OBJECT MANAGEMENT GROUP®

Cloud Service Agreements:
What to Expect and What to Negotiate

Version 3.0

A Discussion Paper from the OMG Cloud Working Group

Document mars/2019-09-07

September 2019

This paper presents a discussion of technology issues considered in a Subgroup of the Object Management Group. The contents of this paper are presented to foster wider discussion on this topic; the content of this paper is not an adopted standard of any kind. This paper does not represent the official position of the Object Management Group.

This page intentionally left blank

Table of Contents

Copyright Notice

What is New in Version 3.0

Version 1.0 of this white paper was published in 2013, and version 2.0 in 2016.

Since then, 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.

Over the last year, data protection issues have become more visible, in part due to the promulgation of the European Union's General Data Protection Regulation (GDPR).

Meanwhile, the respective roles and responsibilities of the parties have been complicated by the proliferation of cloud service resellers. We address this in Step 1, and we added a new Appendix A to provide more details about these new relationships.

Version 3.0 of our paper takes this maturation and evolution of the topic of service agreements into account. Readers are invited to share their comments and feedback by e-mail to cloud-chair@omg.org.

Acknowledgements

The initial development and successive revisions of this discussion paper has been a collaborative effort, bringing together diverse customer-focused experiences and perspectives. The following participants contributed significant expertise and time to this effort:

- Dr. Rizwan Ahmad (Cianaa Technologies)
- Claude Baudoin (cébé IT & Knowledge Management)
- John Bruylant (The Cloud Turbo)
- Marcus Busby (now with Uber)
- Kristin Curran (Cloud Perspective)
- Stephen Cushing (Bendigo Adelaide Bank)
- Hannah Day (The Mayo Clinic)
- Mike Edwards (IBM)
- Jordan Flynn (eFortresses)
- Dominick Grillas (Damo Consulting)
- David Harris (The Boeing Company)
- Rajesh Jaluka (IBM)
- Roberta Mazzoli (now with the Woods Hole Oceanographic Institute)
- John McDonald (CloudOne)
- John Meegan (IBM)
- Sanjay Mundergi (Albertsons)
- Nya Murray (Trac-Car)
- Arvind Radhakrishnen (Tata Consultancy Services)
- Michael Salsburg (Unisys)
- Karolyn Schalk (IBM)
- Lisa Schenkewitz (IBM)
- Anil K. Sharma (IBM)
- Prasad Siddabathuni (Edifecs)
- Rampal Singh (HCL Technologies)
- Annie Sokol (NIST)
- Long Wang (IBM)
- Steven Woodward (Cloud Perspectives)
- John Wooten (ConsultED)

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 offers to cloud service customers 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 on CSAs is available in the OMG Cloud Working Group's *Practical Guide to Cloud Service Agreements* [13].

In general, we have found that the current terms proposed by the providers of public cloud services 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 is the recently published family of international standards on "Cloud Computing – Service Level Agreement (SLA) Framework," ISO/IEC 19086, Parts 1—4. Part 1, "Overview and Concepts," should help provide a common vocabulary for use in CSAs and in associated SLAs. [7]

Specific examples used in this paper only 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 offer 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 (CSPs) to define their CSAs (see references [19] through [71]). 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 (CSCs) can use to compare agreements from different cloud service providers.

CSCs 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 be modified at different times, independently from each other.

Key Abbreviations:

CSA:	Cloud Service Agreement
CSC:	Cloud Service Customer
CSP:	Cloud Service Provider
IaaS:	Infrastructure as a Service
PaaS:	Platform as a Service
SaaS:	Software as a Service
SLA:	Service Level Agreement
OLA:	Operating Level Agreement
AUP:	Acceptable Use Policy
GDPR:	General Data Protection Regulation

Customer Agreement

Since business service management includes the processes and procedures of the CSP, 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 CSC is expected to use the public cloud service. Alternate terminology includes "Terms of Use," "Provision of the Service" and "Services Description."
- *Personal Data Processing.* This provides specific roles and responsibilities of the CSP and CSC with respect to personal data that a CSC may store or process in the cloud. The objective is to provide a single section/document that contains the data handling details that a CSC needs in order to ensure compliance with the multitude of Data Protection laws and regulations such as the EU's GDPR. Given that such new regulations can affect organizations worldwide, the CSC should ensure that the roles and responsibilities specified in the agreement do not conflict with or hinder the CSC's plan and policy on complying with the regulations.

- *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 CSP suspends for a time the use of the cloud service by a specific CSC, 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 CSP may also specify in this section a temporary suspension clause.
- *Indemnification*. This addresses holding the CSP 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 CSC 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 CSPs we reviewed included acceptable use terms for both the CSP and the CSC:

- It is typical for the CSP 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 CSP usually agrees not to violate the intellectual property rights of the CSC.

In most cases, an AUP is published 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 CSC’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) [13].

You can expect an SLA to directly or indirectly describe a support structure. It is important to understand the SLA and its associated support tiers and services. Understanding the process for reporting problems and the turnaround time for resolution is key to speedy recovery of problems the CSC may experience.

Privacy Policies

Most public CSPs 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 CSPs. In general, the privacy policy describes the different types of information collected; how that information is used, disclosed, and shared; and how the CSP 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 CSC, or extends to the personally identifiable information (PII) of which the CSC 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 CSC is termed a *Data Controller,* is the subject of regulations and laws and is of significant concern for many CSC.

What You Can Expect and What You Should Negotiate

OMG's *Practical Guide to Cloud Service Level Agreements* [16] prescribes a series of ten steps that CSC should take to evaluate CSAs in order to compare public CSPs or negotiate terms with a particular CSP. 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 CSCs should negotiate with CSPs. 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

In this step, we will examine two issues: the additional complexity introduced by cloud service resellers, and the responsibilities imposed on the CSC by the AUP.

New Roles

CSCs need to understand the growing trend of CSPs utilizing business partners and/or value-added resellers (VARs) to sell their cloud services. A VARs often offers its own "agreement," which is mostly a "wrapper" around the CSPs' core agreements; this adds a layer of contractual complexity that needs to be navigated. CSCs should take time to examine and compare what is in the VAR's agreement vs. what is in the CSP's core agreement to make sure that (a) the VAR agreement does not weaken the commitments of the CSP, (b) they are fully aware of the VAR's role in notifications, communication, incident reporting, correction of billing errors, collection of penalties, etc. In most but not all of these reseller agreements, the VAR or business partner is largely acting as an agent for one CSP, or as an orchestrator or broker for selecting cloud services, but there is a risk of introducing additional delays or finger-pointing. A VAR often adds nothing more than "enhanced" – perhaps just meaning more attractive – billing and usage reports, and no additional technical or escalation services.

While those seem to be new roles and services, they are often simply the application of non-standard designations to services and roles that have long existed. Appendix A categorizes reseller types and lists the typical responsibilities associated with each. For example, in some countries there are agencies that serve as brokers for government clients under a title such as "Shared Services" (e.g., Shared Services Canada or the U.S. General Services Administration). These agencies assume a fiduciary intermediary role, not a technical one, which adds a layer of isolation and complexity for problem resolution, as well as extra cost. CSCs who are unsure of what a role does, based on its name, should ask the suppliers which of the categories in Appendix A they fall into.

Similarities and Differences in AUPs

The AUP is the primary artifact that should be thoroughly reviewed by CSCs to understand their responsibilities and those of the CSP. 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 CSPs 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, CSCs 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 CSPs' 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 B contains key observations and actual language examples for the most common aspects of public cloud AUPs.

Recommendations

When dealing with a **reseller or business partner** of one or more CSP, CSCs should understand the exact model under which this third party is acting. Refer to Appendix A and obtain clarification of any point that may be unclear, in order to avoid finger-pointing or delays in the provision, execution and troubleshooting of services.

When evaluating the **Acceptable Use Policy** of a public CSP, CSCs 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 CSC to clearly understand all aspects of the AUP. You should ask the CSP 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 CSP 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 CSC's best interest, because it places the burden of proof on the CSC, 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 CSP may deem unacceptable. Absent scope limitations, this might place the user in breach of contract for an action seemingly unrelated to the cloud service. CSCs 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, CSCs should exercise caution and thoroughly review every provision before agreeing to an AUP. It might be helpful for the CSC 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

CSCs 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 CSPs they work with, as well as with the policies that apply to their in-house systems. For example, cooperation between CSPs 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 OMG's *Practical Guide to Hybrid Cloud Computing* [4]. CSCs purchasing public cloud services through a reseller may find that the reseller's policies supersede the CSP's when it comes to data residency, legislative jurisdiction, and possibly the AUP and suspension of services (see Appendix A).

Areas that are typically most relevant to business concerns are:

- Data policies – residency, storage, disposal, migration, personal data protection and privacy
- 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 CSP are perhaps the most critical business-level policies to be evaluated. While these are most often expressed in the overall CSA, there may be CSP policies included in the AUP or elsewhere that need to be subjected to a thorough review.

The obligation that a CSP has to its clients and their data is partly governed by the data protection legislation applicable to PII in the user's geopolitical location – as defined in *ISO/IEC 19944: Cloud services and devices: Data flow, data categories and data use* [11] – as well as the legislation for those locations in which data may reside or may be made available. CSCs should carefully consider these legal requirements and how the CSA deals with such issues as moving data across locations to offer multisite redundancy without violating applicable laws or regulation. For commercial information which is not PII, and therefore not covered by data protection legislation, the CSA should also contain appropriate language.

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

- The CSC is solely responsible for the development, content, operation, maintenance, licensing and use of their content.
- The CSC retains all rights, title, and interest in their content and data.

- The CSC is responsible for its end users' use of their content and of the cloud service, and for their compliance with the terms of the CSA.
- The CSC is responsible for any individual's personal information (or any other confidential information) stored in the cloud. The CSC 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 CSP cannot be held responsible for any potential violations of privacy laws by the CSC.

The responsibility for maintaining appropriate security, protection and backup of the CSC's data may be shared in a way that needs to be reviewed. In the IaaS model, the CSC may be entirely responsible for this, unless an additional service is purchased from the CSP at an extra cost. Even in PaaS and SaaS models, the CSP 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 CSC 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, CSPs with an internationally distributed infrastructure allow CSCs to select where their data should – or should not – be permanently stored. This option is generally offered to government CSCs, but extends to commercial entities as well. It is a critical provision for CSCs 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 CSP 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 CSC's data.

Finally, the CSP must commit to notifying the CSC in advance of any changes in policies or in systems that affect the way in which CSC's data and content are protected.

Law Enforcement Access

The Customer Agreement should explicitly state that the CSP will not access the CSC's content. However, it usually includes an exception in which the CSP states that it will comply with properly formulated requests by law enforcement agencies. In the event of such valid legal or governmental requests, CSCs should require immediate notification from their CSP, 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 "data subjects."

As has been shown in well-publicized lawsuits, *who* can issue a valid order to produce the data can be unclear, and the laws are evolving rapidly. Therefore, the CSP 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 CSC, the nationality of the person whose data is being requested, etc.

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

- Ensure that the agreement allows the CSC to specify the physical location of their security-sensitive content, or content subject to data residency requirements (acceptable locations vary across industries and according to national legislations).
- Ensure that CSP personnel will not access the CSC's data, except when required by law and duly requested by law enforcement authorities.
- Under such circumstances, ensure that the agreement specifies that the CSP will give immediate notice, allowing the CSC an opportunity to file for a stay of the request, where permitted by law.
- Understand what capabilities the CSP offers for redundancy, replication and backup of CSC data, and what actions the CSC 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 CSPs can make such changes. CSCs 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 CSP to give advance notice (typically 30 days) to their CSCs for any such material change. For services, CSPs 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, CSPs 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, CSCs 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 CSP.
- Ensure that the agreement commits the CSP 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 CSP.

Suspension of Services

CSCs 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 should be part of every Customer Agreement and should describe in detail the circumstances under which the CSP can suspend services to a CSC. 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 CSP to liability
- Usage that represents a direct or indirect threat to the CSP'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. CSCs 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**, CSCs should consider the following best practices:

- Ensure that the agreement specifies that advance notice will be given for all suspensions initiated by the CSP (minimum of 30 days), with the possible exception of well-defined emergency situations.
- Ensure that the agreement allows sufficient time to address the reasons for suspension (minimum of 60 days).
- Ensure that the agreement specifies that the CSC'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 CSP incorrectly decided that the CSC was at fault.

Limitations of Liability

Typically, the limitations of liability expressed in a public CSA protect the CSP and greatly limit the compensation offered to the CSC 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 CSP will not be liable for any deletion, damage or destruction of the CSC's content, loss caused by the inability of the CSC to use the service, etc. In addition, the aggregate liability is specified (i.e. the maximum amount the CSP is liable for). This amount varies for different CSPs but is typically capped at the amount the CSC has paid the CSP 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 CSC's favor in case of a conflict, but if the CSP and CSC are from different states or countries, it is important to know in advance which jurisdiction will prevail. This may be 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 CSP makes no warranties that the CSC's content will be secure or not otherwise lost or damaged. The language differs across the public CSPs that were reviewed, but the general intent is to exonerate the CSP in advance, even if it is unrealistic for the CSC 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 CSC and CSP will indemnify, defend, and hold each other harmless from all liabilities, damages, and costs arising from a third-party claim that the technology used to provide the service infringes or misappropriates any patent, copyright, trade secret or trademark of such third party. While the language differs across the public CSPs 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, CSCs should:

- Carefully review the CSP's *aggregate liability,* since this amount differs across CSPs.
- Ensure that the disclaimers exclude cases where the CSP is negligent.
- Assess the limitations in light of potential damage, and consider obtaining insurance for the difference between the damage and the compensation received from the CSP.
- Compare the indemnification and disclaimer clauses to ensure there are not significant differences between the public CSPs being considered.
- Verify that the indemnification clause is reciprocal – it's not just the CSC protecting the CSP, 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 CSC's confidential information, which may contain non-public intellectual property, there are additional potential issues to consider.

In delivering its cloud service, the CSP 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 CSP to suspension of its right to operate, which would cause harm to the CSC. The CSP should therefore commit to notifying the CSCs in case of a third party's claim of violation of intellectual property. The agreement should include an indemnity clause to ensure that CSCs are held harmless in case of such a claim. As mentioned above, Indemnity clauses in CSAs are often written to protect the CSP against the consequence of CSC actions (and this may be legitimate), but the reverse is not as common.

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

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

CSPs who support community education and user support forums for their CSCs 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 CSPs follow one of three major *service models:* Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Each service model implies a different set of responsibilities between the CSC and the CSP (Figure 1).

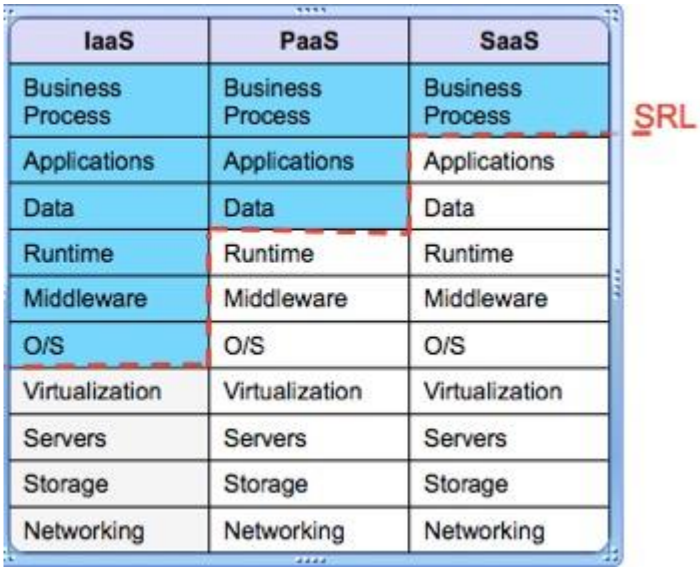

IaaS	PaaS	SaaS
Business Process	Business Process	Business Process
Applications	Applications	Applications
Data	Data	Data
Runtime	Runtime	Runtime
Middleware	Middleware	Middleware
O/S	O/S	O/S
Virtualization	Virtualization	Virtualization
Servers	Servers	Servers
Storage	Storage	Storage
Networking	Networking	Networking

Figure 1 – Service Responsibility Line (SRL). Aspects above the line are typically the CSC's responsibility Details of business continuity, escalation and disaster recovery are where responsibilities blend.

Service models are described in greater detail in the OMG's *Practical Guide to Cloud Computing* [1], the *Practical Guide to Cloud Service Agreements* [10] and the NIST *Cloud Computing Reference Architecture* [13], and therefore do not need to be explained here. What is important is that each model presents significant differences in the types of cloud resources, 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 OMG's *Practical Guide to Cloud Computing* [1], 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 CSPs, CSCs with very stringent requirements should remember than the other deployment models may offer 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 CSC, 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 CSC, including the need to install, configure and run specific software components. A CSP offering IaaS environments that meet with specific compliance types, such as FedRAMP, will have a more detailed CSA for those environments.
- At the opposite end of the spectrum, **SaaS services** offer complete application capabilities, with the CSP usually handling the CSC data that the cloud service uses as part of its operation. Given that the responsibilities of the CSP 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 CSP 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 CSC. However, the CSP is responsible for the installation and operation of substantial software stacks, such as frameworks, database engines, etc. The CSC 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, CSCs 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 CSP. 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 19086-1:2017 standard. [7]

Given the dependency on the integrity of network connectivity between corporate private and public sites, monitoring metrics play an increasing role in meeting critical end-to-end performance objectives. CSCs may not want to rely solely on the CSPs and network providers' assurances of availability and performance. They may also consider:

- real-time cloud infrastructure and network traffic monitoring,
- application performance management tools,
- AI-based monitoring of operations and security anomaly detection.

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 CSPs 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 the 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 CSP.

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

  CSCs might consider also defining and measuring repeated patterns of small (micro-failures) which individually do not reach the accepted threshold of unavailability, but together exceeds an acceptable level of disruption.

- *Credits.* Credits are the sole form of compensation for missed service commitments across all the SLAs that were reviewed, regardless of the service models. CSCs can however introduce specific language to declare repeated and acknowledged service level failures as a breach of service contract, triggering the need for escalation and remedy, above and beyond credits calculations linked to each occurrence of missed Service Level Objectives.
- The calculation of service credits differs significantly from CSP to CSP. 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, the 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 CSC (on the other hand, consider that when a CSC 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 CSC to take specific action to receive credit. The CSC is required to identify, report, and declare 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 CSC 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 CSP reviews the claims and makes a final, unilateral judgment on service credits. In some cases, the CSP processes credits automatically, based on the outages they calculated.
- *Exclusions*. For the most part, exclusions are similar across all of the SLAs that were reviewed. The following events are typically excluded:
  - Factors outside the CSP’s reasonable control
  - Force majeure conditions
  - Failures resulting from any actions or inactions of the CSC or any third party, or from equipment, software or other technology operated by the CSC or a third party

- The CSC refusal to allow the CSP to perform maintenance deemed necessary to maintain the service – whether it is scheduled or emergency maintenance
- Periods of emergency maintenance activities, or a CSC-requested maintenance downtime
- Problems with the CSC's connectivity to the Internet, or other factors outside of the CSP's control
- Outages that last less than a certain amount of time

CSCs should pay attention to the particular case of hybrid cloud services, where a CSP sub-contracts another provider for a component of services, such as telecommunications carrier services. In such a case, there is generally an Operating Level Agreement between the primary CSP and its provider; most primary CSPs will consider a failure from the sub-contracted provider outside of their reasonable control. CSCs might inquire about such OLA dispositions and their impact onto the final service performance.

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.

CSAs offer varying approaches and terms regarding automatic failover and disaster recovery. Some CSPs offer failover across geographic regions, while others rely on multiple zones within a region. CSCs may want to consider how availability is managed when designing for fault tolerance, depending on service delivery requirements, as disaster recovery planning is a critical responsibility.

CSCs should consider requesting response time SLOs for any cloud service API. These objectives are rarely defined today, but a CSC is clearly impacted if an API call takes a long time to complete.

Appendix C 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 offers example language to illustrate the observations.

Appendix D presents more recommendations about the establishment of metrics definitions and a metrics program.

Recommendations

When evaluating the **service level objectives** of a public CSP, or comparing CSPs, CSCs should take the following steps:

- Carefully analyze the service availability guarantees and the associated credits -- including service credit calculations and limits.
- Assess the credits in light of potential business/mission impacts and the associated loss that may be directly or indirectly attributable to the service outage, and consider obtaining insurance for significant differences between credits and loss.
- 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.
- Consider establishing escalation rules for repeated service failures or for patterns of micro-failures within the same time window.
- 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, including sub-contracted OLAs and their impact.
- 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 CSC applications.
- Recognize that the SLA metrics are limited and no standards currently exist, therefore it is ultimately the CSC's 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 OMG's *Practical Guide to Customer Service Agreements* [16], we have chosen to keep these issues together in this discussion paper since they are all covered in Step 5 of the Practical Guide.

Public CSPs 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 CSPs, but

[1] ISO/IEC 19086 Part 2 provides a standard for Service Level Objectives.[8]

inconsistent classification of content by the same CSP, 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 one. Sometimes there is more than one level of indirection. This does not make it easy to compare security statements across CSPs. It also makes it hard for CSCs to understand the total set of statements contained in the agreement. This can lead CSCs 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 CSCs 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 CSCs are unaware of the complexity and implications of this issue. In a global environment, CSPs should indicate which national or regional security and privacy regulations they comply with.

A particular case exists when a CSC uses a “cloud bursting” technique to relieve a temporary resource shortage in its data center (or private cloud) by leveraging a public cloud. When this occurs, the data stored temporarily in the public cloud may no longer abide by the security, privacy and compliance constraints that were in place in the private cloud.

One-Sided Security Obligations

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

- The CSP 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 CSP are typically drastic, namely suspension or termination of the account, without easy recourse or mechanism for complaint submission or dispute resolution.
- The CSP offers no compensation for the CSC’s loss of business if the suspension is found to be unwarranted.
- The jurisdiction clause limits the CSC’s ability to challenge a vague agreement.

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

- Generic language that says that the CSP will protect the CSC’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 CSP accountable in a court of law.

- Language to the effect that the CSP will offer some sort of "help," usually poorly specified, to allow the CSC to maintain its security.
- Vague language about the CSP maintaining certain security measures, usually accompanied with an obligation on the CSC to determine if such measures are adequate or not. There were a couple of exceptions where the CSP included a detailed description of their process.
- On occasions, there was no mention of the CSP's security measures at all.
- "Worse than nothing": in at least one case, not only does the CSP 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 CSP knows and applies serious measures to secure the service. CSCs should look for the cloud service security measures outlined in the OMG's white paper *Security for Cloud Computing, V2.0* [3]. The best CSPs list certifications they have received for their cloud services. Examples include ISO 27001 (with ISO 27017 added in some cases), SOC2, CSA Star, and FedRAMP for U.S. Federal Government agencies. Similar security requirements can be found in industry-specific mandates that are promulgated at the level of a country or region, such as (in the U.S.) HIPAA for healthcare or FDIC and states' cybersecurity requirements for financial services. The advantage of this approach is that assurance is given with respect to a long list of security controls without the CSP having to list them in detail (which could itself be a security risk). It also removes the need for CSCs to perform their own audit.

Transparency of Security Measures

Given the increasing prevalence of cyberthreats and increasingly complex regulatory environments, CSCs need information from the CSP beyond general statements that good security practices are followed.

CSCs should inquire about the following points, and ask where there are corresponding commitments by the CSP:

- Their process of managing risks, especially security-related risks. Standards-based processes include the *NIST Risk Management Framework* [15], ISO 27001, or the Secure Controls Framework (SCF).
- Use of data encryption within the CSP's facilities to protect backup copies, or in transit between data centers.
- Availability of reports following penetration testing or security audits.
- Notification to the CSC 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 CSC perform its own vulnerability testing of the CSP before migrating to the cloud service or when adding a new application?
- If the CSP 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, stored and protected?

Privacy or Protection of Personally Identifiable Information

Privacy is typically a concern regarding both (a) data that CSCs deliberately place in the cloud, (b) data that CSPs collect from CSCs.

CSPs should – and often – tell the CSC what data they will collect from them 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.

However, this is not what most CSCs are concerned about when they think of "privacy in the cloud." They're not so much worried about their own names and addresses, but rather about the personally identifiable information (PII) they hold in the cloud *about others*, who are called "PII principals" in the ISO standards, or "data subjects" in other texts. PII may include:

- 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
- Payment and personal information in an e-Commerce system

From the CSP's perspective, this PII is customer data and it needs to operate its cloud services in such a way that both parties abide by applicable data protection legislation, regulations, or standards such as ISO/IEC 27018, the *Code of practice for protection of personally identifiable information (PII)* [12]. There are differences, however, between IaaS/PaaS and SaaS models:

- The providers of IaaS typically do 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. While some IaaS CSPs acknowledge that their services can be used to store and process such data, they specify that the CSPs are responsible to protect the data.

- The providers of 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, social media hosting services, and many more. In such cases, there is often (and there should always be) an extensive Privacy Policy or Data Protection section either in the CSA or in a separate document.

Regardless of the cloud service model, CSCs should fully understand:

- their own requirements for privacy and other aspects of data handling,
- the CSP's contractual commitments to protect customer data,
- the tools and controls that a CSP may deploy to help protect that data,
- the mandates and protocols to manage, report and disclose breaches and violations, whether hostile or accidental.

CSPs should make it clear what customer data they will be collecting; how this data will be used; and how law enforcement requests for customer data will be handled. In some jurisdictions, the CSP may be ordered not to inform the CSC that the data has been accessed. However, when not prevented by the authorities, the CSP should promptly inform the CSC of the request, and in fact many CSPs indicate that they will do so.

CSCs need to understand how PII is handled across not only the main systems used to deliver the cloud service, but also the many additional systems that the CSP uses *in support of* 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 CSP must provide assurance to the CSC 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 CSP may be ordered not to inform the CSC that the data has been accessed. However, when not prevented by the authorities, the CSP should promptly inform the CSC of the request, and in fact many CSPs 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" [82]. 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 (ISO/IEC 19944).

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.

CSCs legitimately want to know certain things, and are even required by regulations to know them:

- where their data or application resides at a given time,
- whether this location is fixed, or can vary over time at the CSP's discretion (for example, for load balancing or cost reduction reasons), including moving data across borders
- what controls, if any, the CSP may offer in restricting such movement of data,
- what unintended access may result from changes in data location, such as access by a foreign law enforcement or regulatory agency.

CSCs have a responsibility to understand how *sensitive* their data is to its location, as well as what data handling controls are required. For example, does the CSC hold personal information about European Union citizens? In that case, will the controls offered by the CSP enable the CSC to meet the demands of the EU in terms of data protection (GDPR)? The CSP needs to understand the issues and must be able to comply with such requirements, but it is the CSC who knows the data and has ultimate responsibility for it. Using a cloud service does not relieve a CSC of their obligation to protect the sensitive data under their control.

A red flag should be raised if the CSP 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 CSP 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 CSC 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 CSP may replicate customer data, for backup/recovery or "hot standby" purposes, to another data center they operate in a different country.

CSPs 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 CSPs offer no choice about the location(s) where data is stored and processed, while others give control to the CSC – sometimes at an additional cost. In the latter case, the CSC must choose and manage the locations to be used – or to be excluded.

Recommendations

CSCs should request, and CSPs 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 CSC should look for – or demand – information about certifications held by the CSP in relation to security and privacy/data protection. The CSC needs to understand that it is common for such certifications to be scoped to particular cloud services and needs to check the documentation carefully.
- The CSP should commit to specific physical and logical security practices aimed at avoiding disruption to the CSC's business (not just the other way around).
- When a CSP seeks to protect itself by granting itself the right to suspend access to services by a CSC when a security breach is suspected, it needs to provide an emergency mechanism to resolve the issue if the CSC acted in good faith or was actually not responsible for the breach.
- The CSP must investigate any incident with due diligence and inform the CSC about the findings. The CSC 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 CSP believes that it has 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 CSP takes such a measure, which is determined later to not be justified, the CSC should be entitled to compensation for the business disruption suffered.
- If a security attack on the CSP causes the loss of CSC data, the CSP should be obligated to restore the data from a recent, pre-attack backup.
- The CSP should offer or subcontract (at a commercially reasonable cost) a professional security service to help the CSC assess and select the appropriate security mechanisms. That service should also be available in an emergency to help diagnose and repair security issues.
- The CSP should describe what facilities it offers to implement user authentication. In particular, federated identity management (with the CSC'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.

(continued on next page…)

Recommendations (continued):

- The protection of PII contained in customer data (e.g. data about account holders when the CSC is a bank) must be addressed in multiple ways:
    - The CSP 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 CSC; and
    - The CSP should provide advice to the CSC 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.

  (Commonly accepted control objectives, controls and guidelines for implementing measures to protect PII can be found in ISO/IEC 27018 [12])
- The CSP must promptly notify the CSC when data is handed over to a third party or to law enforcement, unless such notification is explicitly and lawfully prohibited.
- The CSP must provide a contact or method to handle privacy issues in accordance with the data protection laws of the CSC's country.
- The CSP should specify where the CSC's data and applications may be stored, including as a result of backup or redundancy measures. If the CSP 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 CSP 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 CSC must understand the location sensitivity of its data, and select a cloud service that will not result in violating data residency laws and regulations, while abiding by disclosure and notification requirements.

The above table shows a particularly long list of desired CSA contents. Some of this content is not offered by many CSPs as part of their standard CSA, especially for IaaS cloud services. Appendices E and F illustrate this with specific examples (and limits) of the language included by typical CSPs in the security- and privacy-related parts of their CSAs. Therefore, CSCs may not be able to use those considerations as hard selection criteria. Instead, this wish list falls into the "what to negotiate" area: it should be openly discussed with CSPs, whose willingness (or not) to make reasonable commitments will 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 CSCs should perform due diligence to ensure that the level of service is managed appropriately by the CSP. CSCs should not expect much to be specified within the standard service agreements, as most public cloud services are provided "as is" with the CSC having sole responsibility to monitor and manage the consumed services.

CSCs 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 validate the level of service from their CSP 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.

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

Cloud Management Platforms

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 multiple best-of-breed cloud services that can be combined to form the optimal solution. Taking this into account, **Cloud Management Platforms (CMP)** are fast becoming an important component in allowing CSCs to successfully leverage and broker hybrid (multi-cloud) environments [77]. Effective cloud service management can therefore involve a CMP, compatible with the range of cloud-based services contracted by the CSC, to provide enhanced cost management, redundancy, as well as more visibility of data about the services contracted from multiple CSPs.

CMPs allow CSCs to better benefit from multiple CSPs while putting in place a formal portal or dashboard to manage tickets and the process interface between the CSC and its growing number of CSPs. 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 [80].

CSCs should clearly understand the roles and responsibilities of the primary CSP and/or of a CMP provider in ensuring the final delivery of combined cloud services to the CSC.

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 CSPs is inferred; in other cases, the

CSP may state in general term that it adheres to the practices of ITIL v3 (Information Technology Infrastructure Library) [1]. In any case, the CSC needs to determine what service management practices the CSP employs. This is crucial to an understanding of the working relationship between CSC and CSP.

CSCs 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 CSPs include these capabilities automatically, and they form a foundation for their service offering.
- Others require the CSC 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 CSP (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." CSCs need to verify that the planned maintenance time is also removed from the total time of reference for the computation of availability.

Another major provision typically states that the CSP may change or remove functionality (including enhancements) at any time, with appropriate notice. Such a change could result in preventing the CSC, or its own clients, from operating a business function. In turn, this makes the CSC incur additional costs, such as having to fail over to another CSP's cloud service. These considerations impact the total cost of ownership (TCO) of a cloud service, hence the cost/benefit calculation. Moreover, an immature public cloud service with frequent releases that modify or remove existing functions may force CSCs to consider changing CSPs.

CSCs need to understand that certain types of maintenance are highly desirable and should be deferred as little as possible – for example, the patching or updating of software with security fixes to address known vulnerabilities. CSCs should look for statements about such maintenance in the CSA, including the maximum acceptable delay to apply a critical or security patch after it has been released. A good general practice has been to deploy critical patches within a month of their official release, but the constant worsening of cybersecurity attacks, and their potential consequences in terms of data breaches, may require shorter delays.

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 CSCs, but seldom outline the services or processes that the CSP utilizes to manage the services it provides. The various agreements are written by the CSPs to protect their own assets rather than the CSC's. In many instances, these agreements state that the agreement itself may be subject to change and termination at the discretion of the CSP. It is most useful for the CSC to negotiate a mirror clause that allows it to switch to another CSP according to specific criteria (exit or reversibility clauses).

Change management and configuration management are very important in the cloud as unused licenses and services can significantly impact the ROI of cloud computing. Most of the responsibility ultimately lies with the CSCs 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

Clear definitions of SLA metrics, such as those in *ISO/IEC 19086 Part 2* [8] are critical, as well as precise definition of how they are monitored, measured and reviewed. While CSPs often use the same names for metrics, the detailed definitions and usage are often different.

To take an example, *availability* is the primary metric most often identified in SLAs, but as the "Service Commitments" section of Appendix C highlights, availability is calculated and used in many different ways. Thus, a 99.5% commitment by one CSP may result in a higher guarantee of service than 100% in another CSP's SLA, due to the ways they respectively calculate and credit outages.

Another issue may present itself when one CSP relies upon other providers to deliver the complete end-to-end service experience. For example, a CSP may offer a SaaS solution that in turn relies on IaaS services from a different CSP. In such a case, it is important to understand whether the first CSP accepts full responsibility for meeting the service level objectives, or attempts to shield itself from that responsibility when the supporting IaaS provider fails to deliver the expected service. These *cascading SLAs* along the supply chain logically depend on each other, but the CSC should not have to deal with parties other than the primary CSP, whose responsibility must include shielding the CSC from the way it assembles the solution it delivers. CSCs should view with suspicion any agreement that exonerates the CSP when it can shift the blame to a third party.

A CSC must understand the CSP's proposed service metrics, how they are derived, and how they are used to measure performance, calculate credits or trigger escalation. CSCs may want to collect their own additional measurements to perform analysis aligned with their business objectives, sometimes involving tools that take measurements at the level of the end user of the service. Some CSPs may be able and willing to supply this information or facilitate its collection, possibly for an additional charge. CSPs 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 *Practical Guide to Cloud Service Agreements* [16].

Service Pricing

The costs of the services clearly need to be discussed, understood, and – when possible – negotiated. Services often incur both non-recurring charges (NRC) and monthly recurring charges (MRC). The NRCs are fixed fees, most often for installation and configuration of the service. The MRCs are variable costs based on consumption, and are applied in accordance to service tiers (e.g., Gold, Silver, Bronze, or Tiny) and the list of selected services. The monthly bill may vary according to consumption and to the dynamic provisioning and de-provisioning of services.

Pricing needs to be directly attached to the specific service units so that invoices are not only clear and justified, but also to support the CSC's internal business chargeback method, if one is in place. Billing reviews are an important part of cloud service management. CSPs are not immune to billing errors, and will usually not detect those that are in their favor.

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

- The price list should be simple and easy to understand.
- The CSP should be accountable to offer evidence of the events that resulted in variable costs.
- Either the use of such resources should be capped in order to prevent accidental overruns, or the CSP should offer a facility to monitor usage and alert the CSC about a potential overrun.

Negotiating the price list item by item can be tedious, but it is effective in streamlining and automating service management tasks, themselves a source of ongoing costs. The list will vary with 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 offered in a CSA are that a CSP is accredited or certified by one or more standard-developing organizations (SDOs) or their certified auditors. Such certifications may not be compulsory, but certification gives assurance that the CSP implements and complies with a required systematic approach. The agreements reviewed mentioned the following instruments:

- 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.

- Federal Risk and Authorization Management Program (FedRAMP).
- Cloud Security Alliance – STAR registry.
- Payment Card Industry Data Security Standard (PCI DSS) certification.
- ISO 27001, 27002, 27017 and 27018 compliance certifications by an “accredited certification body.”
- FIPS (Federal Information Processing Standard) 140-2 validation, related to data encryption.

Most U.S.-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 CSP. Instead, most CSPs align themselves with one of the existing certifications and state that this ensures that the CSC can be HIPAA-compliant as a result. NIST supports this approach in *SP 800-66 Rev. 1, An Introductory Resource Guide for Implementing the HIPAA Security Rule*, which refers to *NIST 800-53*. [14]

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 CSCs 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 (whether by CSCs or independent auditors) are not usually specified in CSAs. The certifications included in many CSAs are usually based on periodic third-party audits, intended to infer credibility without CSCs needing to visit facilities and perform audits. For public CSPs with many customers, allowing CSCs to audit their system is understandably burdensome and is generally not offered.

If the right to audit is an important requirement, the CSC should attempt to negotiate it as part of the contract, but this will be at the CSP’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 CSP, CSCs should consider the following:

- CSCs have the ultimate responsibility to fully understand the agreements, terms, responsibilities, activities and accountability related to service management.
- CSCs must precisely define their objectives and ensure that the CSP 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 cloud services should be considered instead. Integration of cloud-based services from best-of-breed CSPs (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.
- CSCs should understand the service management capabilities available with the cloud service, whether in the form of applications or of APIs.
- CSCs need to consider the CSP's commitments to stability of functionality over time, including APIs and Web services, and how changes can create extra costs or impact users.
- CSCs 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. CSCs may consider contracting an alternative public CSP as a backup solution for the prime CSP's degradation or failure of services. This may lead the CSC to implement a full hybrid cloud solution.
- CSCs 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 CSP's capabilities, and may lead to an agreement to include additional SLAs or commitments in the CSA. For business-critical scenarios, CSCs should consider obtaining an independent examiner's assurance to validate service management maturity, including commitment to renew this assurance process annually. This will ideally include a period of monitoring to ensure that stated practices are really taking place – for example via a SOC 2 Type 2 assurance report.
- CSCs should not totally outsource service management; they need to retain the in-house service management expertise required to monitor and improve cloud performance.
- CSCs should ask for detailed and regular metrics on contracted services. For critical services and/or large contracts, the CSC should request regular operational performance review meetings, in which performance and cost data gathered by both CSP and CSP are reviewed, compared, 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 CSPs, managing these multiple systems becomes a bigger challenge.

In an IaaS model, while CSPs are responsible for the virtualization infrastructure, the platform and software services that are provisioned, configured and running on top of the infrastructure are the responsibility of the CSC. Identifying the potential causes of service problems in advance is essential to ensure service continuity. In view of the complexity of network connectivity, infrastructure, platform and software services on which cloud-based applications depend, it is increasingly important to employ effective operational logging and monitoring capabilities, which may be offered by the CSP or third-party performance monitoring services. Identifying and isolating the root cause of service failures is anything but simple, and requires a trail of data that the CSP must collect.

Operations support requires increasingly specialized, capabilities. Performance monitoring dashboards must be understood and analyzed, particularly where end-to-end functions are delivered by a combination of services from multiple CSPs.

The public CSAs reviewed discuss service commitments, credits, and the credit process in detail. However, when it comes to service failure management capabilities or expectations, the details are sparse. Although not much mentioned, most CSPs follow IT Infrastructure Library (ITIL) or ITIL-compatible practices for managing their cloud services. CSCs 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. Given that the number of security incidents is rising, it may be advisable to set up specific security incident response processes for suspected security breaches or threats. This is a very useful part of endpoint security detection, and establishing automated alerts is clearly an excellent prevention measure. Several Industry or regulatory bodies have mandated specific steps and dispositions in the response process for security incidents, especially if they impact the general public or core services and infrastructure.
- **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. CSCs should find out whether a CSP is employing such preemptive problem identification and resolution techniques. These may be particularly effective to analyze repeated patterns of small incidents, which can often be traced to a common root cause.

CSPs offer multiple mechanisms to notify CSCs of failures of their systems. However, the burden is on the CSC to use this information (and aggregate it from multiple CSPs when applicable) to determine the impact of such failures on their business operations. Further, the financial burden of service failure also falls predominately on the CSC, with compensation from the CSP typically capped at one month of service. The onus may even be on the CSC to identify failures and provide proof that they occurred to the CSP. Finally, there are numerous exceptions for which a CSP does not offer compensation. Refer to *Step 4: Identify Critical Performance Objectives* for details.

Apart from service commitments and credits, CSCs may want to dig into failure metrics – see *ISO/IEC 19086 Part 2* [8] – such as:

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

CSCs need to evaluate the service objectives proposed by CSPs in light of the criticality of the services to their business. Many CSPs give limited assurances of system reliability, which may not satisfy CSCs with customer applications that require guarantees of very high availability and reliability. However, there are techniques to engineer reliable systems using cloud services that are themselves not fully reliable – including the use of redundant components running in physically separated cloud data centers, and hot failover. Some CSPs build such reliability engineering into their offering, others require the CSC to install appropriate additional components to achieve the required results.

Finally, users 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 CSP's commitment as insufficient, even though it is better than what the existing on-premises solution offered.

Recommendations

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

- It is desirable that the CSP offer APIs, webhooks, an RSS Feed, a JSON feed or other electronic means of sending failure and alert data to the CSC's service management system. This enables the CSC 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 CSP (e.g., firewall changes by the CSP may prevent CSC users from accessing cloud services). CSCs must ensure that CSPs offer user interfaces, APIs, or other mechanisms to report failures *to* the CSP.
- The CSP should provide an Expected Time to Resolution (ETR) for any service failure, however detected.
- The CSC must be aware that the elapsed time between failure and recovery may exceed the advertised downtime without breaching the SLA, because the CSP may pause the SLA clock when it is awaiting information if needs from the CSC. Consequently, when a failure occurs, the CSC must mobilize itself in order to minimize such delays.
- CSCs should investigate whether the CSP supports resiliency features such as database replication, clustering with load balancing, and so on.
- CSCs should evaluate cloud services and options that can be used to make cloud applications and systems more resilient. Capabilities such as redundant systems, data replication and failover should all be considered.
- CSCs 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 CSP with a ticketing system that CSCs are allowed to access directly to report failures. This also makes it easier for CSCs to find out the status and ETR of the incident. Notifications by text messages or automated voice calls are also more likely to obtain immediate attention than e-mails.
- 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 CSCs assess MTBF, MTTR, and MTTF to determine expected service downtimes. Evaluate the impact of these downtimes against the nature of your workloads. Consider that the impact of failures may vastly exceed the service credits offered by the CSP, and consider the appropriate alternatives 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 CSP does not absolve CSCs 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. A comprehensive discussion of disaster recovery for cloud workloads [85] can help CSCs understand what a disaster recovery plan might include.

In general, current public CSAs offer 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 CSP 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 CSP’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 CSP makes no warranties that the CSC’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 CSP will not be liable for any deletion, damage or destruction of the CSC’s content.

Given the clauses above, all worded in the CSP’s favor, the onus is clearly on CSCs to define, implement and execute their own disaster recovery plans. Some CSPs explicitly offer capabilities to assist with this. 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 replicate data between those multiple sites in near-real time, and the ability to provision application instances across the sites, with load balancing between them, allowing services to fail over rapidly if one data center is subject to a disaster. In some cases, this is offered as a "Disaster Recovery as a Service"(DRaaS); in other cases, it is up to the CSC to organize the applications and services in an appropriate way to support disaster recovery.

If such a solution is considered, the locations of the multiple data centers should be reviewed to avoid conflicts with data residency requirements (see Step 5 above).

Recommendations

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

- CSCs 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.
- CSCs should ensure that business-critical content is stored redundantly in different geographical locations to help reduce the impact of a disaster. Popular solutions include only running business applications on top of cloud services that have built-in geographical redundancy, or leveraging replication technologies (offered by a third party or by the CSP) to synchronize the states of applications and data with a remote site.
- CSCs should clearly define the Recovery Point Objective (RPO) and Recovery Time Objective (RTO), the two most important metrics of disaster recovery, for each application. 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 period for which recent data updates might be lost due to a disaster; RTO is the maximum time until a business process is restored after a disaster).
- CSCs should ensure an appropriate frequency of backups based on the criticality of content, and should make sure that the backup location is not likely to being affected by the same disasters, and conversely that it is not a location that creates a data residency conflict.
- CSCs should use data and application replication capabilities where provided by the cloud service
- CSCs 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 CSP if needed.

Step 9: Develop an Effective Governance Process

CSCs 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. It may also be the case that an organization's governance framework does not specifically address cloud computing and needs to be updated. OMG's *Practical Guide to Cloud Governance* [17] gives advice and tools for establishing, modifying and sustaining cloud governance – inclusive of business and technical roles – and can facilitate the identification of where audits and monitoring are most necessary.

While the adoption of cloud and complex hybrid and multi-cloud solutions is growing, today's public CSAs contain few provisions for customer—provider management processes. The only formal channels of communication between the CSC and CSP specified in the service agreement are breach of contract type clauses (credit process, suspensions, termination, etc.). None of the agreements that were reviewed specify a commitment to status meetings between the parties. There is seldom a defined escalation process that the CSC can invoke to raise the priority of a service level issue. Where there is a defined escalation process, it is part of a premium support contract from the CSP or a value-added reseller (VAR).

Overall reporting and governance that includes elements such as change management and incident management remain infrequently described in the service agreements. As a result, CSCs 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 that fill the gaps to implement an effective management process. For example, the U.S. HIPAA regulation for healthcare contains the concept of a “business associate agreement,” which extends the obligations of a “covered entity” to its own suppliers. 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

An exit clause should be part of every CSA. It describes the details of the exit process, including the respective responsibilities of the CSP and CSC in case the relationship terminates – prematurely or not. CSCs 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 CSP.

In most cases, details of the exit process are contained in the Termination clause that is part of the Customer Agreement. All Termination clauses define two basic types of termination:

- *Termination for Convenience*. CSCs can typically stop using the cloud service at any time. Likewise, a CSP may terminate the agreement for convenience at any time without liability to the CSC. Advance notice – usually 30 days – is typically specified before termination occurs. In some cases, CSCs 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, such as when security violations are alleged, the CSP typically gives itself the right to suspend services *immediately* in order to protect itself and other CSCs, pending resolution or termination.

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

Cloud computing businesses are regularly acquired, and the CSP, CSC, cloud service broker, cloud service auditor, and cloud service carrier (the five roles defined in the *NIST Cloud Reference Architecture*) can all be affected by the change. The CSC might not be comfortable with the new policies and regulations being enforced by a new owner (which may be a geo-jurisdictional constraint due to the company’s place of business). Similarly, the CSC might not be comfortable if the acquirer is one of their competitors. These issues arise with all service models (IaaS, PaaS, SaaS) and even in the case of cloud hosting providers (for example, co-located data centers). These realities will cause terminations that can be initiated by the CSP or by the CSC. These terminations are not SLA-related, but highlight the importance of such proactive considerations.

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

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

Recommendations

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

- CSCs should ensure their agreement specifies that advance notice will be given for all terminations initiated by the CSP (minimum of 30 days).
- CSCs 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, CSPs should offer assistance to CSCs 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 CSC 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 CSC sufficient time to verify that all data has been correctly migrated to a new service.
    - Only with the CSC's written agreement should data be removed and destroyed before that time.
- At the completion of the exit process, CSCs should receive written confirmation from the CSP that all of the CSC's data, including analytical and statistical information derived from it, has been completely removed from the CSP's systems.

Conclusion

The CSA landscape continues to evolve. While some agreements are still rudimentary in terms of assurances offered to CSCs, it is encouraging to see that more and more CSPs 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, CSPs will continue to offer more specific terms 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 CSPs. 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 or recent 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) [5] provide hope for greater consistency of the terminology used to define service level objectives.

In the meantime, CSCs must carefully evaluate the materials provided about each cloud service they are considering. The recommendations outlined in this document should enable CSCs 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*. https://ww.axelos.com/best-practice-solutions/itil

[2] Cloud Standards Customer Council (2017). *Practical Guide to Cloud Computing, Version 3.0.* https://www.omg.org/cloud/deliverables/CSCC-Practical-Guide-to-Cloud-Computing.pdf

[3] Cloud Standards Customer Council (2017): *Security for Cloud Computing: 10 Steps to Ensure Success, V3.0.* https://www.omg.org/cloud/deliverables/security-for-cloud-computing-10-steps-to-ensure-success.htm

[4] Cloud Standards Customer Council (2016). *Practical Guide to Hybrid Cloud Computing*. https://www.omg.org/cloud/deliverables/practical-guide-to-hybrid-cloud-computing.htm

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

[6] International Organization for Standards (2014). *ISO/IEC 17789:2014 – Cloud Computing – R Reference Architecture.* https://www.iso.org/standard/60545.html

[7] International Organization for Standards (2016). *ISO/IEC 19086-1, Service Level Agreement (SLA) Framework – Part 1: Overview and Concepts.* https://www.iso.org/standard/67545.html

[8] International Organization for Standards (2018): *ISO/IEC 19086-2: Service Level Agreement (SLA) Framework -- Part 2: Metric model.* https://www.iso.org/standard/67546.html

[9] International Organization for Standards (2017): *ISO/IEC 19086-3: Service Level Agreement (SLA) Framework -- Part 3: Core Conformance Requirements*. https://www.iso.org/standard/67547.html

[10] International Organization for Standards (2019): *ISO/IEC 19086-4: Service Level Agreement (SLA) Framework -- Part 4: Components of security and of protection of PII*. https://www.iso.org/standard/68242.html

[11] International Organization for Standards (2017): *ISO/IEC 19944: Cloud services and devices: Data flow, data categories and data use*. https://www.iso.org/standard/66674.html

[12] International Organization for Standards (2014): *ISO/IEC 27018: Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors*. https://www.iso.org/standard/61498.html

[13] National Institute for Standards and Technology (2011): *NIST Cloud Computing Reference Architecture.* Special Publication 500-292. https://www.nist.gov/customcf/get_pdf.cfm?pub_id=909505

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

[15] National Institute for Standards and Technology (2018): *Risk Management Framework Update*. Special Publication 800-37 Rev. 2. https://csrc.nist.gov/publications/detail/sp/800-37/rev-2/final

[16] Object Management Group (2019). *Practical Guide to Cloud Service Agreements, Version 3.0*. https://www.omg.org/cloud/deliverables/practical-guide-to-cloud-service-agreements.htm

[17] Object Management Group (2019). *Practical Guide to Cloud Governance.* https://www.omg.org/cloud/deliverables/practical-guide-to-cloud-governance.htm

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

Cloud Service Agreements

[19] Acquia Cloud Free Agreement: https://www.acquia.com/sites/acquia.com/files/documents/2018-04/AgreementACF.pdf

[20] Amazon EC2 Service Level Agreement: https://aws.amazon.com/compute/sla/

[21] Amazon S3 Service Level Agreement: https://aws.amazon.com/s3/sla/

[22] Amazon Web Services Acceptable Use Policy: https://aws.amazon.com/aup/

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

[24] Amazon Web Services Customer Agreement: https://aws.amazon.com/agreement/

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

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

[27] AppRiver Terms of Subscription: https://www.appriver.com/terms-of-subscription/

[28] AT&T Acceptable Use Policy: https://www.att.com/legal/terms.aup.html

[29] AT&T Privacy Policy: https://about.att.com/csr/home/privacy.html

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

[31] Aveva Cloud Services Agreement: https://sw.aveva.com/legal/cloud-services

[32] BlueHost Terms of Service: https://www.bluehost.com/terms

[33] Centurylink Privacy Policy: https://www.ctl.io/legal/privacy-policy/

[34] Dell Cloud Solutions Agreement: https://www.dell.com/learn/us/en/uscorp1/terms-conditions/art-cloud-solutions-agreement

[35] Dimension Data Privacy Policy: https://www.dimensiondata.com/en/privacy-policy

[36] Dimension Data Public CaaS Service Level Terms (Americas): https://www.dimensiondata.com/-/media/dd/corporate/content-images/pdfs/legal/2-public-caas-service-level-terms-2015.pdf?la=en

[37] Dropbox Security & Privacy certifications: https://help.dropbox.com/accounts-billing/security/standards-regulations

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

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

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

[41] Google Apps Service Level Agreement: https://gsuite.google.com/intl/en/terms/sla.html

[42] Google Cloud Storage, Google Prediction API and Google BigQuery SLA: https://cloud.google.com/storage/sla

[43] IBM SoftLayer Master Services Agreement: http://static.softlayer.com/sites/default/files/assets/page/softlayer_msa.pdf

[44] IBM Cloud Support Plans: https://cloud.ibm.com/docs/get-support?topic=get-support-support-plans

[45] IBM Terms of Use – SaaS Specific Offering Terms (IBM Kenexa Talent Insights): https://www-03.ibm.com/software/sla/sladb.nsf/pdf/6937-01/$file/i126-6937-01_06-2015_en_US.pdf

[46] IBM Kenexa LMS on Cloud – Service Description: http://www-03.ibm.com/software/sla/sladb.nsf/pdf/6511-06/$file/i126-6511-06_11-2018_en_US.pdf

[47] Microsoft Azure Agreement: https://azure.microsoft.com/en-us/support/legal/subscription-agreement/

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

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

[50] Microsoft Trust Center – Privacy at Microsoft: https://www.microsoft.com/en-us/trust-center/privacy

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

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

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

[54] Oracle Privacy Policy: https://www.oracle.com/legal/privacy/

[55] Oracle Cloud Service Contracts: https://www.oracle.com/corporate/contracts/cloud-services/

[56] Progress Sitefinity End User License Agreement: https://www.progress.com/legal/license-agreements/sitefinity

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

[58] Rackspace Acceptable Use Policy: https://www.rackspace.com/information/legal/global/aup

[59] Salesforce Master Subscription Agreement: https://c1.sfdcstatic.com/content/dam/web/en_us/www/documents/legal/salesforce_MSA.pdf

[60] Salesforce Data Processing Addendum (GDPR, Salesforce Processor Binding Corporate Rules, Privacy Shield, and Standard Contractual Clauses): https://a.sfdcstatic.com/content/dam/www/ocms-backup/assets/pdf/misc/data-processing-addendum.pdf

[61] Salesforce Security, Privacy and Architecture: https://www.salesforce.com/content/dam/web/en_us/www/documents/legal/misc/salesforce-security-privacy-and-architecture.pdf

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

[63] Salesforce Privacy Statement: https://www.salesforce.com/company/privacy/full_privacy/

[64] Salesforce.com Premier Success Plans: https://a.sfdcstatic.com/content/dam/www/ocms-backup/assets/pdf/datasheets/DS_SuccessPlans.pdf

[65] SAP Cloud Services Agreements: https://www.sap.com/about/trust-center/agreements.html#cloud-services-agreements

[66] SAP Privacy Statement: https://www.sap.com/about/legal/privacy.html

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

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

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

[70] VMWare Privacy Notice: www.vmware.com/help/privacy.html

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

Papers and Articles

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

[73] 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

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

[75] 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. https://dl.acm.org/citation.cfm?doid=1655008.1655020

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

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

[78] Golden, Bernard: *Cloud Computing: The Truth About What Runs on Amazon.* CIO, September 2010. https://www.cio.com/article/2414854/cloud-computing--the-truth-about-what-runs-on-amazon.html

[79] 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. https://dl.acm.org/citation.cfm?doid=1555336.1555341

[80] Magalhaes, Ricky M. and Monique L. (Oct.-Dec. 2014): *Selecting Cloud Management Platforms.* https://lenta2016blog.wordpress.com/2014/10/31/selecting-cloud-management-platform-part1/ and https://lenta2016blog.wordpress.com/2014/12/08/selecting-cloud-management-platform-part2/

[81] NTT America (2012): *An Evaluation Framework for Selecting an Enterprise Cloud Provider.* http://i.zdnet.com/whitepapers/NTT_Cloud_Evaluation_Framework.pdf

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

[83] Ponemon Institute (2018): *Closing the Cloud Security Gap*. http://s3.amazonaws.com/idgcampaigns/documents/uploaded_data/05c/4f2/1a-/original/closing-the-cloud-security-business-gap.pdf?1542300161

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

[85] 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). https://ieeexplore.ieee.org/document/7266853

Appendix A – Reseller and Business Partner Roles

In this Appendix, related to Step 1 (Understand Roles and Responsibilities), we examine various types of relationships between a cloud reseller and the CSP or CSPs whose services they “wrap” and offer to the CSCs. We distinguish five models:

- the subcontracting model,
- the assignment model,
- the agency model,
- the referral model,
- the orchestrator model.

For each of the five models, we present the following in the table that starts on the next page:

- its definition,
- its advantages and drawbacks,
- who are the parties to the contract (there are often several alternatives),
- who is liable for what,
- who is responsible for providing certifications that may be required (e.g., from whom can the CSC obtain a SOC II type 2, ISO 27001, ISO 27012, FedRAMP, etc., certification?)

Appendix A Table – Reseller and Business Partner Roles

	Subcontracting model	Assignment model	Agency model	Referral model	Orchestrator model
Definition	CSP contracts with Reseller, who in turn contracts with CSC.	Reseller *resells contracts* – like traditional software reseller licensing but adapted to a service model. Reseller *assigns* its rights and obligations under the contract with the CSP to the CSC.	Reseller acts as an *agent* of CSP. Software/services of reseller are in a separate customer agreement with reseller. "Law of agency" allows resellers to enter into contracts on behalf of CSP.	Reseller *refers* CSC to CSP. (Not really a reseller as it does not sell or resell services, but an important element in the marketplace.	Reseller "front-ends' multiple CSPs in an integrated orchestration of services through a single point of contact to the CSC. The orchestration may include classic CSPs or highly specialized services to fit the CSC's needs.
Advantages	Flexibility. Shared Risk. Pricing not fixed. Sub-resellers. White labelling (can be a drawback too).	Support for ancillary contracts. Reseller can "step away" after a period or at end of engagement. Sub-resellers. Contractual "wiggle room" unlike agency model.	Clarity about whom one is working with. CSP is directly contracted with CSC. CSP carries risk. Commercial agent regulations may not apply as those are written for "products." Reseller may provide better cloud support options to CSC.	Direct relationship with CSP. Special discounts may apply. Reseller is usually not contracted directly. CSP carries risk Reseller may represent multiple CSPs.	Highly adaptable Risk-balanced (when applied through assignment or subcontracting). Open and standards-based. Cost-balanced (migration/interoperable). Leverage multiple levels of support. Get more with less.

	Subcontracting model	Assignment model	Agency model	Referral model	Orchestrator model
Drawbacks	Agreements can become complex. Reseller and CSP may breach their mutual agreement. CSP may be hidden.	Does not fit most CSPs. Rights and obligations are a "one-off." Agreements can become inflexible when not used for a specific engagement (such as a migration or interoperability transfer/migration project)	Fiduciary requirements on reseller are in the best interest of CSP, not CSC.		Highly technical. Contract-rich. Reseller must have expertise and understanding in every offering and use case. Susceptible to configuration errors and time to perform the configuration (interoperability takes time the first time).
Contract Types	CSP—Reseller Reseller –Customer Back to back/back to front CSP—Reseller—Customer	CSP—Reseller Reseller—Customer (under rights and obligations defined by CSP)	CSP—Customer Reseller—Customer (for specialized software/services or add-ons)	CSP—Customer	CSPs—Reseller—Customer
Liability	Reseller Reseller & CSP (when back to back/front) Reseller and CSP (preferred)	CSP Some borne by reseller if scoped by a project Statement of Work	CSP – for underlying cloud services Reseller – for add-ons	CSP	CSPs and Reseller
Certification Handling	CSC should request certificates from both. Reseller may offer certificates that are outdated or do not represent the entire CSP ecosystem.	CSP Reseller when engaged (scoped by a Statement of Work)	CSP Reseller (when engaged for additional services)	CSP	CSPs and Reseller – a must

Appendix B – 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 CSP.	"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." "Customer agrees not to, and not to allow third parties to use the Services: to use the Services, or any interfaces provided with the Services, to access any other CSP product or service in a manner that violates the terms of service of such other CSP product or service."
"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 C – 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 CSC to take specific action: • CSC 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. • CSC 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. • CSP 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 CSP's reasonable control. • Force majeure conditions. • Any actions or inactions of the CSC or any third party resulting in the outage. • CSC and/or third-party equipment, software or other technology contributing to the failure. • CSC's refusal to allow CSP to perform maintenance deemed necessary to maintain the cloud 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 D – 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 CSP
  - 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 CSP

Earlier work by the NIST Cloud Audit subgroup identified 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 E – 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 CSC is responsible for protecting the CSP, and must notify the CSP in case of breach, but not the other way around. A few CSPs commit to informing the CSC promptly in case of a security breach, and to provide all information available to them about what happened. Some CSPs, 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	CSPs assume no responsibility for "making the CSC whole" if there is a breach for which they are responsible. Some CSPs include unspecific assurances that they will assist the CSC. Most CSPs 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 CSPs 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 CSPs 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 F – Privacy

This table contains key observations and actual language examples about key privacy issues. This is an area undergoing rapid evolution, with most organizations still in the process of understanding and addressing the consequences of the promulgation of the European Union's GDPR in May 2018. Data protection by social media providers has also become a key concern of society and legislative bodies, and continuing evolution of laws and regulations can be expected.

Subject	Key Observations	Example Language
Information collected about the CSC	Most agreements specify in some detail the kind of information collected by the CSP about the CSC 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 CSP 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 CSP	Many SaaS applications (collaboration, CRM, ERP, Web conferencing, etc.), as well as IaaS storage services, will result in personal information about the CSC's own customers, employees, suppliers, etc., being held by the CSP. Yet most agreements make no mention of any protection given to that data. In some cases, the agreement spells out that the CSC needs to protect its own customers, even though it doesn't say that the CSP 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 CSP 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."