Practical Guide to Cloud Service Agreements

Version 3.0

A Discussion Paper from the OMG Cloud Working Group

February 2019

Document mars/2019-02-01

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

Acknowledgements

The major contributors to the successive versions of this whitepaper have been:

- Claude Baudoin (cébé IT & Knowledge Management)
- Christian Boudal (IBM)
- Hannah Day (Mayo Clinic)
- Beniamino Di Martino (Second University of Naples)
- Chris Dotson (IBM)
- Marlon Edwards (Hoboken Consulting Group)
- Mike Edwards (IBM)
- Dominick Grillas (Damo Consulting)
- David Harris (The Boeing Company)
- Ryan Kean (The Kroger Co.)
- Edwin Lang (Ontario Ministry of Health)
- Yves Le Roux (CA Technologies)
- George Malekkos (Powersoft Computer Solutions Ltd)
- John McDonald (CloudOne Corporation)
- John Meegan (IBM)
- Gerry Murray (Fort Technologies)
- Paddy Padmanabhan (Damo Consulting)
- Massimiliano Rak (Second University of Naples)
- Dave Russell (IBM)
- Karolyn Schalk (IBM)
- Lisa Schenkewitz (IBM)
- Anil K. Sharma (IBM)
- Prasad Siddabathuni (Edifecs)
- Gurpreet Singh (Ekartha)
- Annie Sokol (NIST)
- Joe Talik (AT&T)
- Alex Tumashov (Schlumberger)
- Salvatore Venticinque (Second University of Naples)
- Steven Woodward (Cloud Perspectives)
- John Wooten (CONSULTED)

Revisions

The issues addressed in cloud computing service agreements, and the ways in which they are addressed, have continued to evolve since the Cloud Standards Customer Council (predecessor to the OMG's Cloud Working Group) first published the *Practical Guide to Cloud Service Level Agreements* white paper in April 2012. We addressed some changes in Version 2.0 in 2015. Version 3.0 continues this effort. In particular:

- All ten steps in the *Guide for Evaluating Cloud Service Agreements* section have been updated to reflect current best practices.
- References to cloud computing standards have been updated.
- References have been added to several CSCC white papers published in 2016—2018.
- All references have been consolidated into a single section and updated.

Introduction and Executive Summary

The *Practical Guide to Cloud Service Agreements* provides a practical reference to help enterprise information technology (IT) departments and business decision makers analyze cloud service agreements (CSAs) proposed by cloud service providers. The paper informs decision makers of what to expect and what criteria to use as they evaluate CSAs from multiple potential suppliers.

CSAs are primarily written to set clear expectations for services between the Cloud Service Customer (CSC – the buyer) and the Cloud Service Provider (CSP – the seller), and this Guide focuses on that relationship. However, CSAs should also exist between a customer and the other cloud computing roles and sub-roles shown in **Figure 2**, such as the cloud network provider, the cloud service broker and even the cloud auditor.

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

There may be different requirements for the content of a CSA according to the service delivery model selected: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS). In this Guide, we focus on the requirements that are common across the various service models, and point out the differences when appropriate.

"The Current CSA Landscape" section explains the dynamics that currently exist between cloud customers and providers, and the impact that company size has on the power to negotiate terms. This section also highlights the nuances of CSA development for different service models.

The "Guide for Evaluating Cloud Service Agreements" section is the heart of the paper. It provides a prescriptive series of ten steps that CSCs should take to evaluate CSAs in order to effectively compare multiple CSPs or to negotiate terms with a selected CSP. This section takes into account the realities of today's cloud computing ecosystem and postulates how it is likely to evolve, including the important role that standards will play to improve interoperability and consistency across providers.

The CSA is often the best indicator of how, and how often, the provider expects their service to *fail*. Therefore, CSCs must remember that downtime, poor performance, security breaches and data losses are ultimately their risks to bear. It's important that CSCs select CSPs who will help them with the fine details in supporting their workloads as they transition to cloud computing.

A related document, the *Public Cloud Service Agreements: What to Expect and What to Negotiate* [1], provides additional details on evaluating CSAs from prospective public cloud providers.

The Current CSA Landscape

CSAs are a set of documents or agreements that contain the terms governing the relationship between the CSC and the CSP. Because the cloud computing market is rapidly evolving and maturing, CSCs should be aware that there may be a mismatch between their expectations and the service terms actually offered by the CSPs. For example, a CSA may not specify the geographic location where customer data will be stored. This could be a showstopper for customers subject to export restrictions of certain types of data from the U.S., or the export of "personal data" from the European Economic Area (EEA), where the General Data Protection Regulation (GDPR) is applicable. Additional security- and privacy-related standards and considerations are discussed at length in three separate papers: *Security for Cloud Computing: 10 Steps to Ensure Success* [2], *Cloud Security Standards: What to Expect and What to Negotiate* [3], and *Data Residency Challenges* [4].

It is common for disputes to arise over the structure of the agreements, thus CSCs must pay close attention to the language and clauses of the CSA. Large CSPs can be inflexible with their CSAs, while small CSPs may seem more flexible, but tend to over-promise in order to win clients. In the end, it is the responsibility of the CSC to assess the CSAs and supporting documents and determine if their needs can be met by the services described in these documents.

In general, the CSA is comprised of three major artifacts:

- *Customer Agreement*
- *Acceptable Use Policy (AUP)*
- *Service Level Agreement (SLA)*

This classification is not complete, nor is there a standard nomenclature adopted by all cloud service providers to specify their CSAs. Furthermore, CSPs can modify their contract structure and terms at any time.

The *Customer Agreement* section of the CSA describes the overall relationship between the CSC and CSP. Since service management includes the processes and procedures used by the CSP, explicit definitions of the roles, responsibilities and execution of processes need to be formally agreed upon. The Customer Agreement fulfills this need. Various synonyms such as "Master Agreement," "Terms of Service," or simply "Agreement" may be used by certain CSPs.

An *Acceptable Use Policy (AUP)* is commonplace within a CSA. The AUP prohibits activities that CSPs consider to be an improper or outright illegal use of their service. This is one area of a CSA where there is considerable consistency across CSPs. Although specific details of acceptable use will vary between IaaS, SaaS and PaaS providers, the scope and effect of these policies are the same, and these provisions typically generate the least concerns or resistance.

A typical *Service Level Agreement (SLA)* within the CSA describes levels of service using various attributes such as availability, serviceability or performance. The SLA specifies thresholds and financial penalties

associated with failure to meet these thresholds. Well-designed SLAs include escalation procedures to ensure that issues are resolved through proper governance before they escalate into a conflict.

To guarantee an agreed service level, CSPs must measure and monitor relevant metrics. There is often a mismatch between the metrics collected and monitored by the CSPs and the higher-level functional (or "end-to-end") metric relevant to CSCs. This issue is common across service models, but is more acute for SaaS since customers want service levels to be met at the application level where they can be impacted by many factors, such as network connectivity. This is one reason why CSAs for SaaS usually lack stringent service level guarantees.

Service level guarantees for IaaS are better defined than for SaaS or PaaS, but that does not mean that they always meet the customer's expectations. Most public cloud infrastructure services are available only through non-negotiable standard contracts which strictly limit the provider's liability. As a result, the remedies offered in case of non-compliance do not match the cost to the customer of the potential service disruptions. Furthermore, most IaaS providers put the burden of SLA violation notification and credit request on their customers.

In many cases, cloud SLAs do not offer refunds of charges but rather service credits against future use. Whether the relief is in the form of a credit or a refund, it is usually subject to a cap such as one month's standard billing. Credits against future billing will be of little or no benefit to CSCs that decide to switch providers following unsatisfactory service – and they clearly are meant to encourage the customer to stay with the current CSP.

This rather biased situation is starting to evolve. As CSCs become more knowledgeable and competition increases, CSPs are beginning to offer different service options that better shield CSCs from such risks.

The GDPR includes interoperability requirements, such that unsatisfied CSCs can be more easily "ported" to a different CSP, lessening risks associated with CSP "lock-in."

For CSCs, size also matters. In general, the larger the customer deployment, which translates to higher setup and monthly fees, the more power the customer can exert in negotiating more favorable CSAs, even with SaaS providers. No such improvements may be offered to small and medium businesses, but over time we expect the changes imposed by larger customers to trickle down to smaller ones. Better CSAs are a competitive factor, with several options and pricing plans that try to address some of the CSCs' concerns.

Guide for Evaluating Cloud Service Agreements

Before getting to the point of evaluating any CSA, CSCs must first perform a number of strategic steps (develop a comprehensive business case and strategy, select cloud service and deployment models, etc.) that are detailed in the *Practical Guide to Cloud Computing* [5].

With this strategic analysis as a prerequisite, this section provides a prescriptive series of ten steps that should be taken by CSCs to evaluate CSAs in order to compare multiple CSPs or to negotiate terms with a selected CSP. The following steps will be 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**

Requirements and best practices are highlighted for each step. In addition, each step takes into account the realities of today's cloud computing landscape and postulates how this space is likely to evolve in the future, including the important role that standards will play to improve interoperability across CSPs and the ability to compare them.

Within each step, the expected responsibilities of the customer and the provider in meeting both business level and service level objectives are explained. In order to make sound business decisions, CSCs need to understand what they can expect from their CSPs, and view these expectations from the perspective of their organization's policies, standards and overall risk posture. Additionally, CSCs should understand that they remain the owners of their own data – they may delegate its custody, but not its ownership – as they will be held accountable for that data by their management, their customers, and increasingly by the law. This, in turn, will help them clarify their own responsibilities and assess the true costs and benefits of moving to cloud computing.

Step 1: Understand Roles & Responsibilities

From the CSC perspective, one of the significant areas of risk involved with cloud computing is associated with the division of activities and responsibilities between the CSC and the CSP. A clear understanding of what activities are within the scope of the service ("in-scope") provides an opportunity for the CSC and/or CSP to fill the gap (perhaps with an additional service) and reduce the risk of customer satisfaction issues. One way to assess gaps is to review the capabilities required by the

appropriate cloud service model (SaaS, PaaS or IaaS) as defined by publication *NIST-SP-800-145: The NIST Definition of Cloud Computing* [6]:

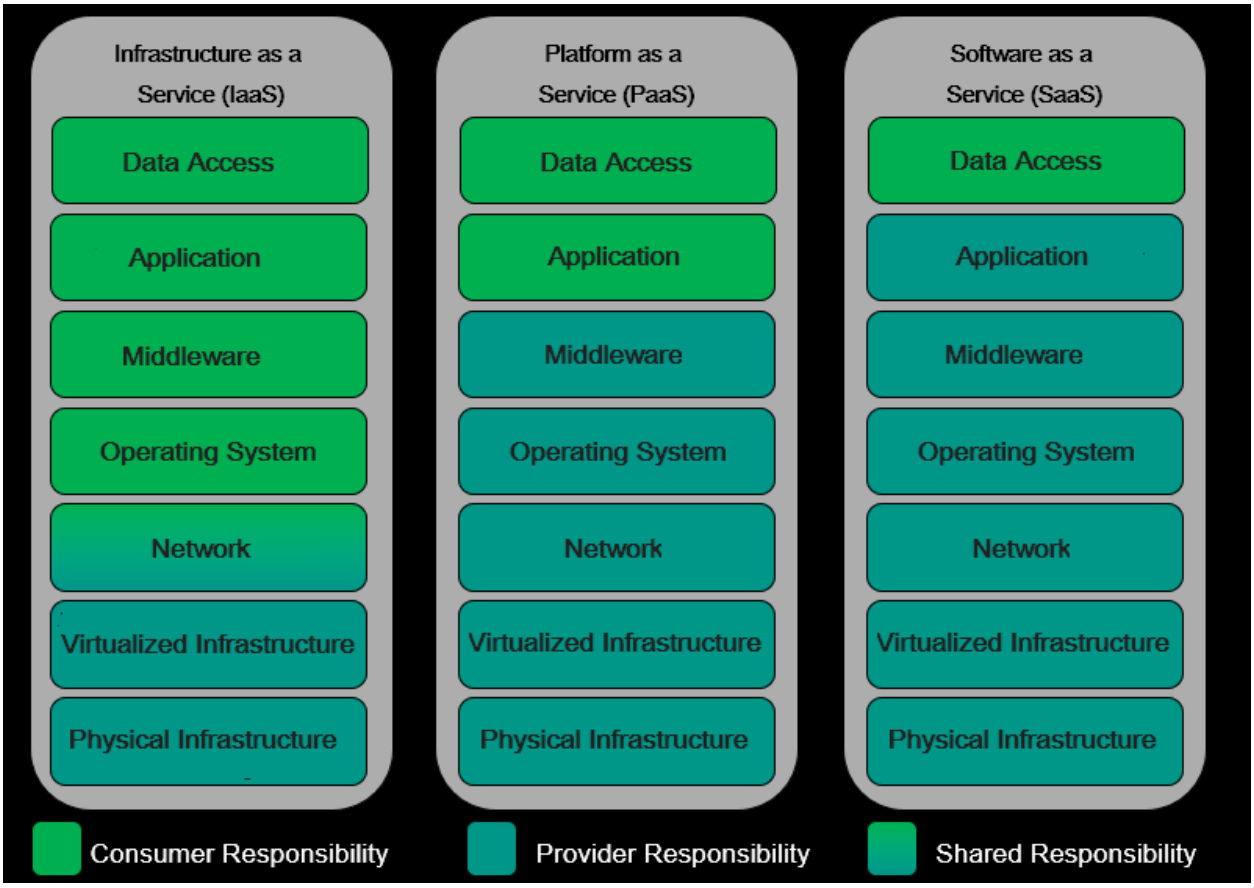

Figure 1 -- Responsibility Allocation by Cloud Service Model

The *ISO/IEC 17789:2014 standard (Cloud Computing – Reference Architecture)* [7] defines three main roles for cloud computing:

- Cloud service customer (CSC)
- Cloud service provider (CSP)
- Cloud service partner (CSN)

The CSP and the CSC are the most significant roles in the provision and use of cloud services, while the CSN is a party engaged in support of the activities of the CSC and/or the CSP.

There are a number of sub-roles of each of the major roles – those sub-roles are shown in Figure 2.

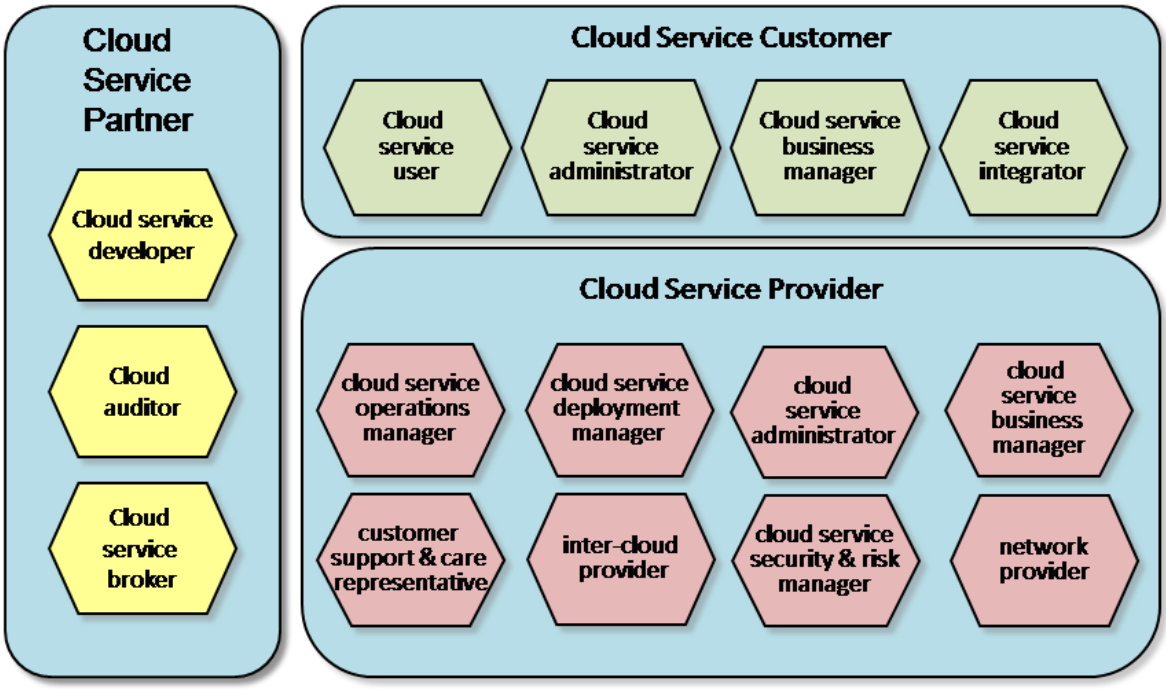

Figure 2 Cloud Computing Roles and Sub-Roles

Each of the sub-roles in Figure 2 has a set of activities and responsibilities, which are described in high-level terms in ISO/IEC 17789. There are also relationships between the sub-roles – for example, the CSC's cloud service administrator may interact with the CSP's customer support and care representative when customer personnel experience problems using the cloud service.

Some of the sub-roles may be explicitly mentioned in a CSA, or they may have a direct or indirect relationship to some aspects of the CSA. The sub-roles of the CSC and of the CSP, in particular, are involved in the delineation of responsibilities that is typical for cloud services – the CSA should make clear statements about those responsibilities. CSCs need to understand the activities and responsibilities of the various sub-roles.

One important area for customers to consider is who is responsible for detecting and then reporting incidents where the cloud service fails to meet some aspect of the CSA or SLA. This includes outages where the cloud service is unavailable, or cases where performance fails to meet stated service levels (for example, response times are too long). It may be the responsibility of the CSC to detect and report such problems, in which case the customer may need to put in place appropriate monitoring technology. However, some CSPs do not allow monitoring technologies to be deployed in their cloud as a matter of standard policy. In such instances, the CSC may attempt to negotiate an acceptable alternative. In any case, it is necessary to be clear about how incidents are then tracked until resolved (see Step 6 on Service Management).

One partner role that is particularly relevant to the CSA and to SLAs is the cloud auditor. It is unlikely that the CSC has direct insight into the operations of the CSP, particularly regarding aspects such as security and the protection of sensitive data such as personally identifiable information (PII). It is typical

for CSPs to offer assurances about these aspects of their cloud services through third-party certifications or attestations. Those are delivered by third party cloud auditors who inspect the cloud service provider's operations at a predetermined frequency; their audit reports are typically based on one or more standards or certification schemes.

Each CSA may be unique based upon the customers' requirements and the cloud services under consideration. CSAs can contain various elements that are not limited to quantitative measures, but can include such qualitative aspects as alignment with standards and data protection. It is strongly recommended that CSCs gain a solid understanding of the spectrum of CSAs currently proposed by CSPs in order to compare the offerings of different providers and assess the tradeoffs between cost and service levels. Refer to the CSCC whitepaper *Public Cloud Service Agreements: What to Expect and What to Negotiate* [1] for details.

We saw in Figure 1 how that the split of responsibilities between the CSP and the CSC is going to differ according to the service model, and this is necessarily reflected in differences in the CSA and SLA. Irrespective of the type of cloud service being evaluated, the CSC should document and prioritize (e.g., must haves, nice to have, etc.) their requirements (functional and non-functional) and expectations.

Step 2: Evaluate Business-Level Policies

When reviewing proposed CSAs, CSCs must consider the policy and compliance requirements relevant to them since there are interdependencies between the policies expressed in the CSA and the business strategy and policies developed across the lines of business or by the Legal or Compliance department. The *data policies* of the CSP, as expressed in the CSA, are perhaps the most critical business-level policies that should be carefully evaluated.

The obligations a CSP has to its customers and their data is governed by a potentially complex combination of:

- CSC requirements
- Any special/specific data handling agreement that they may have with the given CSC
- The data protection legislation applicable to the customer as well as to its individual users (which may not be under the same jurisdiction in a multinational company)
- The laws and regulations applicable where the data resides or is made available
- Third-party use of the data, including the right to resale data
- Indemnification, what a CSP will allow or disallow according to their level of risk tolerance.

CSCs should carefully consider these legal requirements and how the CSA deals with issues such as movement of data when redundancy across multiple sites means subjecting the data to different jurisdictions. The issue of jurisdiction takes on additional complexity when global compliance is taken into consideration and more than one CSP is used. In these instances, the CSC may have to coordinate negotiations between CSPs to ensure continuity of data management. Both the CSC and CSP should view the issue of legal and regulatory compliance as an ongoing activity – one that may require updates to agreements – since privacy and residency laws and regulations are constantly shifting.

Table 1 highlights the critical data policies that need to be considered and included in the cloud CSA.

Table 1—CSA Data Policies

Data Policy	Description / Guidance
Data Preservation and Redundancy	• Timely and efficient capture and preservation of data is critical to maintaining the organizational memory of a business or the general user. CSCs should therefore ensure they have an appropriate data preservation strategy that addresses redundancy. • CSCs should ensure the CSA supports their data preservation strategy that includes sources, scheduling, backup, restore, integrity checks, etc. They should be concerned as to the protections offered or omitted by the CSP, and, where necessary, be prepared to augment the services provided by their CSP with additional services. • It must be possible to test the CSA to demonstrate the required level of service availability.
Data Residency	• CSAs that cover locations under different jurisdictions are challenging. See the CSCC white paper on *Data Residency Challenges* [4] for details. • CSCs should consider how the CSA specifies where their data resides, where it is processed, and how this meets their business requirements as well as the various applicable regulations. CSCs should also understand where the data is viewed or delivered, and whether this results in a transborder data flow with regulatory or tax implications.[1] • For example, can the CSP truly deliver a sound technical solution when sensitive data spans several jurisdictions with conflicting laws? Does the CSP commit, in the CSA, to the specific location(s) where the customer's data will be stored? • If the CSP reserves the right to add new locations, relocate or copy data to a different location, or change data movement policies, will they give the CSC advance notice? Preferably, will they obtain the CSC's permission to relocate its data? • Is there a means to verify the current location of a data set? Can this be done on a periodic basis – perhaps as part of an audit?
Data Seizure	• Law enforcement authorities and other government agencies have the right to seize data under certain circumstances. CSCs should ensure that the CSA provides for sufficient notification of such events. • in the event that the CSP locks access to its systems because of a billing dispute or a security issue, the CSC's data should not be "held hostage" while the issue is being resolved.

[1] To ensure regulatory compliance, the CSA should address the CSC's "right to audit," typical within a competent CSP security audit. For further information, please refer to Step 5 (Evaluate Security and Privacy Requirements), Table 6 (Key Security Considerations for CSAs), *Audit the Cloud Provider's Security CSA Compliance*; and Step 6 (Identify Service Management Requirements), *Auditing.*

Data Policy	Description / Guidance
Data Privacy	• The CSP's data privacy policy should be included in the CSA, and should ensure that the CSP will conduct business in compliance with applicable laws on data protection. • This includes identifying the data sets gathered, data retention policies, how the data is communicated, how personal data is stored and used, etc. • Data privacy in a cloud context is not just about the protection of the information about the CSC's agents in its dealing with the CSP (this is the narrow meaning in many existing SLAs); it also includes the protection of the information that may be stored about the CSC's own customers. • Refer to the Privacy section within Step 5 for more information.
Data Security	The CSC should assess the data security controls offered by the CSP, such as: • Encryption of data at rest • Measures taken to prevent unauthorized user access – including limitations on data access by the CSP's own system administrators • Database encryption • Two- factor/Multi-factor authentication (2FA/MFA) • Monitoring of data access and notification of breaches • Sharing of access logs with the CSC
Data Availability	• Assess whether the CSP's maintenance schedules might interfere with the CSC's business processes that are subject to external constraints, such as financial reporting or the business's hours of operation in certain regions. • CSCs should also ensure there are arrangements in place, include any necessary funding sources, to make their data available in the event that the CSP goes out of business. This may require the purchasing of some form of insurance. • CSCs should assess any "uptime guarantee" provided by the CSP to ensure that those guarantees meet their business needs. A term such as "regular business hours" is imprecise given time zones, regional holidays and week-end observance. A multinational CSC with offices on all continents may require support 24x365. • CSCs should review any liability and indemnification clauses related to data loss or corruption. CSP policies and risk tolerance levels vary greatly. A CSP may provide no coverage, coverage limited to the customer's expense for the faulty service, or up to a specific amount, likely to be minimal compared to the actual impact. A CSC may try to renegotiate these clauses or buy insurance to offset the risk to the business.

Data Policy	Description / Guidance
Change Management and Notification	• The change management and change notification obligations of the CSP should be carefully reviewed, especially the amount of time allowed to prepare for a change. The CSP may also ask the CSC to provide certain change notifications (e.g., a large expected jump in usage), which is a good opportunity to strengthen the customer's own change management policies. • The CSC should assess whether the CSP provides a "sandbox" environment to assess the impact of the proposed changes to their applications and data.

All of these policies impact and influence the CSC's cloud strategy and business case. In many cases, CSA policies are similar across different CSPs and presented as non-negotiable. Since this is not always the case, CSCs should refer to the already cited *Public Cloud Service Agreements: What to Expect and What to Negotiate* [1] for help in improving the CSA when possible.

We now move from data policies to business policies. Table 2 below highlights the critical business-level policies that need to be addressed in the CSA.

Table 2 -- CSA Business Level Policies

Policy	Description / Guidance
Guarantees	• CSA guarantees should be defined, objective and measurable with an appropriately scaled penalty matrix that matches the impact of non-performance by the provider.[2] The CSA should clarify: ▪ What constitutes excused or excluded performance ▪ Escalation procedures ▪ How service-level bonuses and penalties are administered ▪ Remedy circumstances and mechanisms • Guarantees should be expressed as measurable numbers, but the measurement window should also be specified. An availability percentage of 99.99% has a very different effect if it is measured over one year and 0.01% of downtime (almost one hour) occurs at a critical time for the business, as opposed to a reference period of one day. Other guarantees will be expressed in other units, such as time-to-repair in minutes, etc. • What qualifies as an incident varies across CSPs and CSCs should watch for subtle limitations. For example, a CSP may consider a service "unavailable" after 5 minutes of continual outage – so that there may be ten 3-minute halts within a day and the CSP will claim that availability was 100%. The CSC must therefore understand exactly how a metric is defined and measured, and what impacts it has to the CSC's business. • The CSA may also place the onus of reporting issues on the CSC, in which case the CSC needs to understand whether they have the tools to ensure detection and provide timely notification.

[2] Guarantees including measurable metrics will be covered in greater detail in the sections that follow.

Policy	Description / Guidance
Acceptable Use Policy (AUP)	• The AUP will clearly describe how the CSC may use the service and the agreement generally will describe what actions the CSP may take in the event of a breach. • In today's cloud environment, this policy is typically non-negotiable and the terms generally favor the CSP. • CSCs need to understand the impact of such policies if they use the cloud solution to in turn provide a service to end users over whom they have limited control. • Verify that there are no clauses allowing the CSP to resale the customer's data (even if redacted or anonymized) to third parties for any purpose.
List of Services Not Covered	• The CSA will state under what conditions and with which described services the CSC is supported. The CSA may also state what is excluded and what constitutes illegal use. • CSCs should look for explicitly stated exceptions and understand why the CSP has excluded them, and whether these exclusions may be harmful to the CSC's business.
Excess Usage	• CSPs operate to make money. While elasticity is a fundamental benefit of using the cloud, CSPs may charge high incremental rates for usage above a contracted threshold, which can be punitive and disrupt the CSC's budget. This is especially risky if there is no warning mechanism when thresholds are approached or crossed. • CSCs should correctly size their usage requirements, reduce the opportunity for usage creep and consider and understand the "what-ifs" of exceeding their usage thresholds.
Activation	• The time at which the service becomes active must be defined precisely to provide a "reference starting point" for the measurement of performance. This is important to measure certain metrics that are associated with a specific time window (e.g., number of outages per 30-day period). It impacts whether an "event" triggers a penalty clause and when an accumulated measurement is reset for the next subscription period. • From a CSA compliance perspective, it is important for CSCs to understand the trigger points under the CSA so they can independently measure event timing.
Payment and Penalty Models	• The CSA should clarify when/how payment is to be made. CSP payment models vary. Monthly recurring or "pay as you use" models are typical. • There may be credit terms that require advanced payment or payment every 30 days. "Just in time" CSPs are sensitive to poor credit control and are likely to be more diligent in suspending service. • Equally, the CSC needs to be diligent in obtaining service credit payments for outages.
Governance / Versioning	• CSP services evolve. New features may be added, others will be taken off support, and some may persist indefinitely. Where the assumptions or conditions under which the CSA was initially accepted are changed, the CSC should review the impact on their specific situation. • CSCs should ensure that there is a mechanism to inform them of changes and, if not, amend their contract to put the onus on the CSP to provide reasonable advance notice of updates. A good CSP will maintain a proactive policy of advising CSCs of changes to their CSA and practice version control on the documents.

<table>
<tr><th>Policy</th><th>Description / Guidance</th></tr>
<tr><td>Renewals</td><td>

Renewals are an opportunity to bargain for better rates or services levels, or switch to another CSP if necessary.
CSPs may write in their contracts an automatic renewal clause that executes in the absence of a cancellation notice before the contract's anniversary date. It is common for CSCs to overlook this deadline and be obligated to renew the contract without having had a chance to negotiate changes or even cancel the service.
CSCs should also scrutinize the renewal arrangements to find under what conditions a CSP may vary the service terms (or revise prices) upon renewal.

</td></tr>
<tr><td>Transferability</td><td>

CSCs should consider the potential need to transfer an agreement to a successor in the event their business is sold.
Conversely, if the CSP's business is acquired or divested, the CSC may not wish to do business with the new entity, and should have the option to migrate to a new CSP without penalty.
The GDPR has also stipulated that CSCs must be able to change CSPs easily.
CSCs may operate several accounts with a CSP and want to offset account credits between accounts. They should find whether the CSA supports this.

</td></tr>
<tr><td>Support</td><td>

CSCs must follow any agreed rules to report problems in order to ensure that the support terms in the CSA are activated and that "the clock starts ticking" for appropriate escalation and penalties.
An example of a support and escalation matrix related to service availability is provided below. All four target times in the table are associated with the commencement "time stamp" of the service or the notification of a service-affecting event.

<table>
<tr><th>Priority</th><th>Description</th><th>Target Response Time</th><th>Target Update Time</th><th>Target Fix Time</th></tr>
<tr><td>P1</td><td>Production software unusable (SaaS) / Production cloud servers inaccessible (PaaS or IaaS)</td><td>1 hour, CSP executive notified of issue</td><td>1 hr</td><td>Immediate - work commences and continues until issue resolved or workaround deployed</td></tr>
<tr><td>P2</td><td>Partial software functionality unusable / Partial service unavailable</td><td>4 hours</td><td>1 day</td><td>2 days, subject to available maintenance slot</td></tr>
<tr><td>P3</td><td>Cosmetic issue</td><td>1 working day</td><td>1 working day</td><td>Next software release or service update</td></tr>
<tr><td>P4</td><td>Information request</td><td>2 working days</td><td>2 working days</td><td>n/a</td></tr>
</table>

For support issues related to security incidents, the priority may be set according to the severity level of a discovered vulnerability. NIST maintains a database of vulnerabilities and a Common Vulnerability Scoring System (CVSS) [8]. Discovery of a vulnerability with a CVSS score of 6 or higher may be designated as a Priority 1 incident.

</td></tr>
</table>

Policy	Description / Guidance
Planned Maintenance	• All systems require maintenance. Complex systems may be designed to include sufficient redundancy so that maintenances can be carried out without affecting the service. • The CSA may, however, describe "uptime" as an availability percentage (e.g. 99.90%). This is the equivalent of 8.5 hours of downtime per year. CSAs may state that this does not include "planned maintenance." This highlights the importance of defining the measurement window. If the availability percentage is measured each month, this allows 12 outages per year, but each of them cannot last more than 42 minutes without triggering a penalty.
Subcontracted Services	• CSPs sometimes include in their CSA a clause that the CSA of an upstream (subcontracted) CSP will govern the services provided by the subcontractor, and that the only available penalties are those from the upstream provider even though its CSA may be less rigorous. The CSC's expectation, based on reviewing the CSA of their immediate CSP, may thus be violated. • Therefore, the CSC should ensure that the immediate (i.e., "prime contractor") CSP's CSA states unambiguously that its CSA applies to the complete service, regardless whether parts of the service come from third parties. • A CSP should be jointly and severally liable with such a third party for any breaches of the agreement by said third party.
Licensed Software	• Cloud services may include third party licensed software which is sold on a monthly licensed basis under a CSP license agreement. Such software is updated regularly by its manufacturer. • CSPs may opt to pass the responsibility for updating the licensed software over to the CSC once they have started to use the service. This absolves the CSP of the risk of disrupting the CSC's operation through an unforeseen software conflict or bug. • Alternately, the CSP may "push" the update, in which case the CSA should require that the CSC be given advance notice of the update. The CSC should have the ability to opt out, or at least to defer the update. However, the CSP may be unwilling to continue to support older versions indefinitely, and there should be a legitimate exception for updates that correct serious security vulnerabilities. • A test environment in which the CSC can exercise the new version without disrupting its operations is highly desirable.
Industry-Specific Standards	• Regulated sectors such as government, financial services, or healthcare, are subject to specific and often quite onerous standards; those must be addressed in the CSA and implementation plans. • CSCs who operate in these regulated industries should ensure that their legal team is fully involved on the negotiation of the CSA.

Policy	Description / Guidance
Additional Terms for Different Geographic Region or Countries ("Geo-Jurisdictions")	• CSCs should consider the CSP's origins and primary market. Detailed refinements to the home market CSA may be required to properly cover CSCs who are located in other areas. • Data protection legislation is one aspect of this, but CSCs should not limit their examination of the agreement to this sole aspect. • Trade control requirements should be considered "in scope" and may require the CSP to identify the country of origin of its remote cloud users, as well as granular record keeping on object access by jurisdiction.
Impact on the CSC's customer relationships	• A CSC that uses a cloud service to provide (in turn) a service to its own customers may be bound by certain SLAs or service level objectives. Such as CSC should carefully review whether any contradiction may arise between the "upstream SLA" (with the CSP) and the "downstream SLA" (with their customer). An obvious case would be a higher uptime commitment to a customer than the uptime guaranteed by the CSP in the CSA.

Step 3: Understand Service and Deployment Model Differences

As cloud computing continues to extend its reach, it is lending itself to virtually "anything and everything as a Service," dubbed XaaS. Emergent service models such as Containers-as-a-Service (CaaS) and serverless Function-as-a-Service (FaaS) foment fantastic changes and new opportunities in cloud services. Over time, this will probably result in a greater variety of service agreements. However, cloud provider offerings still fall into one of the three major service models (IaaS, PaaS, and SaaS), and the newer models are still not quite distinct enough (e.g., CaaS is usually considered a subtype of PaaS), or mature enough in terms of their impact on agreements, to allow us to flesh them out.

CSAs for IaaS, PaaS and SaaS will reflect significant differences in the levels of cloud resource abstraction, service level objectives, and key performance indicators related to each. In addition, the level of clarity varies significantly across service models. To increase effectiveness, specific components of the CSA should be stated in measurable terms and should include:

- The service to be performed and outcome expectations
- Key Performance Indicators (KPIs) and the level of service that is acceptable for each
- The manner by which service is to be measured
- The parties involved and their responsibilities
- The reporting guidelines and requirements
- Incentives for the CSP to meet the agreed-upon target levels of quality

Table 3 highlights the different CSA considerations for each of the cloud service models. Note that across the three models, metrics are not consistently defined across CSPs; therefore, the CSC must be knowledgeable in order to make informed decisions.

Table 3 -- CSA Considerations by Service Model

Service Model	CSA Considerations
IaaS	Cloud IaaS CSAs are similar to SLAs for network services, hosting, and data center outsourcing. The main issues concern the mapping of high-level application requirements on infrastructure service levels. Metrics are well understood across the IaaS abstractions (computing, network, and storage), and should include: ● Compute metrics: availability, outage length, server reboot time. ● Network metrics: availability, packet loss, bandwidth, latency, mean/maximum jitter. ● Storage metrics: availability, input/output per second, maximum restore time, processing time, latency with internal compute resource. Compute metrics usually exclude service levels for compute performance. Customers are simply guaranteed availability of the compute resources for which they paid. CSCs must distinguish between IaaS development and production environments. The latter will typically require more stringent service level objectives than the former. Network metrics in a cloud SLA generally cover the CSP's data center connectivity to the Internet as a whole, not to any specific CSP or CSC. There are several standardization efforts underway to describe and manage IaaS services.[3] Whenever possible, customers should ensure the CSA commits the CSPs to support open standard interfaces, formats and protocols to increase interoperability and portability.
PaaS	Two main approaches exist for building PaaS solutions: *integrated solutions* and *deploy-based solutions*. When reviewing the PaaS service agreement, CSCs should consider tradeoffs in flexibility, control, and ease of use to determine which approach best meets their business needs. ● Integrated solutions are web-accessible development environments that enable developers to build an application using the infrastructure and middleware services supported by the CSP, which primarily controls the management and execution of the application. Typically, service developers only have access to a CSP-defined set of APIs that offer limited control over the coordination of code execution. ● Deploy-based solutions enable deployment of middleware on top of resources acquired from an IaaS CSP, offering the CSC deployment services that automate the process of

[3] IaaS standards include: DMTF CIMI (Cloud Infrastructure Management Interface), DMTF OVF (Open Virtualization Format), SNIA CDMI (Cloud Data Management Interface), The Open Group's SOCCI (Service-Oriented Cloud-Computing Infrastructure), OGF OCCI (Open Cloud Computing Interface), and ISO JTC1/SC 38 Working Group 3 on Cloud Computing. Open source IaaS offerings are having a profound impact on the market and should be considered (OpenStack, in particular).

Service Model	CSA Considerations
	installation and configuration of the middleware.4 These PaaS solutions offer a rich set of management capabilities, including the ability to automatically change the number of machines assigned to an application, and auto-scaling according to the application's usage. At a minimum, IaaS SLAs should roll into PaaS SLAs. CSCs must distinguish between PaaS development and production environments. The latter will typically require more stringent service level objectives than the former. Standards are emerging to help identify PaaS services offered by CSPs and standard interfaces for communicating with PaaS providers to provision or manage PaaS environments. Standards such as TOSCA [9] have come about to address portability and interoperability across providers. In addition, PaaS open source offerings such as Cloud Foundry and OpenShift are starting to build momentum in the market. CSPs should ensure their CSA includes support for open standards, as they become available, to reduce vendor lock-in.
SaaS	CSCs should insist on flexible CSAs that are measurable against their business objectives, not the CSP's reporting needs. Given the broad variety of SaaS services, it is difficult to provide a comprehensive and representative list of SaaS service level objectives to look for. CSCs should expect general SaaS service level objectives like *monthly cumulative application downtime*, *application response time*, *persistence of customer information*, and *automatic scalability* to be included in their CSA. CSCs should ensure that data maintained on the CSP's cloud resources be stored using standard formats to ensure data portability in the event that a move to a different provider is required.

In addition to service models, service deployment terms should be included in a CSA. These terms should clarify to both parties signing the CSA the information required to verify the correctness of deployment actions. Specifically, these terms should identify:

- Deployment model
- Deployment technologies adopted

The deployment model included in the CSA should clearly specify one of the following options: *Private*, *Community*, *Public*, or *Hybrid*. CSCs must be well educated on the characteristics and differences in each of these deployment models since potential value and risk varies significantly. Refer to the *Practical Guide to Cloud Computing* [5] for considerations on selecting a deployment model.

[4] Deploy-based solutions are supported by commercial providers like IBM, Oracle and Microsoft as well as government-sponsored projects like OPTIMIS, CONTRAIL, Cloud4SOA and mOSAIC in Europe.

Table 4 highlights the different CSA considerations across the deployment models.

Table 4 -- CSA Considerations per Deployment Model

Deployment Model	Considerations
Private (On-site)	CSA considerations for private on-site clouds are similar to those of a traditional enterprise IT SLA. However, given that data center resources may be shared by a larger number of internal users, CSCs must ensure that critical service objectives like availability and response time are met via ongoing measurement and tracking.
Private (Outsourced)	CSA considerations for private outsourced clouds are similar to private on-site ones, except that the cloud services are now being provided by an external CSP. The fact that IT resources from the CSP are dedicated to a single customer mitigates potential security and availability risks. CSCs should ensure the CSA specifies security techniques for protecting the CSP's perimeter and the communications link with the provider. CSCs should consider the criticality of the service being deployed to justify the added expense of this model over the Public model.
Public	CSA requirements for public clouds are greater than for private outsourced ones since the CSP's IT resources are now shared across multiple CSCs ("tenants"). As a result, CSCs must understand how the CSP addresses the added security, availability, reliability and performance risks introduced by multi-tenancy. The ability to measure and track specific service level objectives becomes more important in the Public deployment model. CSCs should also ensure the CSA provides adequate methods and processes for ongoing measurement.
Community	For the purpose of this discussion, the considerations are the same as for the Public model above.
Hybrid	CSA considerations for the Hybrid model are similar to the Public model, with the increased likelihood of unique integration requirements between cloud and enterprise services. CSCs should ensure the CSA adequately covers their service and data integration requirements. It is recommended to use a specific and standard document that describes the nature of the interface (along with quality level metrics and performance characteristics associated with the interface) and any security requirements. For example, if the interface is a web service, there may be authentication and authorization requirements with implications for single sign-on and a directory (e.g., LDAP) mechanism.

In addition to specifying the deployment model, the CSA should clarify how a service is made available to service users on a given CSP, for example:

- A web application is deployed on an application server as a Web application ARchive (WAR) file. [10]
- A grid application is deployed on a grid container as a Grid ARchive (GAR) file.
- A virtual machine is deployed on an IaaS provider as a virtual machine disk image that may be represented in various formats. Adoption and support of standards like the Distributed Management Task Force (DMTF) Open Virtualization Format (OVF) is recommended. [11]

When CSAs are signed, a clear description of the technologies involved in the deployment of services should be specified. Note that there is a close relationship between deployment technologies and the kind of services being offered.

Step 4: Identify Critical Performance Objectives

Performance goals within the context of cloud computing are directly related to the efficiency and accuracy of service delivery by the CSP. Typical performance considerations include availability, response time and processing speed, but they can include many other performance and system quality perspectives, such as accuracy, portability, interoperability, standards compliance, reliability, scalability, agility, fault tolerance, serviceability, usability, durability, and more. From this long list, CSCs must decide which measures are most critical to their specific cloud environments and make sure that these measures are included in the SLA. For example, a CSC may require a CSP to support its own (the CSC's) contractual obligations to its customers.

Performance factors important to the CSC should be measurable and auditable, like all metrics, and documented in the SLA in order to provide for rational discussions between the parties. The relevant performance factors depend on the service model (IaaS, PaaS or SaaS) and the type of services provided within that model (for example, network, storage and computing services for IaaS). In order to assess performance objectively and establish trust between the parties, clear and consistent measurements are required. It must be clear how each metric will be used and what decisions will be made from the measurements to align service performance to specific business and technical goals and objectives.

This section will focus on two performance metrics: *availability* and *response time*. The intention is to provide a basic framework to identify and define meaningful and consistent cloud metrics. This framework can then be applied to other metrics not covered here. While many metrics may already be supported by a CSP, it may interpret the definition differently than the CSC does. An agreed definition in the context of a specific cloud solution is critical. Moreover, the measurements captured by a provider should match the definition included in the SLA.

Industry standards should be used when possible to improve consistency. For instance, IEEE has good measurement definitions and categorizations for activities such as maintenance.[5]

[5] Other standard organizations working on measures relevant to cloud services include the International Function Point Users Group (IFPUG), which formed a Cloud Measurement Interest Group in 2013. Some of these efforts leverage existing software measurement guidelines, such as ISO/IEC 20926, which are used for benchmarking. IFPUG works closely with the International Software Benchmarking Standards Group.

Here are the generally accepted definitions for the two metrics of interest:

- *Availability*. Percentage of uptime for a service in a given observation period.
- *Response time*. Elapsed time from when a service is invoked to when it is completed (typically measured in milliseconds).

Table 5 describes three different example scenarios (network availability, storage availability, and service response time) and the specific performance information required for each.

Table 5 -- Availability and Response Time Examples

	Network Availability	Storage Availability	Service Response Time
Metric Name in SLA	Network Percentage Available during critical business hours	Storage Percentage Available	Service X, Y, etc., response times in a given hour
Constraints	Critical time is defined as 12AM GMT to 12PM GMT Monday through Friday	None	Response times will only be evaluated for services X and Y, which are PaaS reusable services that will be invoked by our applications.
Collection Method	Machine	Machine	Machine
Collection Description	Using the DMTF, OGF[6], or other standard to consistently collect the measures.	Using the DMTF, OGF, or other standard to consistently collect the measures.	Using the DMTF, OGF, or other standard to consistently collect the measures.
Frequency of Collection	The network is "pinged" every one minute.	Specific storage services (read and update) are randomly "pinged" every one minute.	For each service invoked, the response time is collected every five minutes.
Other Information	60 seconds of uptime will be recorded for each successful "ping."	60 seconds of uptime will be recorded for each successful "ping."	Each service will be reported separately. Hourly averages will be calculated.
Clarification	No reference to quality or availability of specific services. This is exclusively a measure of network availability.	No reference to quality or availability of specific services. This is exclusively a measure of storage availability.	No individual service reporting is needed (e.g., listing of all services that exceeded SLA agreed response time).
Usage 1 in SLA	Network availability shall be 99.5% or higher between 12AM GMT and 12PM GMT Monday thru Friday.	Storage availability shall be 99.9% or higher.	Response time for service X shall be less than 500 ms, and for service Y less than 200 ms.

[6] Refer to http://www.gridforum.org/ for more information on the Open Grid Forum.

	Network Availability	Storage Availability	Service Response Time
Usage 2 in SLA	For any day when network availability is less than 99.5%, a 20% discount will be applied for the entire day's network charges.	For any day when storage availability is less than 99.9%, a 50% discount will be applied for the entire day's storage charges.	If in any given hour the response times as stated are not met, all services of that type during that hour will be processed at no charge.

Both hardware and facilities should be considered when assessing critical performance levels in an IaaS context. Hardware includes computers (CPU and memory), networks (routers, firewalls, switches, network links and interfaces), storage components (hard disks, Solid State Disks [SSDs], Flash Storage, Software Defined Storage [SDS], tape), and any other physical computing infrastructure elements. Facilities include: heating, ventilation and air conditioning (HVAC), power consumption and dissipation, communications, backup, and other aspects of the physical plant. In the case of PaaS or SaaS solutions, it can be presumed that the unavailability or sub-par performance of any of these components will affect the overall services, therefore it is not necessary to specify them – the measurements should be "end to end," that is, expressed in terms of the user experience.

Moreover, particularly in the IaaS case, higher-level business objectives may dictate what critical resources fall within the scope of the metrics. For example, the power consumption or the heat dissipation may or may not be included, depending whether the customer has established a corporate carbon footprint objective.

In summary, when considering performance metrics in a cloud SLA, it is recommended that consumers:

- Understand the business level performance objectives – for example, reduce cost and time to market per unit of software functionality.
- Identify the metrics that are critical to achieving and managing the business-level performance objectives.
- Ensure that these metrics are defined at the right level of granularity that can be monitored on a continuous basis (in a cost-effective manner).
- Identify standards that provide consistency in metric definitions and methods of collection.
- Analyze and leverage the metrics on an ongoing basis as a tool for influencing business decisions.

Step 5: Evaluate Security and Privacy Requirements

Security [7]

Security controls[8] apply to both on-premises systems and cloud computing. However, because of the cloud service models employed, the operational models, and the technologies used to enable cloud services, cloud computing may present different risks to an organization than traditional IT solutions. For example, new technologies undergoing rapid adoption, such as containers, introduce certain new risks. The distribution of responsibility and accountability between CSCs and CSPs for both implementation and subsequent management is new, and must be clearly specified and understood. Additionally, the security postures, response policies and control schemes of a CSP need to be integrated into those of the CSC. Refer to the *Security for Cloud Computing: 10 Steps to Ensure Success* [2] white paper for details on security requirements for cloud computing.

A required foundation for security, regardless of whether a cloud solution is used, is a *security classification scheme* that applies throughout the enterprise and is based on the criticality and sensitivity of enterprise assets. This scheme should include details about-asset ownership, definition of appropriate security levels and protection controls, and asset retention and destruction requirements. The classification scheme must be the basis for applying access controls, archiving, and encryption methods.

In order to determine what level of security is required for a specific asset, a rough assessment of an asset's sensitivity and importance is required. For each asset, the following questions should be asked:

"How would the business be harmed if…

1. *The asset became publicly available and distributed?*
2. *The process or function was manipulated by an outsider?*
3. *The process or function failed to provide expected results?*
4. *The information was unexpectedly altered?*
5. *The asset was unavailable for a period of time?"*

Table 6 below highlights the key steps customers should take to ensure their CSA sufficiently addresses their unique security requirements.

[7] The security part of this section is based on the Cloud Security Alliance "Security Guidance for Critical Areas of Focus in Cloud Computing, V3.0" [12] and quotes portions of that document.

[8] NIST Publication 800-53 [] defines security controls as "the safeguards/countermeasures prescribed for information systems or organizations that are designed to: (i) protect the confidentiality, integrity, and availability of information that is processed, stored, and transmitted by those systems/organizations; and (ii) satisfy a set of defined security requirements."

Table 6 – Key Security Considerations in CSAs

Security Considerations	Strategic Activities
Assess asset sensitivity and operational security requirements	• Complete an assessment of the confidentiality, integrity, and availability requirements of the assets. • Complete a threat risk assessment and privacy risk assessment. • Address operational security, availability requirements and privacy requirements in response to identified risks and in line with the organization's data classification, information architecture, information security architecture, and risk tolerance.
Understand legal/ regulatory data residency requirements	Understand the regulatory, contractual and other jurisdictional constraints about the logical and physical locations of data, persons, person types (citizenships), services, assets, records and applications.

Restrict and secure data movement; prevent accidental data disclosure	• Establish policies to restrict the movement of sensitive data to cloud services by individuals or departments without the approval or, at minimum, notification of the Security/Privacy departments.
	• Take steps to detect such unapproved data moving to cloud services:
	  ▪ Monitor for large internal data migrations with database activity monitoring (DAM) and file activity monitoring (FAM)
	  ▪ Monitor for data moving to the cloud with URL filters and data loss prevention.
	• Protect data in transit. All sensitive data moving to or within the cloud should be encrypted.
	• Protect data at rest. Sensitive datasets should be encrypted to limit exposure to snapshots or unapproved administrator access. Sensitive data in object storage should be encrypted, usually with file/folder or client/agent encryption.
	• Identify a set of required counter-measures (security controls) that implements the above choices using standard frameworks such as NIST SP 800-53 [13], NIST SP 800-171 [14], ISO 27001 [15], CIS Security Controls [16], or the CSA Cloud Controls Matrix [17].

Establish and track security metrics	• Metrics and standards for measuring the performance and effectiveness of information security management should be established prior to moving to cloud. • At a minimum, organizations should understand and document their current metrics and how they will change when operations are moved into the cloud and where a provider may use different (potentially incompatible) metrics. • Refer to the following resources for specific information on security metrics: ISO 27004:2009 [18], NIST Special Publication 800-55 [19], and CIS Consensus Security Metrics [20].

Security Considerations	Strategic Activities
Assess the CSP's security capabilities	• Assess the CSP's level of security and its maturity. • If compliance to a normative standard (e.g. ISO 27001/27017/27018, see [21]) is asserted, then verify the compliance certificate's validity, and pay attention to the scope of the certificate to ensure it covers the capabilities that are relevant to the CSC. • Look for verifiable evidence of resource allocation, such as budget and manpower to sustain the compliance program. • Verify that annual 3rd party penetration test is performed and remedial actions are taken.
Assess the CSP's security governance	• Assess the CSP's security governance processes and capabilities for sufficiency, maturity, and consistency with the customer's information security management processes. ▪ The CSP's information security controls should be demonstrably risk-based and clearly support these management processes. ▪ Where a CSP cannot demonstrate comprehensive and effective risk management processes in association with its services, CSCs should carefully evaluate use of the provider as well as the user's own abilities to compensate for the potential risk management gaps. • Determine if the CSP's guarantees adequately address your security requirements. The Cloud Security Alliance's *Consensus Assessments Initiative Questionnaire (CAIQ)* provides a set of questions a cloud consumer and cloud auditor may wish to ask of a cloud provider. [22]
Audit the CSP's security CSA compliance	• Few CSPs will consent to a "right to audit" clause in a CSA, which would give CSCs the ability to audit the CSP. For the CSP, this is intrusive and could lead to a multiplication of audits by multiple customers. For the CSC, it is a costly process and requires specialized skills. On the other hand, audits are critical to traceability and transparency. • In lieu of this right, the CSC may require that the CSA obligate the CSP to undergo regular third-party audits, based on recognized normative standards such as the ISO 27000 series, and share the results of the audits. • Self-assessment certifications, such as recorded in the CSA STAR repository [23], also allow the CSC to check at a detailed level the security controls implemented by the CSP, even if this is only an unproven declaration.

Security Considerations	Strategic Activities
Notification	• CSPs should notify CSCs of the occurrence of any breach of its system, regardless of the parties or data directly impacted. • The CSP should include specific pertinent information in the notification, stop the data breach as quickly as possible, restore secure access to the service as soon as possible, apply best-practice forensics in investigating the circumstances and causes of the breach, and make long-term infrastructure changes to correct the root causes of the breach to ensure that it does not recur. • Due to the high financial and reputational costs resulting from a breach, CSCs should require the CSP to accept liability if the breach was their fault (the indemnification clauses contained in the provider's CSAs are often written the other way around: they are meant to protect the CSP from being sued). If the CSP considers the cost of such liability to be too high, they (or, if necessary, the CSC) should consider cyber-insurance products.

Privacy

Many countries have multiple laws, regulations, and other mandates that require public and private organizations to protect the privacy of persons whose personal data is stored in computer systems.

When data is transferred to a cloud, the responsibility for protecting and securing the data typically remains with the controller or custodian of that data, even if in some circumstances this responsibility may be shared with others. When it relies on a third party to host or process its data, the controller of the data remains liable for any loss, damage, or misuse of the data. It is prudent, and may be legally required, that the PII controller and the PII processor (i.e., the CSP) enter into a written (legal) agreement that clearly defines the roles, expectations of the parties, and allocates between them the many responsibilities that are attached to the data at stake.

If privacy issues are not adequately addressed in the CSA, the CSC should consider alternate means of achieving their goals including seeking a different CSP or not sending sensitive data to the cloud. For example, if the CSC wishes to send HIPAA[9]-covered information to the cloud, the CSC will need to find a CSP that will sign a "HIPAA business associate agreement," or else not send that data to the cloud. In addition, CSCs may consider keeping sensitive data on premises, or encrypted in a fashion that the CSP cannot read, due to the Cloud Act of March 2018 [24]. Under this Act, all cloud providers are obligated to surrender data requested by U.S. law enforcement agencies, regardless of data residence. The European Union's General Data Protection Regulation (GDPR) has been very much at the forefront of mature privacy practices, with major financial penalties applied for non-compliance.

[9] Health Insurance Portability and Accountability Act, United States legislation from 1996 that provides data privacy and security provisions for safeguarding medical information.

Preservation of information, included in some privacy regulations, can require that large volumes of data be retained for extended periods. This has significant ramifications as far as the CSA is concerned. The CSA should address questions such as:

- What happens if the preservation duration requirements outlast the terms of the CSA?
- If the CSC preserves the data in place, who pays for the extended storage and at what cost?
- Can the CSC effectively download the data in a forensically sound manner, so it can preserve it offline or near-line?

The reverse risk may also exist: the backup and disaster recovery policies of a CSP may cause copies of data or code to be retained beyond the retention period intended by the CSC. This may be a problem during the "discovery" phase of litigation. One party may claim that certain data (e.g., copies of old e-mails) has been erased, and the opposing party may discover that it still exists in the backups made by a CSP and subpoena the CSP. In another scenario, a CSP may wish to implement an end user's "right to be forgotten," only to find that it has no ability to selectively delete the backup copy of the user's records. Some of these risks may be mitigated by using encryption techniques where the CSC controls the encryption keys, such as in "bring your own key" offerings.

Refer to the *Security for Cloud Computing: 10 Steps to Ensure Success* [2] whitepaper for an overview of the privacy requirements for cloud computing and worldwide privacy regulations that currently exist.

Step 6: Identify Service Management Requirements

Critical to achieving the business goals of moving to the cloud is a uniform, straightforward, transparent and extensible system for managing and monitoring cloud services. In this section, we outline key CSA considerations in the area of service management.

Every computing system requires internal controls, management, automation, and self-healing in order to operate in today's interconnected world, an area commonly called Application Performance Management, or APM. A move to the cloud still requires these elements – perhaps even more so. Such controls apply to all service layers such as:

- Core infrastructure provision (host services)
- Network carriers and ISP (including failover on excess volume or restricted capacity)
- Data storage, with or without backup and mirroring
- Application (SaaS/PaaS models) with or without data storage.

Scope of Service Management:

Auditing

Monitoring and Reporting

Measurement and Metering

Provisioning

Change Management

Upgrades and Patching

Although there is not yet a standard CSA language for service management, the agreements must include provisions covering the following aspects.

Auditing

Managing cloud services critically relies on an independent or verifiable methodology for auditing and reviewing those services. Asking for evidence of this helps discern between CSPs who are fully capable of deep manageability and those who provide only a simple veneer on someone else's offerings. As stated by many an experienced manager, people "do what you inspect, not what you expect."

The objectives of CSA terms in the area of auditing are:

1. Provide the CSC with an unbiased assessment of its ability to rely on the CSP
2. Assess the depth and effectiveness of the CSP's internal systems and measures
3. Provide tools to compare quality levels with other CSPs
4. Ensure the openness needed to allow continuous review and improvement
5. Uncover issues in the CSC's own ability to interface with the CSP in order to provide uninterrupted services
6. Specify the parties' roles and responsibilities in case of a disruption or failure of the contracted services

Objective 5 deserves a comment. Many documented incidents have come not (or not only) from a CSP's inability to service a customer, but from the ability of the customer's systems to interface properly with the cloud. Therefore, any audit scope should include both the CSP and any internal systems of the CSC exposed to the cloud to ensure a complete "envelope" of integrity.

The scope of an audit protocol must be broader that the contract terms and conditions, and address general issues of management and governance, including necessary resources to mitigate any risks found. For example, it is insufficient to include a provision to regularly audit security and encryption keys, only to neglect addressing any internal resource allocations, scheduling, review and approval processes needed to perform the audit and address any issues stemming from the audit.

Organizations should try to leverage methods of audit and compliance they already apply in their business, and extend those to the cloud rather than creating new ones.

Monitoring and Reporting

Service management requires visibility of the service data. While every CSP offers a different system for visualizing data (Web-based, e-mail based, live, reactive, portal-based), and there are no standards or benchmarks yet, customers should look within the CSA for a minimum set of monitoring and reporting (including visualization) capabilities:

1. *Cloud Performance Management*. Monitoring and reporting of the response times of systems within the cloud and between the cloud and the target user systems.
2. *Peak Load Performance*. Monitoring performance factors when the cloud is under stress, either intentional or unintentional. Systems under a heavy load can perform differently than expected, and the interactions and dependencies of a complex cloud are often unknown in advance.

3. *Hybrid and Multi-cloud Performance*. Clouds that consist of different subsystems, often sourced from different cloud providers, require monitoring the interactions between those hybrid cloud components.
4. *Application Performance*. This concerns the internal processing benchmarks of cloud-based applications, as well as end-user experience measurements.
5. *Problem Notification*. Monitoring and reporting on failures and issues with the cloud system, specifically the notification timeliness, severity level assessment, and escalation processes.
6. *Incident/Problem Management*. This concerns the resolution of the incident, or alternately launching mitigation steps or invoking recovery or continuity procedures. There should be an integrated system to manage and monitor incidents, shared between the CSP, CSC, and any involved third parties (such as subcontractors to the CSP) to avoid multiple entry of problems and inconsistencies in status reporting.
7. *Record keeping* may be required in specific cloud service use cases, to meet requirements on trade controls, protection of PII, etc. The CSA should identify those records the CSP must maintain and make available to the CSC.

The CSP may provide native service monitoring capabilities to the CSC. However, the CSC must establish whether the monitoring and alerting capabilities provided meet their requirements, and if necessary, may need to put in place an additional layer/tools for service alerts, possibly with a "sense and respond" self-learning capability. In the case of a multi-cloud environment, integrating the information (alerts or reports) coming from multiple CSPs is particularly important (see the *Practical Guide to Cloud Management Platforms* [25]).

Measurement and Metering

Since most cloud services are billed dynamically on a usage basis, the CSC must have confidence in the accuracy and transparency of the measurement and metering system employed by CSPs. The cost savings or "cost shifting" benefits of the cloud depend on this.

The CSA should commit the CSP to employ a metering system with the following capabilities:

1. Assurance of accurate billing, and a methodology for handling objections or challenges to any automated metered billing (e.g., missing credit for a service interruption or degradation).
2. Applying different methods to different services. For example, performance testing, analytics, security scanning, backup, and virtual desktops might all be measured differently and metered separately.
3. Handling taxation and compliance issues as appropriate for each location and user. For example, as countries and local governments have implemented different approaches to tax online commerce, the CSP must be able to discern between these sources of use and meter them independently.

Provisioning

Service provisioning is a key enabler of the agility that comes from using the cloud, and must have certain qualities that are reflected in the CSA:

1. *Core provisioning speed*. Baseline expectations of the speed of deployment of new systems, new data, new users, new desktops, or any capability within the scope of the contracted service.
2. *Customization*. It is unusual that any template-based rapid provisioning method can be used "out of the box" without configuration and customization. Without careful management of the expectations and contractual levels for this function, any savings gained by automated rapid provisioning can evaporate in the face of delays in customizations post-deployment.
3. *Testing*. A strong CSA must provide for testing automated deployment and scaling prior to need. This is particularly acute in areas where provisioning is employed in disaster recovery or backup situations.
4. *Demand Flexibility*. The system must accommodate dynamic de-provisioning to match downturns in demand and reduce the customer's cost. The process to dynamic allocate processing units, bandwidth or storage space should be specified, including in emergency situations.

This is not an exhaustive list of considerations, only the basic requirements of any contractual definition of rapid provisioning. Each organization will need to add industry- or application-specific requirements.

Change Management

Activities for requesting, reviewing, testing, and acceptance of changes differ little from those already in use with other managed IT services agreements. However, roles, responsibilities and tooling often change and lead to process realignment. In addition, the deep cultural shift of adopting a cloud strategy may make people more sensitive to change, more suspicious, or less likely to cooperate. In this case, extra care should be taken to manage the process carefully.

Conflicts may easily arise between a CSP and a CSC when the CSP assumes that a change will have no impact, and therefore does not need to be communicated to its CSCs or approved by them, and it turns out that there is an adverse impact after all. A change to a service component (e.g., software upgrade) could trigger a domino effect on connected and interdependent components. A new security disposition or a change to the names and conventions of a data store could disable a SaaS application. All changes to the cloud configuration and architecture, even if they are supposed to have no impact, require both local testing and integrated or regression testing to verify that the interoperability of cloud components was not compromised; the CSC and CSP should agree in advance about the responsibility for, and extent of, such testing.

Customers should be notified of an impending change so that they can check that the change occurred on schedule, be on alert for any unexpected behavior that may be a consequence of the change, and report it immediately.

Upgrades & Patches

Some of the changes to be managed include:

- Upgrades to SaaS applications or to platform software (operating system, middleware, database, management system…)
- Patches, including urgent security updates or severe bug fixes
- Improvements in existing contracted services.

The CSA should outline the basic rules to handle these inevitable needs.

1. **Responsibility for change requests.** The CSA should state which party is responsible for each type of change and will lead the effort. A SaaS software upgrade is usually the responsibility of the CSP, while a database upgrade in an IaaS/PaaS environment may be the CSC's responsibility. The CSA should include a service responsibility matrix.
2. **Planning process.** For changes potentially impacting the CSC, there should be a clearly defined project plan to develop, test and implement the changes made to the cloud environment. This includes the identifying the project manager, resources, schedule, risks, contingencies, and problem resolution methods to ensure success. This is similar to other managed IT services situations, except for the increased anxiety and scrutiny that the cloud draws.
3. **Problem resolution.** Since problem resolution may require close cooperation between CSP and CSC, the CSA upgrade procedures must specify responsibilities and methods for resolving issues. This often leads to re-engineering of change and problem management processes.
4. **Rollback process.** Upgrade plans should include checkpoints and rollback capabilities, and should specify in advance when to "pull the plug" and restore the environment to its initial state in case of an unexpected problem cannot be solved immediately.

Throughout the process, regular communication meetings should occur to keep both parties informed and aligned.

Step 7: Service Failure Management

Service failure management deals with what happens when the satisfactory delivery of a cloud service does not occur. Of course, in order to objectively recognize a failure, the expected cloud service capabilities and performance should be explicitly documented in the CSA, as described in Step 4.

Occasional service outages are inevitable. Well-publicized cloud outages have lasted from a few minutes to several hours, in some cases exceeding a full day, with occasional customer data loss. Preparation should aim at reducing both impact and duration of the failure. Service failure management preparation involves the CSC, the CSP, and any third-party cloud services or integrators.

Service failure falls within the scope of three "layered" contractual instruments. The **Service Level Agreement** is a contractual commitment to provide and measure a service, with remediation, escalation and penalty clauses. It encompasses:

- The **Operating Level Agreement (OLA)**, which describes the interdependencies and responsibilities between Cloud Service components and providers (such as an ISP, a network carrier, an application owner, a data center hosting service, etc.)
- **Quality of Service (QoS)**, an objective specification of the performance of the service, especially in terms of network access (volume, latency, prioritization, failures), including when demand exceeds nominal capacity.

Service Failure management and preparation can vary based on the deployment model: Private (single tenant), Public (IaaS, PaaS, SaaS) or Hybrid (the most common today).

Service Failures

Service failure management begins with the detection and alert that an event has occurred (see "Monitoring" in Step 6). Immediate detection and response to an incident is critical to limit impact or propagation and should be tested in advance. Table 7 shows six different types of a service failures resulting in disruptions.

Table 7 – Types of Cloud Service Failures

Service Component	Responsibility	Comments
Unavailability of all or part of the Cloud Services	CSP or hosting provider	
Failure of network infrastructure component(s)	Network provider, CSP or CSC	Depends on who owns and operates the failing component
Failure to execute service requests within acceptable performance – Application failure	Cloud Service Developer	Can be complex in the case of a Hybrid Cloud
Data Integrity / Availability failure	CSP generally, but see comments	Cascading responsibilities with Hybrid. End users are often a cause of data integrity issues.
Data and Service Restore/Continuity failure	CSP or contracted third party	
Security or compliance breach or failure	Access, location and infrastructure are the responsibility of the CSP. Endpoint protection and user authentication are the CSCs.	

CSCs must ensure that a comprehensive incident management system is triggered – preferably using automation, then manually if needed, upon any service failure occurrence.

Automated systems ensure that the detection is rapid, objective and policy-based. A rapidly growing market has developed in Artificial Intelligence for IT Operations (AIOps). Early localization of the failure is critical to assign troubleshooting to the CSP, ISP or host, network provider, or application owner. Descriptions in the OLA and QoS should document those troubleshooting responsibilities. An AIOPS tool can isolate compromised elements or re-route traffic to alternate cloud components. In order to avoid delays and confusion when a failure occurs, the CSP must publish, and the CSC must understand, the service failure management procedures used by the CSP to:

- Identify, report and notify failures, even when detected by a provider or by the customer
- Address and troubleshoot a reported failure, including the crisis management process
- State timescales for remedial action, by disruption thresholds (e.g., 5 min., 15 min., 1 hr.)
- Detail steps to be followed to improve the provider's operations, avoid repetition of the failure and determine whether an SLA violation occurred.

Preparing for Service Failure

The Preparation for Service Failure is the set of actions and dispositions taken to mitigate the impact of a severe service disruption. Preparedness might drive changes to the CSA, SLA and QoS, possibly the architecture of the overall cloud solution.

The CSC determines the data management disposition for the overall solution as part of its cloud strategy:

- The basic preparation measure is data backup. If the data is in the cloud, and the CSP is ensuring the backup, the CSC must periodically test the restore capability, and must require that the backup location be reasonably protected from suffering the same problem as the original storage location.
- A more complete solution is for the CSC to have a disaster recovery (DR) and business continuity (BC) plan, to be executed if the cloud service failure exceeds pre-determined impact or duration thresholds. Disaster recovery planning is discussed in Step 8 of this paper.

Key preparation steps include:

- Defining or contracting separate backup services from the core services.
- What-If scenarios covering likely service failures, and a mitigation plan involving other providers.
- Periodic testing of the DR/BC plans to validate their currency and relevance.
- Adapting the development and operations strategy to automate re-deployment, self-healing, or failover to another cloud instance.
- Reducing the risk of security attacks and data breaches through authentication, access control, and encryption.

The CSC should ask for the CSP's Disaster Recovery plan. In the United States, most CSPs should also be able to provide an annual SSAE 18 SOC2 report [26]. Compliance with ISO/IEC 17789, Cloud Computing Reference Architecture [7] is recommended, as well as Uptime Institute Tier III rating [27] for the cloud data center or host.

Monitoring and Notification

In the absence of a service monitoring capability, the CSC would just be waiting blindly for end users to notice that a service is unavailable and swamp the helpdesk with inquiries that would come as a surprise. This is the worst option and should be avoided. Instead, the customer should be proactively monitoring service availability and performance, and all parties should receive automated notification of a failure in a timely manner. Automated monitoring can be performed in several different ways:

1. The CSC can put in place a system, or systems, which monitor its usage of the cloud service, using its own tools (i.e., independently from the CSP) to monitor service performance and detect incidents.
2. The CSP's own cloud service monitoring system may offer an interface (API, web service, service portal) enabling the CSC to obtain real-time monitoring data. There are some drawbacks to this approach:
    - the interface may not carry notifications of all incidents of importance to the customer,
    - a failure at the CSP may stop the notifications from occurring,
    - this becomes complex when multiple providers are involved,
    - Problems with network connections may impact the service, even though monitoring information sent by the CSP indicates that the service is normal.
3. If there is a cloud service integrator, it may be entrusted with the overall monitoring, alerting and exception handling across all providers. This will require an additional Operating Level Agreement (OLA) to describe shared processes and procedures.
4. The CSC can use the CSP's monitoring and alerts tools. This presents a lock-in risk, limiting the portability of the monitoring of the workload to a different CSP is needed.

Implementing redundant methods will improve reliability, since each method can be affected by the failure of certain components.

For notification of service failures, the ideal situation is an automated interface between the CSC and the CSP(s), transmitting notifications of a service failure in all directions. This allows either party to first detect the incident and notify the other(s). However, there should be a backup notification method in case this interface itself fails.

In all cases, the customer must retain the ability to report a service disruption by opening a "ticket" in a centralized incident management system.

Regular drills should be performed, just as is done for fire safety, to ensure that the notifications work and that appropriate information is broadcast to the users when the incident starts as well as when it is resolved.

Remediation

Remediation steps are solutions to be deployed when a service outage occurs. Remediation should start immediately upon the declaration of service failure. Remediation goals are prioritized as follows:

1. Resume providing the full service to the users per the CSA and within the SLA normal thresholds. The resumed availability of the cloud components or service and compliance with the CSA determines the end of a failure event. Some solutions to resume service (such as switching to an alternate service in a public cloud) could cause secondary SLA, security or compliance violations. The SLA and OLA should not allow the incident to be considered resolved until there is full, SLA-compliant capability restoration.
2. Recreate a service provision through re-provisioning the cloud solution, invoking the Disaster Recovery Plan or alternate business continuity solutions. Having a distinct cloud backup or storage solution with its own SLA (including a Recovery Point Objective [RPO] and a Recovery Time Objective [RPO]) can be critical to resume service and limit the impact.
3. The ultimate and latest remedy for service failure is the granting of service credits to offset the impact of the failure, typically a percentage of the fees paid by the CSC. However, as pointed out earlier, this remedy usually does not approach the magnitude of the damage suffered by the CSC is the failure was prolonged or affected critical systems. Still, there is room to strengthen a CSA in this respect within the SLA, OLA and QoS.

Limitations

CSAs famously contain liability limitations for certain types of service interruptions. While these may vary between providers, common exclusions include:

- Scheduled or emergency outages
- Acts of force majeure
- Suspension of service due to legal reasons
- Internet access issues outside the control of the provider

Some CSPs also exclude “scheduled downtime” from the SLA’s availability metrics calculation. It may be unclear whether this refers to a predictable window of time (first Sunday of every month, 3:00 to 5:00 a.m.) or whether a downtime event outside of that rule is also excluded by virtue of being announced a few days ahead of time.

Roles and Responsibilities

The roles involved in cloud computing service failure management are described in the ISO/IEC 17789 Cloud Computing Reference Architecture [4]. The cloud operator has the responsibility to drive the incident management process and needs therefore to receive an alert when a service failure is detected.

On the CSC side, additional roles may be involved, including a central or integrated helpdesk and a cloud service integrator (if any). The helpdesk should be aware of the service failure and the likely impact and estimated time to resolution in order to answer questions from users. A cloud service integrator would be engaged to triage the service failure and potentially propose solutions or workarounds to resolve or mitigate the service failure and its impact on the customer’s business.

Step 8: Understand the Cloud Disaster Recovery Plan

Disaster recovery (DR) is a subset of business continuity and focuses on processes and technology for resumption of applications, data, hardware, communications (such as networking), and other IT infrastructure in case of a "disaster" – either a natural or human-originated event that severely and durably disables IT infrastructure or software systems.

CSCs often feel a false sense of security regarding DR planning. Just because a business has outsourced an IT capability (infrastructure, platform or application) to a CSP does not absolve it of the need for serious DR planning. On the contrary, since it may make it harder to verify the measures put in place by the CSP, due diligence is both more difficult and more important. Customers should view developing, documenting, testing, and updating a DR plan as an important step to reduce business risk while moving to the cloud. The DR strategy may need to be reconsidered in light of the move to the cloud.

Well-known key principles of DR planning apply equally to cloud services:

- Since an organization's unique business priorities dictate the relative importance of its infrastructure components and applications, a cloud DR plan is company-specific and driven by business objectives.
- The process starts with identifying and prioritizing applications, services and data
- Recovery time objectives (RTOs) and recovery point objectives (RPOs) are established separately for each service to minimize the business impact – sometimes, faster service restoration is more important than recovery every last transaction entered before an outage, or vice versa.
- The DR plan must be tested regularly.

Most CSAs provide cursory treatment of DR issues, procedures and processes, and offer inadequate guarantees in case of a service outage due to a disaster. However, large and established CSPs have such a stake in maintaining the availability of their services that their extensive DR infrastructure (alternate power supplies, multiple independent network connections, etc.) is better than what most small and medium CSCs have in place internally.

Despite the limitations in CSAs, customers should address key disaster recovery questions/issues with their prospective CSPs early in the process of cloud adoption. Here are key questions to ask:

- What level of redundancy is in place to minimize outages, including replication of services in different geographical regions?
- What actions will be taken in the event of a prolonged disruption or a disruption with a serious business impact?
- What is the process of performing disaster recovery testing, and how often are the tests conducted? Are the reports of the tests provided to clients and are the tests automated?
- What are the key service CSP and CSC contacts (names, phone numbers, email addresses, alternate means of contact in case of communications disruption)?
- What is the contingency plan during a natural disaster?
- What events are excluded from the guarantees in the CSA?

- Does the CSP provide cloud insurance to mitigate user losses in case of failure?

While established cloud vendors are quite resistant to altering existing CSAs, large customers may use these questions as a framework to negotiate stronger DR procedures and guarantees. Smaller customers do not have that leverage, but may benefit from procedures and systems put in place by a CSP to satisfy its larger clients.

Besides fires, earthquakes, electric grid failures and the like, other business risks need to be considered, such as:

- security breaches, malware attacks or denial-of-service attacks,
- failure of the CSP's business,
- law enforcement actions, including seizure of assets,
- failure of the CSP to make the customer's data available for retrieval on demand

Some of these events may hamper a CSC's attempt to overcome a failure by switching to another provider. Cloud users should plan ahead for such contingencies, focusing on restoring access to data and applications in a timely manner.

Risk impact and mitigation measures depend on the cloud service model (IaaS, PaaS or SaaS):

- For an infrastructure service failure, the use of hot/warm recovery sites in different geographical areas or on a completely different cloud may be sufficient, as long as the data is replicated or is securely backed up in a location not affected by the same disaster as the primary resource. While the CSP may provide such services, it is up to the CSC to subscribe to them, balancing the additional cost with the business risk of a failure.
- If a SaaS provider suffers a serious failure, it *may* be possible to retrieve the application data, but software and the business logic customizations will be left behind. Re-deploying the application on premises is a complex task. Despite good planning, in some cases no easy recovery solution is available. Standards for data and meta-data in specific application domains (ERP, CRM, etc.) may be needed in the future to support migration between different SaaS solutions in the event of a disaster. This would also alleviate "lock-in" issues – which is why such an effort will face resistance from SaaS providers.

Any disaster recovery or mitigation plan based on a failover capability to an alternate site must answer the following questions:

- Which components require redundancy, and which do not?
- How should the components be laid out and configured?
- What is the connectivity between the components?
- Can a failover be triggered automatically, or does it have to be manual?
- What will the process be for the operations staff?
- Should the disaster site be enabled and used all the time as a second site (active/active configuration) or should it only be used when the primary site is down (active/passive)?

Existing standards such as ISO 22301:2012 [28], NIST SP 800-34 [29], ASIS ORM.1 [30], ISO/IEC 27031:2011 [31], and ISO 24762:2008 [32] can provide effective frameworks for planning disaster recovery. Asking a CSP which of these standards it follows is also a way to assess its command of the subject.

Step 9: Develop an Effective Governance Process

A CSC is placing some parts of its IT operations – hence part of its business processes – in the hands of one or more CSPs. Adding this relationship creates a need for strong and detailed governance, inside the customer's organization, for the use of the cloud services.

The first part of the governance process involves the control and oversight of Steps 1-8 of this guide. The second part is the ongoing review of the use of each cloud service, to ensure that it meets business requirements and results in user satisfaction. The governance process should also deal with change: changing business and user requirements, changes made by the CSPs to their offerings, availability of new technologies, as well as changes in laws and regulations (affecting, for example, data protection).

The cloud customer should also adopt systems engineering standards and best practices when developing applications and systems -- small or large – for the cloud. These must be taught, enforced, and reviewed to make sure they are adopted. Consistent approaches to adopting IT systems with a focus on reliability, availability and serviceability will reduce the risk of down time.

The governance process should include a review of CSA service levels as they relate to the contractual obligations the CSC already has toward its own customers.

Finally, the governance process is a cross-functional one, not just the responsibility of IT. The CSC must include members of their risk management, legal, compliance, sourcing and other business functions in the design and operation of the governance process.

Table 8 below highlights the key elements required to operate a successful governance process between the CSC and CSP.

Table 8 – Elements of the Governance Process

Element	Description
Assessment of Service Levels	Periodic assessment of actual service levels vs. those guaranteed by the CSA: ● Monitoring reports from the CSP ● Monitoring reports from the CSC's own cloud service administrators ● Explanation of any discrepancies
Compliance Assessment	Where the compliance of the cloud service to specific standards or regulations is important to the customer, it is necessary for the customer's governance process to periodically check that the cloud service still has valid proof of compliance.

Element	Description
Service Failure Reports	Reports of any service failures or incidents which affect • Service availability • Security, particularly security breaches • Protection of personal data
Service Change Notifications	Any change notifications from the CSP that relate to the cloud services being used (change of APIs, change of functionality, change of service level objectives, change or cloud service pricing, change of terms in the CSA).
Problem Reports	Key indicators should be tracked to ensure that the CSA criteria are being met and that the users of the service (employees as well as the CSC's own customers) are experiencing the agreed service levels: • Total number of problems reported in the current reporting period, split by impact level • Statistics about the time to resolution of high-impact problems • Number of open problems at the end of the period, split by impact level, with statistics on the time they have been open • Number of problems closed (sorted by impact) with the time it took to resolve them • Number of problems not resolved within agreed time frames • Trends of the number of problems being reported (current period vs. previous ones) with the resulting resolutions.
Request Reports	Reports on (non-problem) requests made by the cloud service customer to the cloud service provider: • All requests made • Number of open requests • Statistics on the time to satisfy the requests
User Satisfaction Reports	Reports on user satisfaction with the cloud service(s)
Cloud Service Lifecycle Events	Periodic review of the CSP's service roadmap • End of life of used services • Introduction of new services or features • Follow-up of CSC requests to change service functionality

The cloud service customer must periodically review the elements described in table 8 and decide on an appropriate course of action if the cloud services do not meet the terms of the agreement or do not meet business requirements.[10] How the review is performed is a decision for the customer, and it is likely to depend on the size and structure of its organization. A degree of formality and record keeping is advisable since in some cases, evidence may need to be prepared for presenting to the CSP, especially

[10] If the previous steps have been performed correctly, the terms of the CSA and the business requirements of the CSC should align. If the CSA is fulfilled while the business requirements are not met, the CSA should be revisited.

if there are issues that the CSP is likely to be unaware of or to contest, potentially leading to a dispute about payments and penalties.

Issues that arise regarding compliance with the CSA require different courses of action depending on the nature of the issue(s).

- Some breaches of the CSA may trigger remedy terms, including some level of compensation to the customer – but this is often not automatic. The CSC often needs to formally raise a request with the CSP in order to obtain the remedy.
- More serious breaches or incidents will require stronger action on the part of the CSC. Escalation (see below) may trigger discussions between senior management on both sides. Alternately, the CSC may need to trigger the termination process in order to switch to another CSP.
- For problems that require higher management awareness, it is the responsibility of those involved in the governance process to advise their superiors of the status of a particular issue, and brief them about the facts and the desired remedies.

Escalation Process

While most aspects of governance will be handled through regular reports, reviews, and scheduled periodic meetings, some problems will fall outside the normal management process and will need special focus to ensure a timely resolution. An example of such a process exception is when a major outage (i.e., one that causes loss of service) occurs. This cannot wait for a periodic meeting and requires immediate notification of the management chain.

While we use the term *escalation* to describe this, this specifically means upward communication for awareness of a serious situation – not upward delegation of responsibility for the resolution of the problem.

Table 9 below highlights the types of escalations that can be invoked, general criteria to initiate escalation, and the overall objectives of escalation.

Once an escalation has been initiated, the goal is to ensure that both chains of management (the CSC's and the CSP's) understand the problem, its impact, and the agreed action plan for resolution – including containment of the problem, especially if the problem impacts services to a customer of the CSC.

Escalation should not be considered a last resort in the problem management process. Escalation should be used as an early warning activity to raise management awareness of a potential problem before it becomes critical. Escalation is just another tool to ensure that the expected service levels are delivered to the users and customers of the CSC.

Finally, if an escalated problem can still not be resolved, then other terms of the CSA can be brought to bear to force resolution. One of the outcomes of continuous violations of the CSA can be termination of the agreement. Rigorous documentation of what happened as part of the governance process (reports, meetings, escalation actions, correspondence, etc.) is critical to support the termination process.

Table 9 – Escalation Types, Objectives and Guidelines

Consideration	Description

Types	• **Immediate**
	  ○ A critical business impact is identified
	  ○ There is significant impact to a service offered by the CSC to its own customers
	• **As required**. Typically, this occurs after a review when it is discovered that key service levels are not met, such as:
	  ○ Number of open problems
	  ○ Duration of problem resolution
	  ○ Increasing trend of reported problems without a satisfactory resolution

Criteria to Initiate Escalation	• The problem critically impacts either an organization-wide internal service or a customer-facing service – the service is either unavailable or is significantly degraded. • A problem with significant impact has missed the agreed resolution time objective. • Independent of impact, problems are routinely not being closed within the expected time frames. • The number of reported or open problems is increasing, with no agreed plan to reverse the trend. • Requests to the CSP to participate in root cause analysis or problem resolution in an associated system or tool are ignored.
Objectives	• Raise management awareness to avoid surprises (allows senior management to be in control of the situation, especially when external customers are impacted). • Achieve a bilateral agreement to the action plan to resolve a problem. • Gain management agreement to support the plan, including additional resources when required.

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 CSP's and CSC's respective responsibilities in case the relationship terminates, whether prematurely or at the expected end of the contract. A premature initiation of the exit process may be due to the CSP's inability to deliver the required service levels, or to it's going out of business. A regular exit may occur when, before the end of the contract, the CSC "re-bids" the cloud services and selects a new provider.

Regardless of the reason, a *customer exit plan* should always be prepared in advance, when discussing the CSA, of which it an integral contractual annex. A clearly defined exit process includes detailed procedures to securely and rapidly transfer customer data and applications to another service, ensuring business continuity. It should specify metrics to ensure that the cloud provider is effectively implementing these procedures.

The most important aspect of any exit plan is the transmission and preservation of customer data, which is critical to achieving business continuity. In addition to the legitimate need of the customer to take possession without delay of the data it owns, recent laws and regulations such as the EU's GDPR *require* that customer data be transferable to alternate cloud service providers.

In addition, customers must ensure that their data is completely removed from the provider's environment once the exit process is complete. Failure to do this might violate data protection clauses, or expose the CSC to legal discovery challenges in the future. In some cases, data residency issues could also arise – for example, if a domestic CSP is acquired by a foreign company, which might migrate all the data it holds to another jurisdiction as part of a consolidation strategy. [4][11]

Customers should review the following aspects of the exit clause of a CSA:

- The level of CSP assistance with the exit process, and what actions if any may be billed separately. In most cases, there should be no additional cost associated with the exit process.
- The CSP should be responsible for removing customer data from their IT environments, or at least helping the customer extract and erase their data by providing clear and concise documentation. The CSC and CSP need to agree on what level of assurance or verification is possible and required.
- The format of the data transmitted from the provider to the customer should be specified in the CSA and should conform to standard data formats whenever possible, to enable portability to a new service. The transmission of the data from the CSP to the CSC or a new CSP should use standard packaging and data transfer techniques.
- The CSA should specify that all data and information belonging to the customer is maintained for a specific time period after transition, and is completely removed after that time.
    - The typical time period is 1-3 months, which gives the customer sufficient time to find a new provider and to continue receiving service from the current provider in the interim.
    - The time period should be explicitly documented in the CSA and only with the customer's written approval should data be removed or destroyed before that time.
    - At the completion of the exit process, the CSP should provide written confirmation that all customer data has been removed from its facilities.
    - The issue of backup media is sensitive: in a multi-tenant cloud service, it may not be possible to selectively erase customer A's data from a medium that also contains data from customers B and C. The disposition of backup data should be addressed.
- The CSA must provide appropriate business continuity protection during the exit process. Services must not suffer degradation during the transition process (from the invocation of the termination clause to the actual end of the service).

[11] ISO is elaborating a new data handling standard, due for publication in 2019.

The bottom line is that customers should undertake due diligence when evaluating and ultimately selecting a cloud provider. A trustworthy cloud provider should be prepared to provide customers on a fair and effective exit strategy.

Summary of Keys to Success

Table 10 summarizes the critical success factors, extracted from the ten steps presented in this Practical Guide, for any customer organization evaluating and comparing CSAs from competing cloud providers.

Table 10 – Summary of Keys to Success

Key to Success	Summary
Review internal policies and processes	• Identify key processes and policies that will be affected by cloud adoption. • Contracting and reporting are key areas for review.
Develop a strong business case and strategy for cloud computing environment	• Assess the criticality of the services targeted for deployment in the cloud. • Determine functional and non-functional requirements for each service (performance, availability, security, privacy, etc.) • Understand the laws and regulations affecting the data maintained in the cloud. • Identify key performance metrics for each service.
Assess provider's CSA against functional and non-functional requirements	• Based on the criticality of the service being deployed in the cloud, determine if the CSP's standard CSA meets the above requirements. • If not, determine if the CSP is willing to negotiate key aspects of the CSA to align it with the CSC's business strategy. Otherwise, seek alternative providers who more closely address the requirements. If none can be found, consider keeping the service on premises.
Determine how to monitor CSA performance	• Understand the management process defined in the CSA. • Ensure the CSA includes the ability to monitor, assess and react to key performance measurements. • Understand the notification process when service issues arise, including the method and timeliness of notifications, and the CSP's processes to assign severity and priority levels to issues. • Review and discuss the escalation process for high-impact incidents or violations of the SLA. • Understand the remedies offered and liability limitations imposed by the CSA.

Key to Success	Summary
Ensure an adequate disaster recovery plan can be defined and executed	• Review the impact of moving to the cloud on the DR strategy. • Understand the roles and responsibilities documented in the CSA. • The CSC must understand the CSP's ability to support its data preservation strategy, which includes criticality of data, data sources, scheduling, backup, restore, integrity checks, etc., and ensure alignment of the CSC's DR procedures to those of the CSP, to avoid gaps and overlaps. • Understand that in the end, the CSC bears most of the risk of disaster scenarios that severely limit the ability of the CSP to deliver the service, and is responsible for implementing most of the data preservation strategy. • CSCs should specify appropriate recovery time objectives (RTOs) for each level of criticality of data, and make sure the CSA allows these objectives to be met. • Customers should test and verify the disaster recovery plan prior to production deployment. • Cloud customers should consider purchasing additional risk insurance if the costs associated with recovery are not covered under their organization's umbrella policy for IT services or operational risk riders.

Ensure support for an efficient exit process	• The goal of the exit plan is to ensure minimal business disruption for the customer when the relationship with the cloud provide terminates, prematurely or not.
	• The exit plan should be considered in advance, during the evaluation and selection of a CSP.
	• Roles and responsibilities must be clearly documented in the CSA. In many cases, the CSC may be responsible for initiating most of the exit process steps.
	• The CSA should be support provisions needed by the customer, such as:
	  ○ The CSC should be able to terminate the agreement at any time, without penalty, provided sufficient notice is given to the CSP.
	  ○ Customer data should be recoverable using standard formats and transfer methods to facilitate portability to a new CSP.

Emerging standards in the following areas will help improve the ability for customers to evaluate and compare the service levels offered by different providers:

- Standards that create consistent ways to describe services and associated terms, including price.
- Standardized metrics that allow customers to effectively track and compare CSA performance.
- Standardized security and regulatory compliance requirements to identify control points for risk management.
- Standards that enable coordinated end-to-end CSA management for both cloud customers and cloud providers (including in hybrid or multi-cloud environments).

Cloud computing offers a value proposition that is different from traditional enterprise IT environments. Uncertainty about service levels and other implications of a CSA can create fear, uncertainty and doubt, which impede adoption. With a clear checklist and proper focus on key success factors, cloud customers are able to effectively review and compare CSAs from different cloud providers, then manage the ongoing relationship with the selected provider(s), to ensure that the promise of the cloud is realized.

References

[1] Cloud Standards Customer Council (2016). *Public Cloud Service Agreements: What to Expect and What to Negotiate, Version 2.0.* https://www.omg.org/cloud/deliverables/public-cloud-service-agreements-what-to-expect-and-what-to-negotiate.htm

[2] Cloud Standards Customer Council (2017). *Security for Cloud Computing: 10 Steps to Ensure Success, Version 3.0.* *https://www.omg.org/cloud/deliverables/security-for-cloud-computing-10-steps-to-ensure-success.htm*

[3] Cloud Standards Customer Council (2016). *Cloud Security Standards: What to Expect and What to Negotiate, Version 2.0.* https://www.omg.org/cloud/deliverables/cloud-security-standards-what-to-expect-and-what-to-negotiate.htm

[4] Object Management Group, Cloud Standards Customer Council (2017). *Data Residency Challenges: A Joint Paper with the Object Management Group.* https://www.omg.org/cloud/deliverables/data-residency-challenges.htm

[5] Cloud Standards Customer Council (2017). *Practical Guide to Cloud Computing, Version 3.0.* https://www.omg.org/cloud/deliverables/practical-guide-to-cloud-computing.htm

[6] National Institute for Science and Technology (2011): *Special Publication 800-145: The NIST Definition of Cloud Computing.* https://csrc.nist.gov/publications/detail/sp/800-145/final

[7] ISO/IEC 17789:2014: Information Technology - Cloud Computing - Reference Architecture. http://standards.iso.org/ittf/PubliclyAvailableStandards/c060545_ISO_IEC_17789_2014.zip

[8] FIRST: *Common Vulnerability Scoring System v3.0: Specification Document*. https://www.first.org/cvss/cvss-v30-specification-v1.8.pdf

[9] OASIS (2018). *Topology and Orchestration Specification for Cloud Applications (TOSCA).* https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=tosca

[10] Wikipedia: *Web application Archive (WAR)*. https://en.wikipedia.org/wiki/WAR_(file_format)

[11] The DMTF[12] (2018). *Open Virtualization Format (OVF) standard.* https://www.dmtf.org/standards/ovf

[12] Cloud Security Alliance. *Security Guidance for Critical Areas of Focus in Cloud Computing Version 3.0* (2011). https://downloads.cloudsecurityalliance.org/assets/research/security-guidance/csaguide.v3.0.pdf

[12] Formerly known as the Distributed Management Task Force.

[13] National Institute for Science and Technology (2017): *Special Publication 800-53 Rev. 5 (Draft)*[13]*: Security and Privacy Controls for Information Systems and Organizations.* https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/draft

[14] National Institute for Science and Technology (2016): *Special Publication 800-171 Rev. 1: Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations.* https://csrc.nist.gov/publications/detail/sp/800-171/rev-1/final

[15] International Organization for Standardization: *ISO/IEC 27001:2013, Information technology – Security techniques – Information security management systems – Requirements.* https://www.iso.org/standard/54534.html

[16] Center for Internet Security: *CIS Controls™ V7.* https://learn.cisecurity.org/20-controls-download

[17] Cloud Security Alliance (2018): *Cloud Controls Matrix V3.0.1.* https://cloudsecurityalliance.org/artifacts/csa-ccm-v-3-0-1-11-12-2018-FINAL/

[18] International Organization for Standardization: *ISO/IEC 27004:2016, Information technology – Security techniques – Information security management – Monitoring, measurement, analysis and evaluation.* https://www.iso.org/standard/64120.html

[19] National Institute for Science and Technology (2008): *Special Publication 800-55 Rev. 1, Performance Measurement Guide for Information Security*. http://csrc.nist.gov/publications/nistpubs/800-55-Rev1/SP800-55-rev1.pdf

[20] Center for Internet Security: *CIS Consensus Security Metrics.* https://learn.cisecurity.org/benchmarks

[21] International Organization for Standardization: *ISO/IEC 27017:2015, Information technology – Security techniques – Code of practice for information security controls based on ISO/IEC 27002 for cloud services.* https://www.iso.org/standard/43757.html

[22] Cloud Security Alliance (2017): *Consensus Assessments Initiative Questionnaire (CAIQ) V3.0.1*. https://cloudsecurityalliance.org/artifacts/consensus-assessments-initiative-questionnaire-v3-0-1/

[23] Cloud Security Alliance: *Security Trust and Assurance Registry (STAR).* *https://cloudsecurityalliance.org/star/registry/*

[24] United States Congress: *Clarifying Lawful Overseas Use of Data Act (CLOUD Act).* March 2018. https://www.congress.gov/bill/115th-congress/senate-bill/2383/text

[25] Cloud Standards Customer Council (2017). *Practical Guide to Cloud Management Platforms.* https://www.omg.org/cloud/deliverables/practical-guide-to-cloud-management-platforms.htm

[26] American Institute of Certified Public Accountants (AICPA, 2018): *SOC 2® Reporting on an Examination of Controls at a Service Organization Relevant to Security, Availability, Processing Integrity, Confidentiality, or Privacy.* https://www.aicpastore.com/SOC/reporting-on-controls-at-a-service-organization-re/PRDOVR~PC-0128210/PC-0128210.jsp

[27] Uptime Institute: *Tier Classification System.* https://uptimeinstitute.com/tiers

[13] Publication of the final version is expected by Summer 2019.

[28] International Organization for Standardization: *ISO 22301:2012, Societal Security – Business continuity management systems – Requirements.* https://www.iso.org/standard/50038.html

[29] National Institute for Science and Technology (2010): *Special Publication 800-34 Rev. 1: Contingency Planning Guide for Federal Information Systems.* https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/final

[30] ASIS International: *ANSI/ASIS ORM.1-2017, Security and Resilience in Organizations and Their Supply Chains.*[14] https://www.asisonline.org/publications/sg-security-resilience-in-organizations-and-their-supply-chains---requirements-with-guidance-standard/

[31] International Organization for Standardization: *ISO 27031:2011, Information technology – Security techniques – Guidelines for information and communication technology readiness for business continuity.* https://www.iso.org/standard/44374.html

[32] International Organization for Standardization: *ISO 24762:2008, Information technology – Security techniques – Guidelines for information and communication technology disaster recovery services.* https://www.iso.org/standard/41532.html

[14] Supersedes ASIS SPC.1.