MTBF.

When considering the service objectives proposed by cloud service providers, customers need to evaluate them in light of the criticality of the services to their business. Many cloud services provide limited assurance regarding system reliability and as a result, they cannot be used in a straightforward way for customer applications that require guarantees of very high availability and reliability. However, it is often possible to engineer reliable systems using cloud services that are themselves not fully reliable. The techniques for achieving this include the use of redundant components running in physically separated cloud data centers and hot failover techniques. Some cloud services build this kind of reliability engineering into their offering, others require the cloud service customer to install appropriate additional components to achieve the required results.

Finally, customers who consider migrating to cloud services from an in-house solution should understand their current performance and failure management practices. It is a common mistake to consider a provider's commitment as insufficient, even though it is better than what the existing on-premises solution offers.

Recommendations

When evaluating **service failure** management, customers should consider the following:

- It is desirable that the provider offer APIs, webhooks, an RSS Feed, a JSON feed or other electronic means of sending failure and alert data to the customer's service management system. This enables the customer to manage all services (on-premises or cloud) in a uniform and consistent manner. The description of such interfaces may not be part of the CSA, but may appear instead in separate technical documentation.
- Conversely, some failures may go undetected by the provider (e.g., firewall changes by the provider that prevent customers from accessing cloud services). Customers must ensure that the provider offers user interfaces, APIs, or other mechanisms to report failures *to* the provider.
- The provider should provide an Expected Time to Resolution (ETR) for any service failure, however detected.
- The elapsed time between failure and recovery may exceed the advertised downtime but may still not breach the SLA. This happens because the service provider can pause the SLA clock when their support organization needs some information from the customer.
- Cloud service customers should investigate the cloud services offered to see if they support resilient features such as database replication, clustering with load balancing and so on.
- Cloud service customers should evaluate the cloud services to understand how they can build resilient applications and systems using those services, even where those services can suffer from point failures. Capabilities such as redundant systems, data replication and fail over should all be considered.
- The customers must clearly understand responsibilities and hand-off procedures. In most service agreements we reviewed, the alerting and notification method was by e-mail to the address in the agreement. This can be a big risk, even for non-critical systems, resulting in loss of productivity or missing a key milestone. Instead, we recommend selecting a public cloud service provider with a ticketing system that customers are allowed to use for reporting failures. This also makes it easier for customers to find out the ETR.
- When reviewing the data privacy part of the SLAs or AUPs, be sure to confirm that the monitoring capabilities of the cloud's service failure management systems do not violate the data privacy stipulations.
- We also recommend that customers assess MTBF, MTTR, and MTTF to determine expected service downtimes. Evaluate the probability of these downtimes against the nature of your workloads. Consider that the impact of failures may outweigh the service credits offered by the service provider, and make the appropriate decisions if this is the case.

Step 8: Understand the Disaster Recovery Plan

Disaster recovery is a subset of business continuity and focuses on processes and technology for resumption of applications, data, hardware, data communications, and other IT infrastructure in case of a man-made or natural disaster (fire, flooding, hurricane, tornado, earthquake, etc.). Outsourcing infrastructure, platforms, or applications to a cloud service provider does not absolve customers of the need for serious disaster planning. Every company is unique in the importance it assigns to specific infrastructure and applications; therefore, a cloud disaster recovery plan must be tailored to each organization, and business objectives play an important role in determining the specifics of disaster recovery planning.

In general, current public CSAs provide inadequate guarantees in case of a service outage due to a disaster. Most cloud SLAs provide cursory treatment of disaster recovery issues, procedures and processes. Instead, the CSAs that were reviewed focused on limiting the liability of the cloud service provider in disaster events, and consistently covered the following areas:

- *SLA Exclusions*. This section of the cloud SLA contains language that excludes service credits for outages caused by factors outside of the provider's reasonable control, including any force majeure event, Internet access problems, or similar issues.
- *Disclaimers*. This section of the Customer Agreement contains language stating that the service offerings are provided "AS IS" and that the provider makes no warranties that the customer's content will be secure or not otherwise lost or damaged.
- *Limitations of Liability*. This section of the Customer Agreement contains language stating that the provider will not be liable for any deletion, damage or destruction of the customer's content.

Given the clauses above, the onus is clearly on cloud service customers to define, implement and execute their own disaster recovery plans, leveraging the services of the providers in the best possible manner (i.e., backup services, geographically dispersed redundancy services, etc.). A comprehensive discussion of disaster recovery for cloud workloads [78] is very helpful for customers to understand what a disaster recovery plan might include.

Some cloud service providers explicitly offer capabilities to assist with disaster recovery. For example, the cloud services can be made available in multiple geographically separated data centers, with customer control over the placement of data and application instances. There may be the ability to copy data between those multiple sites in real time and the ability to provision application instances across the sites, with load balancing between them. Such capabilities provide the basis for rapid failover should one data center be subject to a disaster. In some cases, this is offered as a service ("Disaster Recovery as a Service"); in other cases, it is up to the customer to organize the applications and services in an appropriate way to support disaster recovery.

Note that the use of multiple data centers for the purpose of disaster recovery may conflict with data residency requirements (see Step 5 above).

Recommendations

Despite the limitations in current public CSAs, cloud service customers should address **key disaster recovery procedures** early in the process of cloud adoption:

- Customers should devise a disaster recovery plan by identifying and prioritizing applications, services and data, and determining for each one the amount of downtime that is acceptable before there is a significant business impact.
- Customers should ensure that business critical content is stored redundantly in different geographical locations to help reduce the impact of a disaster. Popular solutions include building business applications on top of those cloud services that have built-in geographical redundancy, or leveraging replication technologies (provided by a third party or by the cloud service provider) to synchronize the states of applications and systems with a remote site.
- Customers should clearly define Recovery Point Objective (RPO) and Recovery Time Objective (RTO), the two most important metrics of disaster recovery, for the devised disaster recovery plan to be practical and effective. Then the proper disaster recovery technologies for redundant storage, replication, orchestration, and other necessary automation can be determined based on the RPO and RTO values (RPO is the maximum targeted period for which data might be lost from an IT service due to a disaster; RTO is the maximum targeted duration of time within which a business process must be restored after a disaster).
- Customers should ensure an appropriate frequency of backups based on the criticality of content.
- Customers should use data and application replication capabilities where provided by the cloud service
- Customers should implement a mechanism to promptly detect and quantify outages in order to begin mitigation and/or recovery processes as soon as possible and to facilitate reporting and proving failure to the provider.

Step 9: Develop an Effective Governance Process

Customers legitimately expect an effective management process for any problems that may arise with their public cloud usage. Cloud services are now used for mission-critical functions, not just for low-impact ones; therefore, these services need to be integrated, managed, reported and governed appropriately.

However, today's public CSAs contain few provisions for customer-provider management processes. The only formal channels of communication between the customer and provider specified in the service agreement are breach of contract clauses (credit process, suspensions, termination, etc.). None of the agreements that were reviewed specify status meetings between the parties. There is seldom a defined escalation process which the customer can invoke to raise the priority of a service level issue.

Overall reporting and governance that includes elements such as change management and incident management remain infrequently described in the service agreements. In some cases, separate agreements are required to fill the gaps. For example, in the U.S. healthcare regulations known as

HIPAA, there is a concept of "business associate agreement" that extends the obligations of a "covered entity" to its own suppliers.

As a result, customers must carefully consider the types of applications they deploy to a public cloud service. Mission-critical business services and data that require careful monitoring and fast resolution of issues may require supplemental agreements specifying an effective management process. At minimum, a single point of contact for service issue escalation should be designated. Ultimately, private or hybrid cloud approaches may be more appropriate for such business services.

Step 10: Understand the Exit Process

In most cases, details of the exit process are contained in the Termination clause that is part of the Customer Agreement. Customers must fully understand the impact that termination will have on their data and business services, and develop a plan to ensure minimal business disruption during the resulting migration to another provider. All Termination clauses define two basic types of termination:

- *Termination for Convenience*. Customers can typically stop using the cloud service at any time. Likewise, cloud service providers may terminate the agreement for convenience at any time without liability to the customer. Advance notice is typically given before termination occurs (usually 30 days). In some cases, customers may be required to pay a penalty if they terminate an agreement for convenience.
- *Termination for Cause.* Either party may terminate the agreement if there is a material default or breach of agreement by the other party, and that party fails to cure the breach within a certain time period after receipt of notice (typically, 30 days). In some cases, for example when security violations are alleged, the provider gives itself the right to suspend services *immediately* in order to protect itself and other customers, pending resolution or termination.

Termination due to the closing of the provider's business is usually not defined. Providers obviously do not like to mention the risk that they might fail and cease operations. The customer must have a clear understanding of what would occur if the provider business failed, including both service and data recovery implications.

The effect of termination is that all rights under the agreement expire at the end of the notice period. The customer must pay all fees and charges incurred through the effective date of termination. Any provider content the customer has in its possession must be immediately returned or destroyed.

There must be a period of time, and a defined process, for the customer to recover data held in the cloud service. The level of assistance given by the provider during the termination phase varies significantly – clearly, the provider is not greatly motivated to do more (or faster) than what the Customer Agreement specifies. In all cases, the onus is on the customer to copy their content, and to verify that the copy is usable before the original is deleted.

Recommendations

When evaluating the **termination policies**, customers should consider the following best practices:

- Customers should ensure their agreement specifies that advance notice will be given for all terminations initiated by the cloud service provider (minimum of 30 days).
- Customers must put in place contingency plans and procedures to find a new cloud service (or bring the applications and data back in-house), extract and reload their data, and switch to the new cloud service within this time window.
- As part of the termination process, providers should offer assistance to customers to facilitate data extraction (e.g., clear and concise migration documentation, or assistance from a professional services department).
- The agreement should specify that all data and information belonging to the customer will be maintained for a specific time period after transition (in case it takes some time to discover a problem with the initial extraction process), and then be completely removed immediately after.
    - The typical data retention period is 1 to 3 months which gives the customer sufficient time to verify that all data has been correctly migrated to a new service.
    - Only with the customer's written notice should data be removed and destroyed before that time.
- At the completion of the exit process, customers should receive written confirmation from the cloud service provider that all of the customer's data has been completely removed from the provider's systems.

Conclusion

The CSA landscape continues to evolve. While some agreements are still rudimentary in terms of assurances offered to cloud service customers, it is encouraging to see that more and more cloud service providers offer extensive CSAs. Some of the best examples specify comprehensive security capabilities and measures for the protection of personally identifiable information.

Unquestionably, as the cloud computing market continues to mature, providers will continue to offer more specific terms to their customers in the CSA. However, the inconsistent terminology and the scattering of information among many different documents remain problematic. This makes it hard to compare offerings from multiple cloud service providers. In fact, some of the most useful information may not be in the CSA at all, but contained in the general technical documentation for the cloud service. This particularly applies to capabilities such as resilience and redundancy, especially for IaaS offerings.

New initiatives, such as the development of the ISO/IEC 19086 standard or the European Union's Service Level Agreement Legal and Open Model project (SLALOM) [6] provide hope for greater consistency of the terminology used to define service level objectives.

In the meantime, cloud service customers must carefully evaluate the materials provided about each cloud service they are considering. The recommendations outlined in this document should enable cloud service customers to build an evaluation matrix or to understand the questions they should ask about missing materials and ambiguous commitments. Cloud computing has much to offer – customers just need to be clear about what they are actually getting.

References

Foundation Materials, Standards and Regulations

[1] Axelos: *Information Technology Infrastructure Library.* www.axelos.com/best-practice-solutions/itil

[2] Cloud Standards Customer Council (2014). *Practical Guide to Cloud Computing Version 2.0.* www.cloud-council.org/deliverables/CSCC-Practical-Guide-to-Cloud-Computing.pdf

[3] Cloud Standards Customer Council (2015). *Practical Guide to Cloud Service Agreements Version 2.0*. www.cloud-council.org/deliverables/CSCC-Practical-Guide-to-Cloud-Service-Agreements.pdf

[4] Cloud Standards Customer Council (2015): *Security for Cloud Computing: 10 Steps to Ensure Success Version 2.0.* http://www.cloud-council.org/deliverables/CSCC-Security-for-Cloud-Computing-10-Steps-to-Ensure-Success.pdf

[5] Cloud Standards Customer Council (2016). *Practical Guide to Hybrid Cloud Computing*. www.cloud-council.org/deliverables/CSCC-Practical-Guide-to-Hybrid-Cloud-Computing.pdf

[6] European Union: *Service Level Agreement Legal and Open Model project (SLALOM).* http://slalom-project.eu/

[7] International Organization for Standards (2014). *ISO/IEC 17789 Cloud Computing Reference Architecture.* http://standards.iso.org/ittf/PubliclyAvailableStandards/c060545_ISO_IEC_17789_2014.zip

[8] International Organization for Standards (In preparation): ISO/IEC *19086 Part 1: Service Level Agreement (SLA) Framework.*

[9] National Institute for Standards and Technology (2011): *NIST Cloud Computing Reference Architecture.* www.nist.gov/customcf/get_pdf.cfm?pub_id=909505

[10] National Institute for Standards and Technology (2013): Special Publication 800-53, *Security and Privacy Controls for Federal Information Systems and Organizations*, Rev. 4. http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf

[11] U.S. Department of Health and Human Services: *Sample Business Associate Agreement Provisions.* www.hhs.gov/hipaa/for-professionals/covered-entities/sample-business-associate-agreement-provisions/index.html

Cloud Service Agreements

[12] Acquia Cloud Free Agreement: www.acquia.com/acquia/agreement-ft

[13] Amazon EC2 Service Level Agreement: http://aws.amazon.com/ec2/sla/

[14] Amazon S3 Service Level Agreement: http://aws.amazon.com/s3/sla/

[15] Amazon Web Services Acceptable Use Policy: http://aws.amazon.com/aup/

[16] Amazon Web Services CloudFront Service Agreement: http://aws.amazon.com/cloudfront/sla

[17] Amazon Web Services Customer Agreement: http://aws.amazon.com/agreement/

[18] Amazon Web Services RDS Service Agreement: http://aws.amazon.com/rds/sla

[19] Amazon Web Services Route 53 Service Agreement: http://aws.amazon.com/route53/sla

[20] AppRiver Terms of Subscription: https://www.appriver.com/services/secure-hosted-exchange/

[21] AT&T Cloud Architect Acceptable Use Policy: www.corp.att.com/aup/

[22] AT&T Cloud Architect Privacy Policy: www.att.com/gen/privacy-policy?pid=2506

[23] AT&T Cloud Services License Terms: www.synaptic.att.com/clouduser/html/home/ATT_Cloud_Services_License_Terms.htm

[24] BlueHost Terms of Service: https://www.bluehost.com/terms_of_service.html

[25] Centurylink Privacy Agreement: www.ctl.io/legal/privacy/

[26] Dell Cloud Solutions Agreement: http://i.dell.com/sites/doccontent/shared-content/solutions/en/Documents/Cloud-Solutions-Agree ment-UK-EN.pdf

[27] Dimension Data Privacy Policy: www.dimensiondata.com/en-US/Policies/Pages/Privacy-Policy.aspx

[28] Dimension Data Service Level Agreement: http://cloud.dimensiondata.com/am/en/about/legal/service-level-agreement

[29] Dropbox Security & Privacy certifications: www.dropbox.com/en/help/238

[30] Future Hosting Service Level Agreement: www.futurehosting.com/legal/dedicated-service-level-agreement/

[31] GoGrid Service Level Agreement: www.gogrid.com/legal/service-level-agreement-sla

[32] Google Cloud Platform Terms of Service: https://cloud.google.com/terms

[33] Google App Engine Service Level Agreement: https://cloud.google.com/appengine/sla

[34] Google Apps Service Level Agreement: www.google.com/apps/intl/en/terms/sla.html

[35] Google Cloud Storage, Google Prediction API and Google BigQuery SLA: https://developers.google.com/storage/docs/sla

[36] IBM SoftLayer Master Services Agreement: http://cdn.softlayer.com/SoftLayer_MSA.pdf

[37] IBM SoftLayer Support page: www.softlayer.com/m/support

[38] IBM Kenexa Term of Use: https://www-03.ibm.com/software/sla/sladb.nsf/pdf/6340-01/$file/i126-6340-01_02-2014_en_US.pdf

[39] IBM Kenexa Cloud Services Agreement: https://www-03.ibm.com/software/sla/sladb.nsf/pdf/6512-01/$file/i126-6512-01_02-2014_en_US.pdf

[40] Microsoft Azure SLAs: https://azure.microsoft.com/en-us/support/legal/sla/

[41] Microsoft Azure Disaster Recovery: https://azure.microsoft.com/en-us/services/site-recovery/

[42] Microsoft Trust Center: www.microsoft.com/en-us/TrustCenter/Privacy/default.aspx

[43] Navisite Acceptable Use Policy: www.navisite.com/legal/acceptable-use-policy

[44] Navisite Privacy Policy: www.navisite.com/legal/privacy-policy

[45] Netsuite Privacy Policy: www.netsuite.com/portal/privacy.shtml

[46] Netsuite Service Level Commitment: www.netsuite.com/portal/pdf/netsuite-service-level-commitment.pdf

[47] Oracle Cloud Services Agreements: www.oracle.com/us/corporate/contracts/cloud-services/index.html

[48] Progress Sitefinity Cloud Services Agreement: www.sitefinity.com/editions/cloud-services-agreement

[49] Rackspace Service Level Agreement: www.rackspace.com/information/legal/cloud/sla

[50] Rackspace Acceptable Use Policy: www.rackspace.com/information/legal/aup/

[51] Salesforce Master Service Agreement: www.sfdcstatic.com/assets/pdf/misc/salesforce_MSA.pdf

[52] Salesforce Data Processing Addendum (Privacy & Security policy): http://www.sfdcstatic.com/assets/pdf/misc/data-processing-addendum.pdf

[53] Salesforce Security, Privacy and Architecture: https://help.salesforce.com/servlet/servlet.FileDownload?file=0150M000003KgdjQAC

[54] Salesforce Heroku Enterprise Acceptable Use Policy: www.heroku.com/policy/aup

[55] Salesforce Heroku Enterprise Privacy Policy: www.heroku.com/policy/privacy

[56] Salesforce.com Premier Success Plans: http://www2.sfdcstatic.com/assets/pdf/datasheets/DS_SuccessPlans.pdf

[57] SAP Cloud Services Agreements: http://go.sap.com/about/agreements.sap-cloud-services-customers.html

[58] SAP HANA PaaS Privacy Policy: http://go.sap.com/about/legal/privacy.html

[59] Schneider Electric Cloud Services Agreement: http://software.schneider-electric.com/legal/cloud-services/

[60] Twilio Acceptable Use Policy: https://www.twilio.com/legal/aup

[61] Twilio Privacy Policy: https://www.twilio.com/legal/privacy/developer

[62] Twilio Service Level Agreement: https://www.twilio.com/legal/service-level-agreement

[63] VMWare vCloud IaaS Privacy Policy: www.vmware.com/help/privacy.html

[64] VMWare vCloud IaaS Service Level Agreement: www.vmware.com/be/support/vcloud-air/sla.html

Papers and Articles

[65] Baudoin, Claude R.: *Cloud Ecology: Surviving in the Jungle.* Cutter IT Journal, March 2013, pp. 19-25. www.cutter.com/article/cloud-ecology-surviving-jungle-417111

[66] Betts, Dominic et al.: *Building Elastic and Resilient Cloud Applications.* Microsoft Patterns & Practices series, 2012, 252 pages. https://www.amazon.co.uk/Building-Resilient-Applications-Microsoft-practices-ebook/dp/B00GRKM0Y6

[67] Cain, Christopher: *Basic Understanding Can Clear Fog Surrounding Cloud Computing Agreements.* In Business, 2010, www.ibmadison.com/Blogger/Open-Mic/February-2010/Basic-Understanding-Can-Clear-Fog-Around-quotCloud-Computing-quot-Agreements-submitted-by-Christopher-C-Cain

[68] Chow, Richard et al. (2009). *Controlling data in the cloud: outsourcing computation without outsourcing control.* In Proceedings of the 2009 ACM workshop on Cloud computing security (CCSW '09), ACM, New York, pp. 85-90. http://doi.acm.org/10.1145/1655008.1655020

[69] European Commission Article 29 Data Protection Working Party: *Opinion 05/2012 on Cloud Computing.* http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp196_en.pdf

[70] Gartner: *Cloud Management Platforms.* IT Glossary. www.gartner.com/it-glossary/cloud-management-platforms/

[71] Golden, Bernard: *Cloud Computing: The Truth About What Runs on Amazon.* CIO, September 2010. www.cio.com/article/618385/Cloud_Computing_The_Truth_About_What_Runs_on_Amazon

[72] Kertesz, Attila et al. (2009): *An SLA-based resource virtualization approach for on-demand service provision.* Proceedings, 3rd international workshop on Virtualization Technologies in Distributed Computing (VTDC '09). ACM, New York, pp. 27-34. http://doi.acm.org/10.1145/1555336.1555341

[73] Magalhaes, Ricky M. and Monique L. (November 2014): *Selecting Cloud Management Platforms.* www.cloudcomputingadmin.com/articles-tutorials/architecture-design/selecting-cloud-management-platform-part1.html

[74] NTT America (2012): *An Evaluation Framework for Selecting an Enterprise Cloud Provider.* www.us.ntt.com/resources/white-papers/an-evaluation-framework-for-selecting-an-enterprise-cloud-provider.html

[75] Object Management Group (April 2016): *Addressing Data Residency Challenges.* Webinar presentation. www.omg.org/data-residency/OMG-Webinar-Addressing-Data-Residency-Challenges-4-14-16.pdf

[76] Ponemon Institute (2011): *Security of Cloud Computing Providers Study*. www.ca.com/~/media/Files/IndustryResearch/security-of-cloud-computing-providers-final-april-2011.pdf

[77] Pucciarelli, Joseph (July 2011): *IT Cloud Decision Economics: 10 Best Practices for Public IT Cloud Service Selection and Management.* http://www.hrbrief.com/content18064

[78] Wang, Long et al. (June 2015): *Experiences with Building Disaster Recovery for Enterprise-Class Clouds*. In Proceedings of 45th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2015).

Appendix A – Analysis of AUP Content

This table contains key observations and actual language examples contained in public cloud AUPs.

Subject	Key Observations	Example Language
Content-Based Prohibitions	Every AUP analyzed had some form of prohibition of unacceptable content. Some AUPs described in detail specifically prohibited content types, while others were general policies that put the determination of acceptable content under the subjective control of the cloud service provider.	"You will not distribute, publish, send, or facilitate the sending of unsolicited mass e-mail or other messages, promotions, advertising, or solicitations (like 'spam'), including commercial advertising and informational announcements. You will not alter or obscure mail headers or assume a sender's identity without the sender's explicit permission."
Security-Related Prohibitions	Most AUPs contained wording that specifically prohibits activities that would compromise the security of the service itself or the security of another organization, or both.	"You may not use the Services to violate the security or integrity of any network, computer or communications system, software application, or network or computing device (each, a "System"). Prohibited activities include: Unauthorized Access; Monitoring of data or traffic; Falsification of Origin."
Service Integrity Prohibitions	Most AUPs included specific prohibitions against doing harm to the service itself. These were mostly related to performance (such as network abuse or attack), but sometimes they included attempts to bypass service limitations which could jeopardize the quality of the service for others.	"You may not make network connections to any users, hosts, or networks unless you have permission to communicate with them. Prohibited activities include: Monitoring or Crawling; Denial of Service (DoS); Intentional Interference; Avoiding System Restrictions."
"Rights of Others" Prohibitions	Many, but not most, of the services contain some level of prohibition against violating the rights of other people. This is separate and distinct from violating the service levels of others, and reaches into their own legal rights as fellow humans.	"Customer agrees not to, and not to allow third parties (including End Users) to use the Services to violate, or encourage the violation of, the legal rights of others (for example, this may include allowing End Users to infringe or misappropriate the intellectual property rights of others in violation of the Digital Millennium Copyright Act)."
Other Prohibitions	There was a wide range of additional prohibited activity unique to some of the AUPs. In many cases those items fell into general category, prohibiting things such as "Abuse" in general, or "Other activities."	"Prohibited uses and activities include, without limitation, any use of the Services in a manner that, in our reasonable judgment, involves, facilitates, or attempts advocating or encouraging violence against any government, organization, group, individual or property, or providing instruction, information, or assistance in causing or carrying out such violence, regardless of whether such activity is unlawful."

Appendix B – Analysis of Cloud SLAs

This table contains key observations and actual language examples specific to Cloud SLAs.

Subject	Key Observations	Example Language
Service Commitment	All of the cloud service commitments reviewed focused exclusively on uptime/availability. • Uptime/availability is expressed as a percentage • Typical percentages included 95.0%, 99.9%, 99.95%, and 100%. • The uptime/availability percentage is typically measured on a monthly basis (one SLA measured it on a yearly basis) Uptime/availability is measured differently across the SLAs that were reviewed: • Based on the total minutes the service is unavailable over a billing cycle (e.g., per month) • Based on the total number of errors divided by the total number of requests during a specific time interval • Based on the elapsed time from when a case is filed until the service is reinstated.	"Customer will receive a service credit for the period of time starting when a Case is filed requesting assistance in accessing Customer data until the service is reinstated." "'Monthly Uptime Percentage' means total number of minutes in a month, minus the number of minutes of Downtime suffered from all Downtime Periods in a month, divided by the total number of minutes in a month." "'Downtime' means more than a ten percent Error Rate for any Eligible Application." One document contains a chart that replaces, but is equivalent to prior language that read as follows" "If in any month the availability percentage is less than 99.9%, Consumer is eligible to receive a Service Credit."
Credits	Service credits are the sole form of compensation for missed service commitments across all the SLAs that were reviewed. • Calculation of service credits differs significantly, including tiered credit of 10%, 25%, and 50%; prorated credit based on unavailability; 5% of fees for each 30 minutes of downtime. • In all cases, the maximum credit cannot exceed 100% of the monthly service charge. In some cases, the maximum credit is lower (50% maximum in one instance). • In most cases, if more than one SLA is impacted by an incident, only one SLA service credit can be claimed.	"If the availability percentage is less than 99.9%, Consumer is eligible to receive a Service Credit in an amount equal to the prorated sum of the per hour charges for the base compute resource for all Instances for the number of the Qualified Outage Minutes." "The aggregate maximum number of Financial Credits to be issued to Customer for any and all Downtime Periods that occur in a single billing month shall not exceed 50% of the amount due by Customer for the Application for the applicable month." "The minimum period of Failure eligible for a credit is 15 minutes, and shorter periods will not be aggregated. The maximum credit for any single Failure is one month's Service fees."

Subject	Key Observations	Example Language
Credit Process	All of the SLAs that were reviewed required the customer to take specific action: • Customer is required to identify and report failures. • The timeframe for reporting failures varied significantly: 48 hours, 5 days, 7 days, 30 days, 10 business days after the end of the billing cycle in which the errors occurred, fifth day of the month following the month in which the failure was observed, etc. • Customer must provide "proof" of breach including dates/times, server request logs, network trace routes, full description of service interruption, the duration of the Incidents, and, in the case of PaaS SLAs, the names of affected databases, failed operations, etc. • Cloud service provider reviews claims and makes final, good faith judgment on service credits.	"To properly claim an SLA credit due, the Customer's master administrative user must open an SLA ticket located inside the Customer portal within seven (7) days of the claimed outage. Customer must include service type, IP Address, contact information, and full description of the service interruption including logs, if applicable." "To submit a Claim, Customer must contact Customer Support and provide notice of its intention to submit a Claim. Customer must provide to Customer Support all reasonable details regarding the Claim, including but not limited to, detailed descriptions of the Incident(s), the duration of the Incident, network traceroutes, the URL(s) affected and any attempts made by Customer to resolve the Incident."
Exclusions	For the most part, exclusions are similar across all of the SLAs that were reviewed. The following events are typically excluded: • Factors outside of the provider's reasonable control. • Force majeure conditions. • Any actions or inactions of the customer or any third party resulting in the outage. • Customer and/or third-party equipment, software or other technology contributing to the failure. • Customer's refusal to allow provider to perform maintenance deemed necessary to maintain the Service, whether scheduled or emergency.	"Other activities, customer directs, denial of service attacks, natural disasters, changes resulting from governmental, political, or other regulatory actions or court orders, strikes or labor disputes, acts of civil disobedience, acts of war, acts against parties, and other force majeure events." "The SLA does not apply to any errors: (i) caused by factors outside of provider's reasonable control; (ii) that resulted from Customer's software or hardware or third party software or hardware, or both; (iii) that are result of abuses or other behaviors that violate the Agreement."

Appendix C – Metrics Programs

To be successful in procuring, transitioning and operationalizing cloud services, an organization must have clear requirements expressed in measurable terms. Successful metrics programs start small and expand progressively, always justifying the introduction of new metrics based on what decisions they enable.

Metrics can be classified according to the stages of cloud adoption or migration:

- **Procurement**
  - Evaluating, selecting and procuring cloud services
  - Contracts: Defining and enforcing service level agreements (SLAs)
- **Transition**
  - Time, cost and required resources to migrate application capabilities to cloud
- **Development & Operations (DevOps)**
  - Accountability of cloud service provider
  - Auditability of service
  - Agility (How fast services could be deployed)
  - Assurance (likelihood of service to work as expected)
  - Monitoring of cloud services
  - Performance and Quality of Service (QoS)
  - Security and privacy
  - Total Cost of Ownership (TCO)
  - Usability (ease of use)
- **Retirement**
  - Cost to retire services from cloud
  - Cost to transition to another cloud service provider

At the time this paper is being finalized, the NIST Cloud Audit subgroup is in the process of finalizing a set of recommendations that contains the following "Top 13 metrics":

- Availability (consumer perspective) and Resource Utilization (service provider perspective)
- Cost (Total Cost of Ownership)
- Functionality Responsiveness (speed of functionality/ services being made available)
- Level of Interoperability and Automation
- Level of automation for Scalability and Monitoring
- Level of integration for Billing and Cross charge
- Quality of Service (QoS)
- Reliability
- Resiliency and Fault Tolerance
- Performance ex: Computation, Responsiveness, Bandwidth, Throughput, Latency
- Security and Privacy Controls
- Time-to-Value (speed of the overall solution being made available)
- Usability (Ease of Use)

Appendix D – Security

This table contains key observations and actual language examples about key security issues.

Subject	Key Observations	Example Language
Responsibility for security of the other party	Most agreements are asymmetrical: the customer is responsible for protecting the provider, and must notify the provider in case of breach, but not the other way around. A few providers commit to informing the customer promptly in case of a security breach, and to provide all information available to them about what happened. Some providers, as part of a higher-tier support agreement, assign a contact person with responsibility to administer security (e.g., manage user accounts).	"...we and our affiliates are not responsible for unauthorized access to your account. You will contact us immediately if you believe an unauthorized third party may be using your account or if your account information is lost or stolen." "This SLA does not cover (without limitation): … failures due to denial of service attacks." "[We are] not responsible for the privacy or security practices of our customers, which may differ from those set forth in this privacy statement." "We do not promise that the Services will be uninterrupted, error-free, or completely secure"
Business risk and liability	Providers assume no responsibility for "making the customer whole" if there is a breach for which they are responsible. Some providers include unspecific assurances that they will assist the customer. Most providers shield themselves from liability, in more or less explicit terms. The language at right is one of the bluntest expressions of this liability limitation.	"...Under no circumstances… shall [provider] or its suppliers be liable to customer or any other person for any indirect, special incidental, exemplary, punitive or consequential damages of any kind…"
Restoration of lost data	Most providers ignore the issue of restoring data that may have been deleted as a result of a security breach. Some explicitly deny having to do anything.	"… Under no circumstances will [provider] be responsible for the restoration of any data to cloud storage or for the loss of any data."
Physical security measures	Most providers are silent about their physical security measures, or about the personnel screening measures they perform to avoid insider attacks. The language at right is a positive exception.	"[Provider] will ensure the presence of a professional security guard in the computer server hosting facilities at all times, charged with enforcing [provider's] security policies."

Appendix E – Privacy

This table contains key observations and actual language examples about key privacy issues.

Subject	Key Observations	Example Language
Information collected about the customer	Most agreements specify in some detail the kind of information collected by the provider about the customer itself, and necessary to conduct business, including contact information and billing information. These agreements go on to justify this practice, and to define what the provider may or may not do with this information.	"We may use your Confidential Personal Information to provide you with and manage the services you request, communicate with you …, personalize the content we deliver, conduct industry or consumer surveys, manage, improve and troubleshoot our network and services, enforce our Terms of Service, or for any purpose otherwise permitted or required by law." "Each party will: (a) protect the other party's Confidential Information with the same standard of care it uses to protect its own Confidential Information; and (b) not disclose the Confidential Information, except to Affiliates, employees and agents who need to know it and who have agreed in writing to keep it confidential."
Personal data that may be stored by the cloud service provider	Many SaaS applications (collaboration, CRM, ERP, Web conferencing, etc.), as well as IaaS storage services, will result in personal information about the customer's own customers, employees, suppliers, etc., being held by the provider. Yet most agreements make no mention of any protection given to that data. In some cases, the agreement spells out that the Customer needs to protect its own customers, even though it doesn't say that the Provider is doing so itself (the third example at right is the most egregious in this respect).	"Customer agrees to protect the privacy and legal rights of its End Users under all applicable laws and regulations." "The Customer acknowledges and agrees that the Customer is solely responsible for any personal information that may be contained in the Content…" "[Provider] cannot commit to particular confidentiality obligations regarding any Content or Customer confidential information."
Location information	Some agreements explicitly acknowledge that the provider may know where the user is located when they interact with the service. There is no assurance that this information will not be exploited.	**"**When you download or use apps created by [provider] or our subsidiaries, we may receive information about your location and your mobile device.**"**