Practical Guide to Cloud Service Agreements
Version 2.0

April, 2015

Contents

Acknowledgements

The major contributors to this whitepaper are: Claude Baudoin (cébé IT & Knowledge Management), Beniamino Di Martino (Second University of Naples), Marlon Edwards (Hoboken Consulting Group, LLC), Mike Edwards (IBM), David Harris (Boeing), Ryan Kean (The Kroger Co.), Yves Le Roux (CA Technologies), George Malekkos (Powersoft Computer Solutions Ltd), John McDonald (CloudOne Corporation), John Meegan (IBM), Gerry Murray (Fort Technologies), Massimiliano Rak (Second University of Naples), Dave Russell (IBM), Karolyn Schalk (Garden of The Intellect LLC), Gurpreet Singh (Ekartha), Annie Sokol (NIST), Joe Talik (AT&T), Salvatore Venticinque (Second University of Naples), Steven Woodward (Cloud Perspectives).

Revisions

Much has changed in the realm of cloud computing service agreements since the original *Practical Guide to Cloud Service Level Agreements* whitepaper was published in April, 2012. Version 2.0 of the document includes the following updates:

- Terminology changes have been made; specifically, the term service level agreement (SLA) has been replaced by cloud service agreement (CSA) to reference the broad agreement that is established between cloud customers and providers. The term SLA is now used to reference that part of the broader CSA that deals specifically with service levels.
- The *Current CSA Landscape* section has been updated substantially to reflect current market dynamics.
- All ten steps in the *Guide for Evaluating Cloud Service Agreements* section have been updated to reflect current best practices. Significant changes have been made to steps 1, 5 and 9.
- References to cloud computing standards have been updated.
- References have been added to several CSCC whitepapers that have been recently published.

Introduction

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

CSAs are primarily written to set clear expectations for service between the cloud customer (buyer) and the cloud provider (seller), but should also exist between a customer and other cloud entities, such as the cloud carrier, the cloud broker and even the cloud auditor. This Guide focuses primarily on the CSA details between the cloud customer and cloud provider.

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.

"The Current CSA Landscape" section explains the dynamics that currently exists 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 steps that cloud customers should take to evaluate CSAs in order to compare multiple cloud providers or to negotiate terms with a selected provider. 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.

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 cloud customer and the cloud service provider. Because the cloud computing market is still developing, cloud customers should be aware that there may be a mismatch between their expectations and the cloud providers' actual service terms. 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).

It is common for disputes to arise over the structure of the agreements, thus cloud customers must pay close attention to the language and clauses of the CSA. Large cloud providers can be inflexible with their CSAs, while small cloud providers may seem more flexible, but tend to overpromise in order to obtain clients.

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 it uniformly adopted by the cloud industry: no standard nomenclature is used across the various cloud providers to specify their CSAs. Furthermore, cloud providers can modify their contract structure and terms at any time.

The *Customer Agreement* section of the CSA describes the overall relationship between the customer and provider. Since service management includes the processes and procedures used by the cloud provider, 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 providers.

An *Acceptable Use Policy (AUP)* is commonplace within a CSA. The AUP prohibits activities that providers 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 cloud providers. Although specific details of acceptable use will vary among IaaS, SaaS and PaaS providers, the scope and effect of these policies is 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 violations of these thresholds. Well-designed SLAs can significantly contribute to avoiding conflict and can facilitate the resolution of an issue before it escalates into a dispute.

To guarantee an agreed service level, service providers must measure and monitor relevant metrics. There is often a mismatch between the metrics collected and monitored by the service provider and the higher-level functional (or "end-to-end") metric relevant to customers. 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. 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 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 customers that decide to

switch providers following unsatisfactory service – and they clearly are meant to encourage the customer to stay with the current provider.

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

For cloud customers, 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 all other customers. Better CSAs will inevitably become a competitive factor. Eventually, customers of all sizes will be able to choose from a range of service terms that are more favorable and more flexible.

Guide for Evaluating Cloud Service Agreements

Before getting to the point of evaluating any CSA, customers 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* [2].

With this strategic analysis as a prerequisite, this section provides a prescriptive series of steps that should be taken by cloud customers to evaluate CSAs in order to compare multiple cloud providers or to negotiate terms with a selected provider. The following steps are discussed in detail:

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

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 and comparability across providers.

Step 1: Understand Roles & Responsibilities

From the cloud service customer perspective, one of the significant areas of risk involved with cloud computing is associated with the division of activities and responsibilities between the cloud service customer and the cloud service provider. It is necessary to have a full understanding of who is responsible for which activities to ensure that there are no gaps which could lead to problems when using cloud services.

The ISO/IEC 17789 cloud computing reference architecture standard[1] has 3 main roles for cloud computing:

- Cloud service customer
- Cloud service provider
- Cloud service partner

The cloud service provider and the cloud service customer are the most significant roles in the provision and use of cloud services while the cloud service partner is a party engaged in support of the activities of the cloud service customer and/or the cloud service provider.

There are a number of subroles of each of the major roles – the subroles are shown in Figure 1:

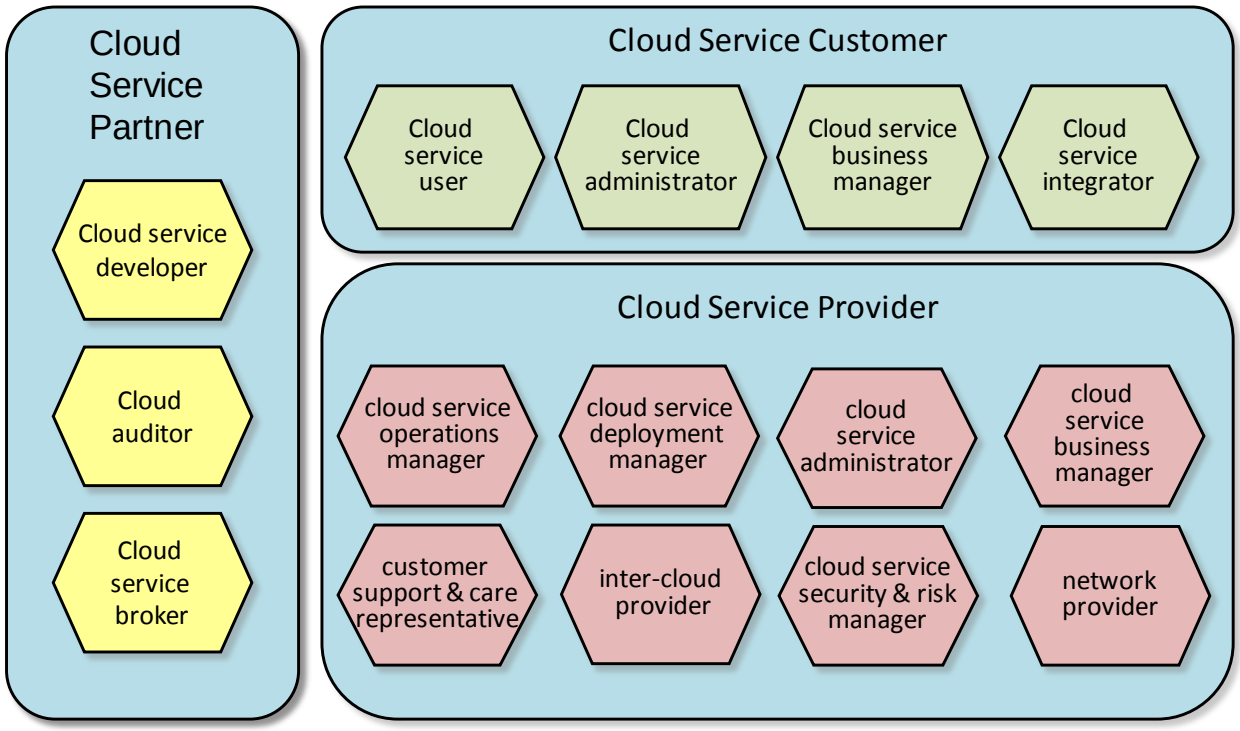

Figure 1: Cloud Computing Roles and Subroles

[1] See http://standards.iso.org/ittf/PubliclyAvailableStandards/c060545_ISO_IEC_17789_2014.zip for the ISO 17789 standard.

Each of the subroles in Figure 1 has a set of activities and responsibilities which are described in high-level terms in ISO/IEC 17789. There are also relationships between the subroles - for example, the cloud service administrator of the customer may interact with the customer support and care representative of the provider in cases where customer personnel experience problems using the cloud service.

Some of the subroles may appear in a CSA, or they may have a direct or indirect relationship to some aspects of the CSA. The subroles of the cloud service customer and the cloud service provider, in particular, are involved in the split of responsibilities that is typical for cloud services - the CSA should make clear statements about those responsibilities. Cloud service customers need to understand the activities and responsibilities of the various subroles and ensure that the CSA and its associated SLA contains appropriate commitments and service level targets to address those activities and responsibilities for the cloud service(s) covered by the CSA.

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 can include outages where the cloud service is unavailable, or may include cases where performance fails to meet stated service levels (for example, response times are too long). How such incidents are detected must be established - it may be the responsibility of the customer and the customer may need to put in place appropriate monitoring technology. It is also necessary to be clear about how incidents are then reported and tracked until resolved.

One partner role that is particularly relevant to the CSA and to SLAs is the cloud auditor. It is unlikely that the cloud service customer has direct insight into the operations of the cloud service provider, particularly regarding aspects such as security and the protection of sensitive data such as personally identifiable information (PII). It is typical for cloud service providers to offer assurances about these aspects of their cloud services through certifications or attestations which are provided by third party cloud auditors who inspect the cloud service provider's operations and issue reports 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 and are not limited to quantitative measures, but can include other qualitative aspects such as alignment with standards and data protection. It is strongly recommended that cloud service customers gain a solid understanding of the spectrum of CSAs that currently exist for cloud service providers in order to compare cloud services offered by different providers and assess 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.

It is important to recognize that the content of a CSA and associated SLA is likely to vary depending on the category of the cloud service. Considerations for an IaaS service offering compute and storage infrastructure are likely to be very different from those for a SaaS service that offers complete application functionality for some business functions. At the very least, the split of responsibilities between the provider and the customer are going to be different for these different cases, and this is necessarily reflected in differences in the CSA and SLA.

The following sections, which cover the cloud CSA evaluation steps in detail, each elaborate on the expected responsibilities of the customer and the provider for both business level and service level objectives. In order to make sound business decisions, it is important that customers understand what to expect from their cloud service provider. This, in turn, will help them clarify their own responsibilities and help them assess the true cost of moving to cloud computing.

Step 2: Evaluate Business Level Policies

Customers must consider the policy and compliance requirements relevant to them when reviewing a CSA since there are interdependencies between the policies expressed in the CSA and the business strategy and policies developed across the lines of business. The *data policies* of the cloud provider, as expressed in the CSA, are perhaps the most critical business level policies and should be carefully evaluated.

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

- customer requirements,
- 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.

Customers 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 at different times. The issue of jurisdiction takes on additional complexity when global compliance is taken into consideration and more than one cloud provider is used. In these instances the customer may have to coordinate negotiations between providers to ensure the necessary data management.

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 capturing and preservation of data is critical to maintaining the organizational memory of a business or the general user. Customers should therefore ensure they have an appropriate data preservation strategy that addresses redundancy within the system. • Cloud customers 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 service provider. • It must be possible to test the CSA to demonstrate the required level of service availability.

Data Policy	Description / Guidance
Data Location	• CSAs that cover locations under different jurisdictions are challenging. • Customers should consider how the CSA specifies where their data resides, where it is processed, and how this meets the various applicable regulations. Customers should also understand where the data is viewed or delivered, and whether this results in a transborder data flow with regulatory or tax implications. • For example, can the provider truly deliver a sound technical solution when sensitive data spans several jurisdictions with conflicting laws? Does the provider commit, in the CSA, to the specific location(s) where the customer's data will be stored? • If the provider reserves the right to add new locations or change data movement policies, will they give the customer advice notice? Preferably, will they obtain the customer's permission to relocated its data? • Is there a means to verify the current location of a data set?
Data Seizure	• Legal powers enable law enforcement and other government agencies to seize data under certain circumstances. Customers should ensure the CSA provides for sufficient notification of such events. • Customers should also ensure there are arrangements in place to make their data available in the event that their provider goes out of business. • in the event that the provider locks access to its systems because of a billing dispute or a security issue, the customer's date should not be "held hostage" while the issue is being resolved.
Data Privacy	• The provider's data privacy policy should be included in the CSA, and should ensure that the provider will conduct business in compliance with applicable laws on data privacy 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 customer's agents in its dealing with the provider (this is the narrow meaning in many existing Service Level Agreements), it also includes the privacy of the information that may be stored about the customer's own customers. • Refer to the Privacy section within Step 5 for more information.
Data Availability	• Assess whether the provider's maintenance schedules might interfere with business processes subject to external constraints, such as financial reporting or the business's hours of operation in certain regions.
Change Management and Notification	• The change management and change notification obligations of the provider should be carefully reviewed, especially the amount of time allowed to prepare for a change. The provider may also ask the customer to provide certain change notifications, which is a good opportunity to strengthen the customer's own change management policies.

In addition to data policies, there are a number of other business level policies expressed in the CSA that require careful evaluation. Uptime and availability are another area where customer requirements and policies may not match up with the language of the vendor, and where location and jurisdiction may come into play. For example, if the uptime guarantee is for "regular business hours," then organizations with multiple locations in different time zones need to clarify whether the guarantee covers only the headquarters location or all regions. Similarly, "week-ends" or "holidays" have different meanings in different countries. For some multinational customers with offices in all continents, the sun literally never sets on their empire, and the provider may not be ready to commit to supporting them 24x365.

All of these policies will impact and influence the customer's cloud strategy and business case. In many cases, these policies, as defined in the CSA, are non-negotiable and are similar across different cloud providers. However, there will be instances where some of these policies can be negotiated and/or some of these policies differ sufficiently across different cloud providers to warrant careful consideration from customers.[2]

Table 2 below highlights the critical business level policies that need to be considered and 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.[3] The CSA should clarify: o What constitutes excused or excluded performance o Escalation procedures o How service-level bonuses and penalties are administered o Remedy circumstances and mechanisms • Guarantees should be expressed as a measurable number, for example a percentage such as 99.999% for service availabliity, denoting the amount of time the service is guaranteed to be working. Other guarantees will refer to matrics in other units, such as time-to-repair in minutes, etc. • Availability measures need to include the measurement window.

[2] Refer to the already cited *Public Cloud Service Agreements: What to Expect and What to Negotiate.* [1]

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

Policy	Description / Guidance
Acceptable Use Policy	• The acceptable use policy will clearly describe how the customer may use the service and the agreement generally will describe what actions the provider 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 cloud provider. • Customers 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,
List of Services Not Covered	• The CSA will state under what conditions and with which described services the customer is supported. The CSA may also state what is excluded and what constitutes illegal use. • Customers should look for explicitly stated exceptions and understand why the provider has excluded them.
Excess Usage	• Providers operate business models to drive revenue. While elasticity is a fundamental benefit of using the cloud, customers may find that usage above their contracted thresholds will incur high incremental rates which can be punitive and disrupt their budgets. • Customers 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, in order 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. • From a CSA compliance perspective, it is important for customers 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. Provider 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" service providers are sensitive to poor credit control and are likely to be more diligent in suspending service. • Equally, the customer needs to be diligent in obtaining service credit payments for outages.
Governance / Versioning	• Provider services evolve. New features may be added, others will go out of warranty, and some may persist indefinitely. Where the assumptions or conditions under which the CSA was initially accepted are changed, the customer should review the impact on their specific situation. • A good provider will maintain a proactive policy of advising customers of changes to their CSA and practice version control. • Customers should ensure that there is a mechanism to inform them of changes and, if not, amend their contract to put the onus on the provider to provide reasonable advance notice of updates.

<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 relocate to another provider if necessary.
Providers may write in their contracts an automatic renewal clause that kick in in the absence of a 90-day cancellation notice before the contract’s anniversary date. It is common for customers to overlook this deadline and be obligated to renew the contract without having had a chance to negotiate changes or even cancel the service.
Customers should read the terms and conditions of the renewal arrangements, and consider the conditions under which a provider may vary the service terms (or revise prices) upon renewal.

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

Customers should consider the potential need to transfer an agreement in the event their business is sold.
Conversely, if the provider’s business is acquired, the customer may not wish to so business with the new owner, and should have the option to migrate to a new service without penalties.
Customers may operate several accounts with a provider and want to offset account credits between accounts. Is this supported in the provider’s contract terms?

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

Customers must follow the provider’s rules to report problems, in order to ensure that the support terms specified 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/Production cloud servers inaccessible</td><td>1 hour, Provider’s executive notified of issue</td><td>1hr</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>1day</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/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>
</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 downtime per year. CSAs may state that this does not include "planned maintenance." Thus, the provider may have a service outage for 8.5 hours. plus maintenance time, and the customer is not entitled to compensation under the CSA. This highlights the importance of defining the measurement window. If the availability percentage is measured each month, this allows 12 outages nut each of them cannot last more than 42 minutes without triggering a penalty.
Subcontracted Services	• Providers sometimes include in their CSA a clause that the CSA of an upstream (subcontracted) provider will govern the services provided by the sibcontracot, and that the only available penalties are those from the upstream provider even though its CSA may be less rigorous. The customer's expectation, based on reviewing the CSA of their immediate provider, may thus be violated. • Therefore, the customer should ensure that the immediate provider CSA states unambiguously that its CSA applies to the complete service, regardles whether parts of the service come from third parties.
Licensed Software	• Cloud services may include third party licensed software which is sold on a monthly licensed basis under a service provider license agreement. Such software is updated regularly by its manufacturer. • Providers may opt to pass the responsibility for updating the licensed software over to the customer once they have started to use the service. This absolves the provider of the risk of disrupting the customer's operation through an unforeseen software conflict or bug. • Alternately, the provider may "push" the update, in which case the CSA should require that the customer be given advance notice of the update. The customer should have the ability to opt out, or at least to defer the update. However, the supplier may be unwilling to continue to support older versions indefinitely, and there should be a legitimate exception for updates that correct serious security vulnerabilities.
Industry Specific Standards	• Regulated industries, like government, financial services, and healthcare, are subject to specific and often quite onerous standards which must be addressed in the CSA and implementation. • Customers who operate in these regulated industries should ensure that the their legal team is fully involved on the negotiation of the CSA.
Additional Terms for Different Geographic Region or Countries	• Customers should consider the provider's origins and primary market. Detailed refinements to the home market CSA may be required to properly cover customers who are located in remote markets. • Data protection legislation is one aspect of this, but customers should not limit their examination of the agreement to this sole aspect.

Step 3: Understand Service and Deployment Model Differences

Services offered by cloud providers typically fall into one of the three major groups of service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). For each category, there are significant differences in the levels of cloud resource abstraction, service level objectives, and key performance indicators that will potentially be included in a CSA. In addition, the level of clarity varies significantly for each service model. 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 service provider to meet the agreed upon target levels of quality

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

Table 3 highlights the different CSA considerations for each of the cloud service models.

Table 3: CSA Considerations for Service Models

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 services levels. • Metrics are well understood across the IaaS abstractions (compute, network, and storage). Customers should expect to find a subset of the following metrics in their cloud SLA. o Compute metrics: *availability, outage length, server reboot time* o Network metrics: *availability, packet loss, bandwidth, latency, mean/max jitter* o Storage metrics: *availability, input/output per second, max 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. • Customers must distinguish between IaaS development environments and IaaS production environments when reviewing their cloud IaaS service agreements. IaaS production environments will typically require more stringent service level objectives than IaaS development environments.

	• Network metrics in a cloud SLA generally cover the cloud provider's data center connectivity to the Internet as a whole, not to any specific provider or customer. • There are several standardization efforts within the IaaS space which help describe and manage the services offered at this level.[4] Whenever possible, customers should ensure the CSA includes provisions requiring their cloud providers 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, customers should consider tradeoffs in flexibility, control, and ease of use to determine which approach best meets their business needs. o Integrated solutions are web accessible development environments which enable developers to build an application using the infrastructure and middleware services supported by the cloud provider. Management of the application and its execution is primarily controlled by the cloud provider. Typically, service developers only have access to a provider-defined set of APIs which offer limited control on the coordination of code execution. o Deploy-based solutions enable deployment of middleware on top of resources acquired from an IaaS cloud provider, offering deployment services to the customers which automate the process of installation and configuration of the middleware.[5] These PaaS solutions offer a rich set of management capability including the ability to automatically change the number of machines assigned to an application, and self-scaling according to the application's usage. • At a minimum IaaS SLA's should roll into PaaS SLA's. • Customers must distinguish between PaaS development environments and PaaS production environments when reviewing their cloud PaaS service agreements. PaaS production environments will typically require more stringent service level objectives than PaaS development environments. • Standards are emerging to help identify PaaS services offered by cloud providers and standard interfaces for communicating with PaaS providers to provision or manage PaaS

[4] 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).

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

	environments. Standards, like OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA)[6] 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. • Customers should ensure their CSA includes support for open standards, as they become available, to reduce vendor lock in.
SaaS	• Customers should insist on flexible CSAs that are measurable against their objectives, not the cloud providers' reporting needs. • Given the wide variation of services provided at the SaaS level, it is difficult to provide a comprehensive and representative list of SaaS service level objectives for customers to look out for in their CSAs. • Customers 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. • Customers should ensure that data maintained on the provider'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*. Customers 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* [2] for considerations on selecting a deployment model.

Table 4 highlights the different CSA considerations across the deployment models.[7]

[6] Refer to http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=tosca for details on TOSCA.

[7] The Community deployment model is not called out explicitly in the table since it is very similar to the Public deployment model.

Table 4: CSA Considerations for Deployment Models

Deployment Model	CSA Considerations
Private (On-site)	• CSA considerations for Private (On-site) 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, customers 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) are similar to Private (On-site) except cloud services are now being provided by an external cloud provider. The fact that IT resources from the provider are dedicated to a single customer mitigates potential security and availability risks. • Customers should ensure the CSA specifies security techniques for protecting the provider's perimeter and the communications link with the provider. • Customers should consider the criticalness of the service being deployed to justify the added expense of this model over the Public model.
Public	• CSA considerations for the Public model are greater than the Private (Outsourced) model since the provider's IT resources are now shared across multiple customers. • As a result, customers should carefully review the CSA to understand how the provider 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. Customers should also ensure the CSA provides adequate methods and processes for ongoing measurement.
Hybrid	• CSA considerations for the Hybrid model are similar to the Public model with the increased likelihood for unique integration requirements between cloud and enterprise services. • Customers 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 on the 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 cloud provider, for example:

- A web application is deployed on an application server as a Web application ARchive (WAR) file.[8]
- 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 one of many different formats. Adoption and support for standards like the Distributed Management Task Force (DMTF) Open Virtualization Format[9] (OVF) is recommended.

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 cloud provider. Typical performance considerations include availability, response time and processing speed, but they can include many other performance and system quality perspectives.[10] Cloud customers must decide which measures are most critical to their specific cloud environments and ensure these measures are included in their SLA.

Performance statements that are important to the cloud customer should be measureable 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 potential metrics not covered in this document. While many of the metrics may already be supported by your cloud provider; they may interpret the definition differently than you do. An agreed definition in the context of a specific cloud solution is critical. Some calibration may be required if a measurement captured by a provider does not exactly match the definition included as part of the SLA.

---

[8] Refer to http://java.sun.com/j2ee/tutorial/1_3-fcs/doc/WCC3.html for more information on the WAR (Web application ARchive) file.

[9] Refer to http://www.dmtf.org/standards/ovf for more details.

[10] Additional system quality measures that could be included in service performance include accuracy, portability, interoperability, standards compliance, reliability, scalability, agility, fault tolerance, serviceability, usability, durability, etc.

Industry standards should be used when possible to improve consistency. For instance, IEEE has good measurement definitions and categorizations for activities such as maintenance.[11]

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 (example)	Storage Availability (example)	Service Response Time (example)
Metric Name in SLA	Network Percentage Available Critical Business Hours	Storage Percentage Available	Service XXX Response Time in a Given Hour; Service YYY Response Time 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 XXX and YYY, which are PaaS reuseable services that will be invoked by our applications.
Collection Method	Machine	Machine	Machine
Collection Description	Using the DMTF, OGF[12], 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 XXX and YYY service invoked, the response time is collected every five minutes.

[11] 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.

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

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 service. This is exclusively a measure of network availability.	No reference to quality or availability of specific service. This is exclusively a measure of storage availability.	No individual service reporting is needed (for example, 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 to 12PM GMT Monday thru Friday.	Storage availability shall be 99.9% or higher	Response time for XXX service shall be less than 500 ms, YYY service less than 200 ms.
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), 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" 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 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[13] and Privacy Requirements

Security controls in cloud computing are, for the most part, no different than security controls in any IT environment. 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 from traditional IT solutions. Refer to the *Security for Cloud Computing: 10 Steps to Ensure Success* [3] whitepaper for details on security requirements for cloud computing.

There are two asset categories that require security and privacy consideration for cloud computing:

- Information (which belongs to the customer but has been moved into the provider's cloud)
- Applications, functions or processes (being executed In the cloud to provide the required service to the customer)

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

In order to determine which 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. An employee of our cloud provider accessed the asset?
3. The process or function was manipulated by an outsider?
4. The process or function failed to provide expected results?
5. The information was unexpectedly altered?
6. 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.

[13] 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" and quotes portions of this document. The document is available at http://www.cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf.

Table 6: Key Security Considerations for CSAs

CSA Security Considerations	Strategic Activities
Assess asset sensitivity and the operational security requirements of applications	• Complete an assessment of the confidentiality, integrity, and availability requirements of the assets. • Complete a threat risk assessment and privacy risk assessment. • Address application 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.

Put in place restrictions against unauthorized asset movement and accidental 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:
	  o Monitor for large internal data migrations with database activity monitoring (DAM) and file activity monitoring (FAM)
	  o 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 volumes 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.

Establish and track security metrics	• Metrics and standards for measuring performance and effectiveness of information security management should be established prior to moving to cloud computing. • 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.[14]
Assess the cloud provider's security capabilities	• Assess the cloud provider's level of security and its maturity. • If compliance to a normative standard (e.g. ISO 27002/27017[15]) is asserted, then verify the compliance certificate and its validity. • Look for verifiable evidence of resource allocation, such as budget and manpower to sustain the compliance program • Verify internal audit reports and evidence of remedial actions for the findings
Assess the cloud provider's security governance	• Assess the provider's security governance processes and capabilities for sufficiency, maturity, and consistency with the customer's information security management processes. o The provider's information security controls should be demonstrably risk-based and clearly support these management processes. o Where a provider cannot demonstrate comprehensive and effective risk management processes in association with its services, customers 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 provider's guarantees adequately address your security requirements.[16]

[14] Refer to the following resources for specific information on security metrics: ISO 27004:2009 (see http://www.iso.org/iso/catalogue_detail.htm?csnumber=42106), NIST Special Publication 800-55 Rev.1, Performance Measurement Guide for Information Security (see http://csrc.nist.gov/publications/nistpubs/800-55-Rev1/SP800-55-rev1.pdf), and CIS Consensus Security Metrics v1.1.0 (see http://benchmarks.cisecurity.org/en-us/?route=downloads.show.single.metrics.110).

[15] Visit http://www.iso27001security.com/html/27017.html for details on ISO/IEC 27017.

[16] The Cloud Security Alliance Consensus Assessments Initiative Questionnaire (CAIQ) provides a set of questions a cloud consumer and cloud auditor may wish to ask of a cloud provider. Refer to http://cloudsecurityalliance.org/research/cai/ for details.

Audit the cloud provider's security CSA compliance	• A "right to audit" clause in a CSA gives customers the ability to audit the cloud provider, which supports traceability and transparency. • Use a normative specification in the "right to audit" clause to ensure mutual understanding of expectations. • In time, this right should be supplanted by third-party certifications (e.g., driven by ISO/IEC 27002/27017). If the cloud provider is not willing to subject itself to a customer audit, it should propose to use a trusted third-party that follows a normative standard.

Providers should notify consumers of the occurrence of any breach of its system, regardless of the parties or data directly impacted. The provider 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, consumers should require the provider to indemnify them 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 cloud provider from being sued for the consequence of customer actions).

Privacy

In many countries throughout the world, numerous laws, regulations, and other mandates require public and private organizations to protect the privacy of personal data 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 cloud service provider) 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 cloud customer should consider alternate means of achieving their goals including seeking a different provider, or not sending sensitive data to the cloud. For example, if the customer wishes to send HIPAA-covered information to the cloud, the customer will need to find a cloud service provider that will sign a HIPAA business associate agreement or else not send that data to the cloud.

Preservation of information, included in some privacy regulations, can require that large volumes of data be retained for extended periods. What are the ramifications of this under the CSA? What happens if the preservation requirements outlast the terms of the CSA? If the customer preserves the data in place, who pays for the extended storage and at what cost? Does the customer have the storage

capacity under its CSA? Can the customer effectively download the data in a forensically sound manner so it can preserve it offline or near-line? These are some of the privacy related questions that need to be addressed in the CSA.

The reverse risk may also exist: the backup and disaster recovery policies of a cloud provider may cause copies of data or code to be retained beyond the retention period intended by the customer. 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 cloud provider and subpoena the provider. In another scenario, a cloud customer 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.

Refer to the *Security for Cloud Computing: 10 Steps to Ensure Success* [3] whitepaper for an overview of the privacy requirements for cloud computing and worldwide privacy regulations that currently exist.

Step 6: Identify Service Management Requirements

The fundamental goals of any cloud computing environment are to reduce cost, improve flexibility and increase reliability of the delivery of a service. Critical to meeting these goals is a uniform, straightforward, transparent and extensible system for managing and monitoring cloud services. In this section we will outline some key things to consider in the area of service management when entering into a service agreement with a cloud computing provider.

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. Although the standards for CSA language for service management are evolving, it is of upmost importance to include provisions for the considerations outlined below in your agreements.

Auditing

First and foremost in ensuring manageability of cloud services is a methodology for auditing and reviewing those services. This helps discern between providers 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 objective of any CSA terms in the area of auditing is multi-fold:

1. Provide you with an unbiased assessment of your ability to rely on the service provided
2. Assess the depth and effectiveness of the provider's internal systems and measures
3. Provide tools to compare quality levels with other competing providers
4. Ensure the openness needed to allow continuous review and improvement
5. Uncover issues in your own organization's ability to interface with the provider and provide uninterrupted services

This last objective is especially important. Many documented challenges have come not from a cloud provider's ability to service a customer, but the ability of the customer's systems to interface properly

with the cloud. Therefore any audit scope should include both the provider and any internal systems exposed to the cloud to ensure a complete "envelope" of integrity.

When considering the scope of any auditing protocol, you must step beyond contract terms and conditions and ensure that you are addressing general issues of management and governance, including necessary resources to mitigate any risks found. For example, it's 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. Consider carefully the importance of leveraging methods of audit and compliance that already exist in your organization, and look to extend those to the cloud vs. creating new ones.

Monitoring & Reporting

Transparency of the service level is extremely important to a successful service management protocol. While every cloud vendor offers different systems for visualizing data and its implications (web based, e-mail based, live, reactive, portal-based), customers should demand from any CSA a minimum set of capabilities:

1. *Cloud Performance Management*. This domain focuses on the response times for systems within the cloud architecture and between the cloud and the target user systems.
2. *Peak Load Performance*. This domain focuses on measurements and timings for when the cloud is under stress, either intentional or unintentional. As systems can perform differently when under different loads, and the interactions and dependencies of a complex cloud are often unknown in advance, it's important to visualize data both in a steady state as well as under load.
3. *Hybrid and Inter-cloud Performance*. As many clouds consist of different subsystems, often sourced from different cloud providers, it's critical to visualize data about the interactions between those hybrid cloud components.
4. *Application Performance*. This domain focuses on the applications executed from the cloud, particularly internal processing benchmarks as well as end-user experience measurement.
5. *Problem Notification*. This domain focuses on monitoring and reporting on failures and issues with the cloud system. Addressed are issues with prioritization, notification and severity level assessment.

Although the benchmarks in each of these areas are evolving, ensuring your CSA includes the ability to see, assess and react to measurements in these areas will help keep your cloud infrastructure running smoothly.

Measurement and Metering

A core characteristic of many cloud services is an on-demand model, where services used are billed as they are consumed, on a time or capacity basis. Therefore it is important to have confidence and transparency in the measurement and metering system employed by cloud providers, as embodied in the CSA you negotiate. At a minimum, you must ensure that metering systems employed by your cloud providers include:

1. Assurance of accurate billing, and a methodology for handing objections or challenges to any automated metered billing
2. The ability to segregate different services into different methods of billing: for example, performance testing, analytics, security scanning, backup, and virtual desktops might all be measured differently and metered separately.
3. Ability to handle taxation issues from geography to geography, and from user to user. As each country and municipality has implemented different approaches to taxation of online commerce, your provider must be able to discern between these sources of use and meter them independently.

Provisioning

While auditing, monitoring, measuring and metering relate primarily to the cost savings features of the cloud, provisioning is a key enabler of the improved flexibility that comes from the cloud. However, it's not without its own unique qualities that must translate into your CSA:

1. *Core provisioning speed*. As part of a CSA, there should be baseline expectations of the speed of deployment of new systems, new data, new users, new desktops or any function that's core to the service provided by the cloud vendor.
2. *Customization*. It's unusual that any templated method of rapid provisioning 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*. Important to any strong CSA are provisions 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*. It does no good to have a technical solution to rapid provisioning if the system is incapable of dynamic de-provisioning to match downturns in demand.

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 their own particular additional topics, particularly for different industries or IT applications running in the cloud.

Change Management

Change is an inevitable part of any IT system, and the cloud is no different. Fortunately, there is little that is special about the cloud in regards to considerations for change management. Procedures for requesting, reviewing, testing, and acceptance of changes differ little from those already in use with other IT subcontractor contracts and outsource agreements. The only unique issue is the sensitivity that many have to changes that have potentially radical implications, such as the cloud. In this case, extra care should be taken to manage the process carefully.

Upgrades & Patching

A subset of change management is upgrades or improvements in existing contracted services, such as when an upgrade or patch is needed, or when a new version of an underlying management system or

SaaS application is rolled out. In these cases, it's important to outline in your CSA a set of basic steps for these inevitable needs.

1. **Responsibility to develop requested changes.** There should be a clearly defined responsibility set for which party is in the lead for different types of upgrades. For example, if the upgrade is dependent on many subsystems or people internal to an organization, not in the cloud, it might be advisable to center the responsibilities on the contracting organization vs. the cloud provider. On the other hand, if the majority of the upgrade happens with cloud-provider personnel within the cloud space, it's likely the provider would assume primary responsibility.
2. **Process for identifying a timeline to develop, test and implement the change.** There must be a clearly defined "chain of command" and project plan for all changes made to the cloud environment, properly resourced and timed to ensure reasonable contingencies and problem resolution. Here too, little is different regarding a cloud solution vs. a traditional IT solution, with the exception of the increased anxiety and scrutiny that the cloud draws today. This is in many ways simply a special case extension of change management policies which should already be in place.
3. **Process for resolving problems resulting from change.** Since problems can often be compounded and result from multiple factors both within and outside the cloud, a CSA-based outline of upgrade procedures must include a clearly defined set of responsibilities and methods for resolving issues introduced by any upgrade.
4. **Back-out process if the changes cause major failures.** Even the best-laid plans often run aground on the rocks of reality. Cloud service providers should automatically embed rollback checkpoints throughout an upgrade plan in order to "pull the plug" and restore any upgrade to its initial state should an unexpected and unsolvable problem crop up during the upgrade procedure. Throughout the process, regular communication meetings should occur to keep both parties in sync.

Step 7: Prepare for Service Failure Management

Service failure management outlines what happens when the expected delivery of a cloud service does not occur. Cloud service capabilities and performance expectations should be explicitly documented in the CSA, as described in Step 4. It is important to note that the term "service failure" can cover a number of different things, from the complete unavailability of the cloud service, through response times that are longer than those promised in the SLA, through error responses to valid service requests made by the users. Service failure management covers activities both of the cloud service customer and also the cloud service provider.

Service failure management begins with the detection and alerting that a failure has occurred. The cloud service customer must ensure that cloud service failures can be detected. The cloud service provider may provide service monitoring capabilities to the cloud service customer and may in addition provide alerts to the customer when cloud service failures occur. However, the cloud service customer must establish whether the monitoring and alerting capabilities provided (if any) meet the customer's requirements. The cloud service customer may often need to put in place their own set of cloud service

monitoring and alerting capabilities to ensure that all the potential cloud service failures of significance to the customer are detected.

Once a cloud service failure is detected, then the cloud service customer must ensure that a management system is in place to alert appropriate customer staff, to report the failure to the cloud service provider (assuming that the failure was not already detected and reported by the provider) and to put into action any processes to mitigate the failure. For some cloud services and for some types of service failure, the cloud service customer may need to provide suitable evidence to the cloud service provider that a failure has occurred. The cloud service customer must track the progress of each reported failure and if the failure is not rectified within stated timescales, an escalation process must be followed.

The cloud service customer must understand the cloud service provider's service failure management procedures:

- The process for reporting failures detected by the customer
- The process which the provider will follow to address a reported failure
- The timescales for remedial action
- The process that the cloud service provider will follow subsequent to a failure to improve the provider's operations to avoid the failure occurring again

Planning for cloud service failures on the part of the cloud service customer will also often involve having a disaster recovery plan, which will be brought into action if the cloud service failure is likely to have a significant impact on the business. Disaster recovery planning is discussed in Step 8 of this paper.

Remedies

The primary remedy for service failure is service credits. These are typically based upon a percentage of the fees paid by the cloud service customer during the billing cycle. The actual percentage will vary depending on the cloud service provider and on the nature of the cloud service itself. However, it is common that these service credits will not exceed 100% of the paid fees. This can result in service credits not being in proportion to business cost or risk to the cloud service customer.

Limitations

Within each cloud service provider's service agreement there may be liability limitations for certain types of service interruptions. While these may vary dependent upon the provider, a sampling of several major providers shared the following exclusions:

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

In addition to common, shared limitations, there are cloud service providers who may also cite scheduled downtime as being excluded from the CSA metrics.

Roles / Responsibilities

The roles of cloud computing service failure management are described in the ISO/IEC 17789 Cloud Computing Reference Architecture [4]. The cloud service administrator has the responsibility to drive the incident management process and so needs to receive an alert when a service failure is detected. Assuming the service failure is impacting the customer use of the service, the cloud service administrator will engage the cloud service provider's incident management process, as described in the service agreement.

On the cloud service customer side, additional roles may be involved, including the help desk and the cloud service integrator. The help desk should be aware of the service failure and the likely impact and estimated time to resolution so as to be able to answer questions about the cloud service from cloud service users. The cloud service integrator would be engaged to triage the service failure and potentially propose solutions or workarounds to reduce the impact on the customer's business.

Monitoring and Notification processes

Monitoring of a service failure can be done in one of two ways.

1. The cloud service customer puts in place system(s) which monitor the customer use of the cloud service. The concept is that the customer does not rely on any capabilities of the cloud service provider and instead places instrumentation of some form in the customer-side components that use the cloud service. This might, for example, involve routing all customer requests to the cloud service through an instrumented component such as an Enterprise Service Bus (ESB). Requests made to the cloud service can then be monitored for success or failure, for their response times and for any other characteristics of importance to the customer. A set of rules can be put in place which will determine if there is a service failure and an alerting process invoked when a service failure is detected.

2. The cloud service provider has in place a cloud service monitoring system which has an interface enabling the cloud service customer to monitor the behavior of the cloud service and to receive alerts in the case of a service failure. These alerts should be integrated into the cloud service customer's alerting system. An alert would be sent to the cloud service customer when a service failure occurs – but the cloud service customer must understand what type of failures are notified through this process and it may well be the case that not all service failures of importance to the customer will be notified. Upon receiving a notification, the cloud service customer should follow their established service failure management process.

For a typical cloud service, it is likely that the cloud service customer will use both approaches to the monitoring of the cloud service. In some cases, the cloud service provider monitoring will be absent or will be inadequate for the customer. In other cases, there may be factors that can only be monitored by the customer, such as the effect of the internet connection to and from the cloud service.

For the notification of a service failure, in the ideal situation there should be a two-way automated interface between the cloud service customer and the cloud service provider that is used to transmit notifications of a service failure in both directions. This caters for either party becoming aware of the service failure. In the case where two way notification is not provided, the cloud service customer should expect there to be a facility for the customer to report a service failure to the cloud service provider – and a process for the customer to track what is happening in regard to each reported service failure.

Step 8: Understand the Disaster Recovery Plan

Disaster recovery is a subset of business continuity and focuses on processes and technology for resumption of applications, data, hardware, communications (such as networking), and other IT infrastructure in case of a disaster. By the term disaster we mean either natural disaster or man-made events that have an impact of availability of IT infrastructure or software systems.

It is common to see a false sense of security among cloud customers regarding disaster recovery planning. Just because businesses are outsourcing the infrastructure (IaaS), applications (SaaS), or platforms (PaaS) to cloud service providers does not absolve them of the need for serious disaster planning. Every company is unique in the importance it assigns to specific infrastructure/ applications, thus, a cloud disaster recovery plan is specific to each organization, and business objectives should play an important role in determining the specificity of disaster recovery planning.

The process of devising a disaster recovery plan starts with identifying and prioritizing applications, services and data, and determining for each one the amount of downtime that's acceptable before there is a significant business impact. Service priority, required recovery time objectives (RTOs), and recovery point objectives (RPO's) will determine the overall disaster recovery approach. For example, in some applications maintaining uptime may be more important than having the data precisely replicated as of the last time of failure. Further, while 99%+ uptime SLAs are common in cloud computing (approximately 4 days of down time a year), it may not be adequate for specific application and business needs.

In general, current CSAs provide inadequate guarantees in case of a service outage due to a disaster. Most CSAs provide cursory treatment of disaster recovery issues, procedures and processes. That being said, it is rare for SMBs to internally develop the extensive disaster recovery infrastructure of large and established cloud providers.

Despite the limitations in CSAs, cloud adopters should address key disaster recovery questions/issues with their service providers early in the process of cloud adoption. The key areas to address with cloud providers are:

- How is service outage defined?
- What level of redundancy is in place to minimize outages including co-location of services in different geographical regions?
- Will there be a need for scheduled down time?

- Who has the burden of proof to report outages? This can be difficult to prove in case of conflicts with the cloud providers.
- What is the process that will be followed to resolve unplanned incidents?
- How will unplanned incidents be prevented or reduced?
- When does the time clock start on lack of service availability in order to measure service credits?
- How will incidents be documented or logged?
- 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 is the problem escalation process?
- Who are the key service provider and customer contacts (name, phone number, email address)?
- What is the contingency plan during a natural disaster?
- How is the customer compensated for an outage? It must be noted that cloud providers have limits on the maximum compensation provided in case of an outage, and the compensation is an insignificant remedy in case of serious outage.
- Does the cloud vendor provide cloud insurance to mitigate user losses in case of failure? Although this is a new concept, some major cloud vendors are already working with insurance providers.

Answers to the questions above will be highly specific to particular organizations, and their specific disaster recovery needs. For large enterprises the questions mentioned above can be used as a framework to seek a stronger disaster recovery component in a negotiated CSA. It is important to emphasize that this is only possible for large enterprises with large contracts. Established cloud vendors are quite resistant to altering existing CSAs.

There are large numbers of events that can have negative impact on the availability of cloud services provisioned by customers. Although, detailing all of them is out of the scope of this section, some of the important areas that cloud customers should consider are in areas of security/intrusion detection, denial of service, viability of a cloud provider, data ownership and recovery. As an example to highlight the above, consider a company using SaaS for critical applications, such as order management, billing, or ERP. The cloud user will face major technological hurdles in shifting to another provider in case of a disaster like a financial failure of the cloud provider. Cloud users should make it a priority to address key contingencies in case of such an event. Issues such as access to data and the application in a timely manner are critical to clarify.

While, in most cases, companies will be able to retrieve the application data from an established SaaS provider, the business logic and software systems will be left behind. One solution is to deploy the SaaS software onsite and run it internally – clearly a difficult and risky solution to implement. So, despite good planning, in some cases no easy solutions are available for negative events. Development of data and meta-data standards in specific application domains could provide a considerable benefit for customers and allow them to migrate to different SaaS solutions in the event of a disaster. The

development of such standards though is in direct conflict with the interests of many providers, and will take time to materialize.

It is also important to understand that risk mitigation related to disaster recovery for cloud solutions will also depend upon the specific cloud type (IaaS, SaaS etc). Compared to the SaaS example above, in the case of a negative event for an application running on an IaaS, the client can implement a different set of solutions. One example solution would be to architect the application to continue performing in the face of individual resource failure (e.g., server failure, storage failure, network failure, etc), or in the case of a significant infrastructure failure use hot/warm sites in a different geographical zone or on a completely different cloud. The key point to understand is that risks and solutions associated with negative events will be different for SaaS, IaaS and PaaS.

When it comes to disaster recovery the public cloud presents a due-diligence paradox. While there are myriad options for implementing disaster recovery, and the cloud may simplify enterprise IT by abstracting away a lot of the complexity, it also increases the difficulty of performing comprehensive due diligence including testing of disaster recovery procedures. Lack of such diligence accompanied by weak CSAs represents a potential risk in the area of business continuity and disaster recovery. Thus, companies should view developing and testing a disaster recovery plan as an important part of moving to the cloud. Companies can consider using business continuity/disaster recovery standards as part of their planning efforts. Existing standards such as BS 25999:2007, NFPA 1600:2010, NIST SP 800-34, ASIS SPC.1-2009, ISO 27031, and ISO 24762 can provide an effective starting point for planning disaster recovery.

Step 9: Develop an Effective Governance Process

The use of cloud services by a cloud service customer means that the customer organization is placing some parts of its IT operations – and hence part of it business processes - in the hands of outside suppliers in the form of one or more cloud service providers. As a result of the interface(s) between the customer and the provider, there is a need for strong and detailed governance of the use of the cloud services on the customer side.

The first part of the governance process involves the control and oversight of the previous steps outlined in this practical guide, which provide the necessary underpinnings for the selection and use of cloud services. The second part of the governance process is the regular ongoing review of the use of each cloud service, to ensure that it meets business requirements and to ensure both internal and external user satisfaction with the cloud services and the applications built on them. The governance process should also deal both with changing business and user requirements and also with any changes to the cloud service(s) that may be made by the cloud service provider.

Table 8 below highlights the key elements required to operate a successful governance process.

Table 8: Governance Process

Element	Description

Periodic assessment of achieved cloud service levels against agreed CSA	• Reports from cloud service provider of cloud service levels • Monitoring reports on cloud service usage created by customer cloud service administrators
Periodic assessment of compliance of cloud service	• 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.
Service failure reports	• Reports of any service failures or incidents which affect o Service availability o Security, particularly security breaches o Protection of personal data
Notification of changes from the cloud service provider	Any change notifications from the cloud service provider which 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)
Key indicator reports	Four key indicators should be tracked to ensure that the CSA criteria are being met and that the downstream users of the service (either internal or external to the enterprise) are experiencing the level of service that has been agreed to: • High impact problems and time to resolution • Number of open problems and their respective impact • Total view of problems not resolved within agreed to time frames • Trends of number of problems being reported with the resulting resolutions
Problem reports	In order to ensure CSA compliance, a set of reports needs to be produced: • Reports that focus on the current reporting period addressing: o All problems reported (sorted by impact) o Problems closed (sorted by impact) o Duration of open problems (sorted by impact)

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 • Time to action requests
User satisfaction reports	Reports on user satisfaction with the cloud service(s)

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. How the review is performed is a decision for the customer and it is likely to depend on the size and structure of the customer organization. A degree of formality and record keeping is advisable since in some cases, evidence may need to be prepared for presenting to the cloud service provider, especially if there are matters under dispute between the customer and the cloud service provider.

What constitutes an appropriate course of action will depend on the nature of the issue(s). Some breaches of the CSA terms may trigger remedy terms which imply some level of compensation to the customer – but it may often be the case that the customer must formally raise a request to the provider in order to trigger the terms of the remedy. More serious breaches or incidents are likely to require more significant action on the part of the customer. This may take the form of discussions between senior management from the customer with their counterparts from the cloud service provider. Alternatively, it may take the form of the customer deciding to switch their use of cloud services to another cloud service provider, triggering the termination process.

For problems that require higher management awareness, it is the responsibility of those involved in the governance process to advise their respective management chains on the status of a particular issue.

Escalation Process

Inevitably, there will be problems which fall outside the normal management process and will need additional focus to ensure a timely resolution. An example of the exceptional process is a major outage, i.e. loss of service, which cannot wait for a periodic meeting and requires an immediate notification of the management chain.

While we use the term escalation, the escalation process is really upward communication for awareness for a particular situation and not an upward delegation of responsibility for the resolution of the problem.

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

Table 9: Escalation Considerations

Consideration	Description
Objectives	• Raise management awareness to avoid surprises (gives the perception that senior management is in control of the situation). • Gain agreement for action plans to resolve a problem. • Develop either a plan and gain agreement for additional resources, when required.
Guidelines	• Problem has a critical impact to the overall business to either an internal service or a customer facing service. • Service is still available but is significantly degraded; potential impact to a customer facing service. • Problem is of a significant impact and has missed the agreed to targets for resolution. • Independent of impact, problems are not being closed within the expected guidelines. • Number of problems is increasing with no agreed to resolution to reverse the trend. • Requests to the cloud provider to participate in root cause analysis or problem resolution in an associated system or tool are ignored.
Types	• **Immediate** o A critical business impact is identified. o Significant impact to a customer facing service. • **As required**. Typically after a review when: o The duration of problem resolution is not being met. o Number of open problems exceeds expectations. o Trend for reported problems is increasing without a satisfactory resolution plan being offered.

Once an escalation has been initiated, the goal is to ensure that both chains of management understand the problem, its impact, and the currently agreed to action plan for resolution including containment of the problem, especially if the problem impacts an external customer service.

If a resolution of an escalated problem cannot be reached through the escalation process then the terms of the CSA can be brought to bear to force resolution. One of the outcomes of continuous breaches to the CSA can be termination of the agreement with the provider for the contracted service(s). It should

be noted that the minutes generated from the management process is an important set of documentation to support the termination process.

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 a tool to manage the services and ultimately provide the best services to the users of the service(s), whether the users are internal or external to the organization.

Step 10: Understand the Exit Process

An exit clause should be part of every CSA and describes the details of the exit process including the responsibilities of the cloud provider and consumer in case the relationship terminates prematurely or otherwise.

There are numerous potential scenarios that could cause the termination of service between customer and provider which would result in the execution of the exit process. For example, a provider may be unable to deliver the required levels of performance and availability specified in the SLA, or it may be the case that the provider is going out of business. Regardless of the reason, a clearly defined exit process that ensures secure and speedy transfer of customer data and applications is essential.

A customer exit plan should always be prepared at the outset of the CSA and is an integral contractual annex. This plan should ensure minimal business disruption for the customer and ensure a smooth transition. The exit process should include detailed procedures for ensuring business continuity and it should specify measurable metrics to ensure the cloud provider is effectively implementing these procedures.

The most important aspect of any exit plan is the transmission and preservation of cloud service customer data, which is critical to achieving business continuity. In addition, customers must ensure that their data is completely removed from the provider's environment once the exit process is complete. Customers should look out for and be aware of the following details when they evaluate the exit clause included in a CSA.

- The level of provider assistance in the exit process and any associated fees should be clear in the CSA. In most cases, there should be no additional cost associated with the exit process.
- Providers should be responsible for removing customer data from their IT environments, or at least aid the customer in extracting and erasing their data by providing clear and concise documentation.
- The format of the data transmitted from the provider to the customer should be specified in the CSA and should leverage standard data formats whenever possible to ease and enhance portability.
- The CSA should specify that all data and information belonging to the customer is maintained for a specific time period after transition and then be 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 and/or destroyed before that time.
- Customers should ensure that the CSA provides appropriate business continuity protection during the exit process.
- At the completion of the exit process, customer should receive written confirmation from the provider that all of the customer data has been completely removed from the provider's IT environment. The written confirmation should also state that the provider agrees not to use the customer data for any reason in the future, including using the data for statistical purposes.

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 keys to success for any customer organization evaluating and comparing CSAs from different 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 a move to cloud services. • Purchasing and reporting are key areas for review.
Develop a strong business case and strategy for cloud computing environment	• Assess criticalness of services being deployed in the cloud. • Determine functional and non-functional requirements for each service (performance, availability, security, privacy, etc.). • Understand legal and regulatory requirements concerning 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 criticalness of the service being deployed in the cloud, determine if the cloud provider's CSA is sufficient to address the functional, non-functional, legal, and regulatory requirements of the service. • If not, determine if the cloud provider is willing to negotiate on the key aspects of the CSA that are not in line with your business strategy. • If the cloud provider is not willing to negotiate on these critical points, seek alternative providers who more closely address your requirements. • If a cloud provider who addresses your requirements cannot be found, strongly consider keeping the service within your enterprise IT environment.
Determine how to monitor CSA performance	• Assuming a cloud provider is found that meets your service requirements; understand the management process defined in the CSA. • Ensure your CSA includes the ability to see, assess and react to key performance measurements that will help keep your cloud infrastructure running smoothly. • Understand the notification process when service issues arise including method and timliness of notifications along with prioritization and severity level assessment of issues. • Be aware of remedies and liability limitations offered by the cloud provider when service issues arise.
Ensure an adequate disaster recovery plan can be defined and executed	• The cloud customer bears the risk of disaster scenarios that severely limit the ability of their cloud provider to deliver service. • Cloud customers must understand the provider's ability to support their data preservation strategy which includes criticalness of data, data sources, scheduling, backup, restore, integrity checks, etc. • Roles and responsibilities must be clearly documented in the CSA. In many cases, the cloud customer may be responsible for implementing most of the data preservation strategy. • Based on the criticalness of the data, cloud customers should clearly define recovery time objectives. • 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 should the relationship with the cloud provide terminate prematurely.
	• The exit plan should be taken into account during the assessment phase of potential cloud providers.
	• The provider's CSA should be carefully reviewed to ensure the customer defined exit plan is capable of being implemented.
	  o The customer should be able to terminate the agreement at any time, without penalty, provided sufficient notice is given to the provider.
	  o Data maintained on the provider's cloud resources should be stored using standard formats to ensure data portability.
	  o Transmission of data from the provider's cloud resources should leverage standard packaging and data transfer techniques.
	• Roles and responsibilities must be clearly documented in the CSA. In many cases, the cloud customer may be responsible for initiating most of the exit process steps.

In addition, 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.

Cloud computing offers a value proposition that is different from traditional enterprise IT environments. With proper focus on the key success factors, customers are able to effectively review and compare CSAs from different cloud providers to ensure the promise of the cloud is realized.

Works Cited

[1] Cloud Standards Customer Council (2013). *Public Cloud Service Agreements: What to Expect and What to Negotiate*
http://cloud-council.org/resource-hub.htm#migrating-applications-to-public-cloud-services

[2] Cloud Standards Customer Council (2014). *Practical Guide to Cloud Computing, Version 2.0.*
http://cloud-council.org/resource-hub.htm#practical-guide-cloud-computing-v2

[3] Cloud Standards Customer Council (2015). *Security for Cloud Computing: 10 Steps to Ensure Success, Version 2.0.*
http://cloud-council.org/resource-hub.htm#security-for-cloud-computing-10-steps-to-ensure-success

[4] ISO/IEC 17789: Cloud Computing Reference Architecture
http://www.iso.org/iso/catalogue_detail?csnumber=60545

Additional References

Cloud Security Alliance. *Security Guidance for Critical Areas of Focus in Cloud Computing Version 3.0 (2011).* http://www.cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf

This document provides an actionable, practical road map to managers wanting to adopt the cloud paradigm safely and securely.