Analysis of Public Cloud Service Level Agreements – An Evaluation of Leading Software as a Service Providers

Michael Seifert[1]

[1]Martin Luther University Halle-Wittenberg, Chair for Information Management, Universitaetsring 3, 06108 Halle (Saale), Germany
michael.seifert@wiwi.uni-halle.de

Abstract. Public cloud and software as a service (SaaS) are two of the largest growing IT markets in recent years. Cloud customers need to assess whether the predefined service level agreements (SLAs) of public cloud providers are suitable for their business requirements. Due to the lack of a standard SLA formulation, cloud consumers have significant effort in analyzing SLAs against their compliance, which could be supported by semi-automated SLA management.
SLAs of five leading SaaS providers with comparable public cloud business applications were examined as an as-is analysis. Using 18 derived parameters, the SLAs were formalized and evaluated in terms of matchmaking. The percentage of formalization and matchmaking among the five providers was found to vary between 20% and 73,3% across four SLA categories. Several contributions could be made for practitioners, but also for researchers on how to address the high degree of heterogeneity in public cloud SaaS SLAs.

Keywords: public cloud, software as a service, service level agreements

1 Introduction

The entire cloud market has been increasing continuously for years [1, 2]. Especially the market of public cloud [1] as well as of software as a service (SaaS) [1, 2] is growing significantly. An increasing number of companies are deciding to consume their business applications from the cloud instead of providing them by themselves [3]. At the same time, it enables software vendors to provide their solutions to a wide range of customers [4]. SaaS adoption is receiving increasing attention in practice [5]. The possibility of fast implementation and a higher innovation cycle makes SaaS attractive for businesses [4]. The billing model – from capital expenditure to operational expenditure – is also a valid argument for adopting SaaS in comparison to traditional application service consumption [6].
But cloud providers are also affected by typical IT challenges. For example, cloud providers may require downtime to perform maintenance on their IT infrastructure [7,

16th International Conference on Wirtschaftsinformatik,
March 2021, Essen, Germany

8]. At times, even large cloud providers, and therefore cloud users, experience unplanned downtime [9]. These planned and unplanned downtimes are usually defined and described by cloud providers. This information is agreed and documented with the customer in so-called service level agreements (SLA) [10]. Many cloud providers (usually public cloud providers) even publish their SLAs before signing a contract, which makes it possible for potential customers to analyze them in advance [7, 8].
Accordingly, as a potential cloud customer, an upcoming decision to adopt cloud services should always be based on the customer's own business criticality (e.g. for possible unavailability of the service) to assess risk and service level compliance [11, 12].
The main challenge here is the lack of cloud SLA standards. For potential cloud customers, this means to evaluate the SLA individually against their own business requirements. In addition, certain information that one provider describes in its public SLA may be documented differently, or not at all, by another provider [7].
When evaluating the SLAs of cloud providers on the customer side, it should also be noted that new cloud services often have to be integrated or composed with existing IT services (e.g. for master data exchange) [13, 14]. This means that the cloud customer must not only evaluate the components of the SLA for themselves but aggregate them with SLA parameters of existing IT services to evaluate whether the composition of services continues to meet their business requirements or at least does so at acceptable risk [12, 13].
In research, established models and methods are already proposed for the two scenarios, (I) cloud service selection [15, 16] and (II) cloud service composition [17, 18]. There are also numerous ontologies and meta-models published for standardization and semi-automated SLA-aware selection and composition of cloud services [19, 20]. To enable evaluation and enhancement of models and methods in research, as well as to provide an overview to cloud customers and providers, the state of current cloud SLAs is identified. This study was conducted with the following research questions (RQ), as an as-is analysis of present-day public cloud SaaS SLA.

- RQ1: How can public cloud SaaS SLAs be formalized and categorized in a consistent way?
- RQ2: How much can the formalization of SLA components be used to compare or aggregate (named matchmaking) content from different providers?
- RQ3: What can be derived for research and practice from the results of this study?

To answer these research questions, the next section introduces the fundamental cloud terminology and necessary concepts as a theoretical background. Next, the definition of the study scope is provided by presenting the choice of the study sample and the criteria for analysis. In addition, related work is presented in section 3 and compared with the study scope at hand. Section 4 outlines the data collection of five leading public cloud SaaS provider and their SLAs to make the research comprehensible. Furthermore, the collected data is formalized and categorized here according to RQ1 in context of a moderated focus group as a qualitative research method [21]. In section 5, the five cloud provider SLAs are instantiated according to the formalization. The parameters are then analyzed in terms of their matchmaking to provide an answer to RQ2. The evaluation and discussion of the analysis in section 6 is followed by a consideration of threats to validity of this research in section 7. The article ends with the conclusion in which the contributions to research and practice of this paper are summarized.

2 Theoretical Background

The study is grounded on cloud terminology following the National Institute of Standards and Technology (NIST). Three cloud service models and four cloud deployment models can be distinguished [22].

The service models are differentiated into Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS) [22]. IaaS describes a cloud service in which the provider delivers a complete IT infrastructure ready to use to the consumer [22]. Thereby, computational as well as network and storage resources are composed. With PaaS, further services on top of the composed IT infrastructure are delivered to the customer, enabling application development, for example [22, 23]. With PaaS, customers get the opportunity to develop their own applications in the cloud. With SaaS, usable software is provided to the cloud consumer on top of IT infrastructure [13, 22]. SaaS is usually used by organizations when the cloud application already meets the functional business requirements or when in-house operation of the application is not preferred.

The cloud deployment models are divided into Private Cloud, Public Cloud, Community Cloud, and Hybrid Cloud [22]. Private clouds are services that are deployed by the provider for a specific consumer organization [24]. The private cloud provider is usually in close interaction with the customer in order to consider their business requirements. In contrast, the public cloud is about the provider making the cloud service available to many users and in general [22, 24]. Due to the identical composition, the cloud provider can deliver its service to a large number of customers at the same time. Community cloud is similar to private cloud, but is in contrast provided to be consumed by multiple organizations with similar concerns [22, 24]. Community cloud is used, for example, when several universities consume the same cloud service, but want to have their respective customizing considered. Hybrid cloud is defined as the combination of at least two cloud deployments [22, 24]. The most common type of hybrid cloud is the combination of public cloud and private cloud. The increasingly popular hybrid cloud driven by public cloud SaaS adoption has significant impact on IT management [25].

In order to ensure the contractual relationship between the cloud provider and the customer, service level agreements are signed. A Service Level Agreement (SLA) is a contract for an agreed IT service between a provider and a consumer [10]. The details of the SLA must be underpinned by measurable parameters before and during the service lifecycle in order to be comprehensible for the provider and the customer. To measure and evaluate agreed performance levels of cloud services, qualities of service (QoS) are commonly used [26].

3 Study Scope and Related Work

The study presented in this paper has a two-sided target audience, (I) research and (II) practice. For researchers, the study aims to provide new practical insights for further adaptation and evaluation of existing SLA management models and concepts (e.g. smart contracts), as well as for cloud selection and composition methods and artifacts (e.g. QoS aggregation). For practitioners, the study is of interest because it provides the cloud consumer with an overview of what aspects of public SLA they can align their business needs with. For cloud providers, the analysis serves as a guide to what other providers present in their SLAs and how the survey sample focuses on various SLA parameters and categories.

For the analysis of the SLAs, two evaluation criteria formalization and matchmaking are applied. Formalization is understood as the distillation of the described semantic in the SLA as comparable parameters, as in [7] and [8]. Formalization is therefore where (I) the aspects are included in the respective SLA of the providers and (II) can be assigned to the respective parameters defined.

In order to support the SLA management, we use matchmaking as a second evaluation criteria. Matchmaking is known as a method in QoS-compliant selection of Web services [19]. As an approach to examine constraint satisfaction problems, metrics are checked for semantic and unit-specific equivalence [27]. The SLA parameters are checked by matchmaking to ensure that they are operable among providers, i.e. (I) comparable in terms of cloud service selection or (II) aggregable in terms of cloud service composition.

The study is conducted with focus on one cloud model, namely SaaS. The decision was made because it is expected that IT departments will increasingly need to evaluate SLA compliance in the context of business requirements based on functional or strategic preferences of a specific cloud application [5].

It was also decided to focus on public cloud as the deployment model of the study. Both the decision for the cloud model and the cloud deployment of the study are supported by the high market relevance [1, 2].

Another decision regarding the scope of the study is the focus on business applications (compared to cloud applications for private use). The reason for this is that the commercial risk of insufficient service levels is significantly higher in the business context. In order to ensure the best possible comparability of the SLAs of different providers, business application cloud services were analyzed that are not industry specific.

The scope of the study aims to achieve the highest possible generic coverage, while at the same time ensuring the highest possible transparency of the selection, in order to be able to use the results of the study as broadly and specifically as possible. In the context of the presented study scope, two related work studies are presented in Table 1.

Table 1. related work

	Baset (2012) [7]	Guila and Sood (2013) [8]
title of publication	Cloud SLAs: Present and Future	Comparative Analysis of Present Day Clouds using Service Level Agreements
cloud models	IaaS, PaaS	IaaS, PaaS, SaaS
cloud providers	Amazon, Azure, Rackspace, Terremark, Storm	Rackspace, Engine Yard, Google
SLA parameters	service guarantee, service maintenance, service credit, service violation measurement & reporting	service commitment, definition, credit request/claim, service credit, SLA exclusions

The articles of Baset [7], and Guila and Sood [8] are from the years 2012 and 2013. Due to the passing time in between, it can be assumed that there are changes in the common public cloud SLA. Accordingly, our investigation provides a refresh regarding current cloud SLAs.

In contrast to our fixed scope on SaaS, at least two different cloud models are considered in each of the studies. Both studies also examined public cloud SLAs, so publicly available SLAs served as the foundation.

One issue of criticism in both studies is the comprehensibility of the vendor selection. Neither article explains how the selection is made for each cloud model considered. The article at hand will therefore describe the selection of vendors to be analyzed in section 4 based on the maximum possible generalizability of our results.

Last, this study differs from related work in the depth of analysis. With the motivation of semi-automated processing of the contents of SLA, the capability of matchmaking for the parameters is examined in section 5. The studies by Baset [7], and Guila and Sood [8] each stop at the formalization of the SLA aspects, and thus do not examine subsequent machine processing.

4 Data Collection and SLA Formalization

The data collection starts with a search for market study on the valuation of cloud computing. After screening the two studies on the cloud computing market of Gartner Incorporated (Gartner) [1] and Synergy Research Group (SRG) [2] the leading vendors of SaaS were selected. The selection of our sample for public cloud SaaS business applications goes back to the breakdown of the "Worldwide Market Share of Enterprise SaaS" by SRG [2] and is shown in Table 2.

The cloud services depicted are all public cloud SaaS and, to ensure comprehensibility, non-industry-specific IT applications for business context. Based on the top five enterprise SaaS providers, administrative business applications were selected that could potentially be used in a variety of organizations. Content management systems (CMS) as well as customer relationship management (CRM) and enterprise resource planning (ERP) systems are used in almost all organizations and thus provide a suitable basis for an analysis.

Table 2. selected cloud providers and applications

vendor	application product	application type	links
Adobe	Adobe Experience Manager	Content Management System (CMS)	[31-32]
Microsoft	Microsoft Dynamics 365 Business Central	Enterprise Resource Planning (ERP)	[30]
Oracle	Oracle Fusion Enterprise Resource Planning Cloud	Enterprise Resource Planning (ERP)	[29]
Salesforce	Salesforce Customer 360	Customer Relationship Management (CRM)	[27-28]
SAP	SAP S/4HANA Public Cloud	Enterprise Resource Planning (ERP)	[26]

The data collection process also required further assumptions. First, the formalization of the SLA was performed in each case with reference to corresponding productive system of the cloud service. This represents the aspiration to reflect the business risk in the case of downtimes of the economically relevant systems. Second, in order to formalize certain SLA components, a specific region in which the system is hosted had to be assumed. For this we assumed to be from Germany and chose a region as close to Germany as possible.

The formalization was performed with the following sequence and in context of a moderated focus group [21]. The focus group consisted of 6 researchers and 4 practitioners, each of whom was included in the discussion at both stages of formalization (phase one and phase two).

In phase one, the SLA documents [28–34] were reviewed completely for each of the five providers in sequence and recorded in tabular form for each SLA-relevant parameter. The naming of the tabular documentation of the parameters was inspired by the respective SLA documents and the naming of the parameters of the related work (shown in Table 1). Once an aspect was identified in a subsequent SLA that did not semantically fit into existing parameters, a new parameter was created. Accordingly, the dataset of formalized parameters in Table 3 represents a union of the aspects of the five SLAs. Even if this means that not every parameter can be instantiated or mapped for each of the five SLAs of a provider. Instead, this satisfies the objective of an overview of possible aspects of a present-day public cloud SaaS SLA.

In phase two, after going through all the SLA documents, minor adjustments were made to improve the understandability and comprehensibility of the parameters and categories. For example, the distinction between maintenance and major release upgrades was formulated consistently according to their three relevant aspects. Potential shortcomings in the generation of the formalization are discussed in section 6 with respect to the validity of this study.

As the result, a formalization of SLA corresponding to four categories, each with associated two to six parameters (18 parameters in total), has been generated which is shown in Table 3.

Table 3. formalized SLA parameters and categories

category	no.	parameter	metric, description
service commitment	1.1	target service uptime	percent of minutes per month
	1.2	downtime	definition of downtime
	1.3	exclusion	definition of exclusion from downtime calculation
	1.4	service timetable	time when the service is available
	1.5	recovery time objective (RTO)	maximum time (in hours) between decision to active recovery process and the point at which you may resume operations
	1.6	recovery point objective (RPO)	maximun period (in hours) of data loss from the time the first transaction is lost
service maintenance	2.1	region	where the service is hosted
	2.2	system maintenance announcement	time of announcement of maintenance
	2.3	system maintenance date	maintenance starting time (time zone)
	2.4	system maintenance duration	maximum duration in hours for the maintenance
	2.5	major/release upgrades announcement	time of announcement of the upgrade
	2.6	major/release upgrades date	upgrade starting time (time zone)
	2.7	major/release upgrades duration	maximum duration in hours for the upgrade
service credit	3.1	credit calculation	service credit in relation to monthly payment
	3.2	credit notification	time to report a violation to the provider
	3.3	maximum credit volume	maximum service credit to be paid per month (as a percentage of the monthly fee)
service contract	4.1	termination clause	condition for exceptional termination of the order
	4.2	end of life	notification before the service is no longer generally available (in month)

The first category, *service commitment* (cf. Table 1, Guila and Sood), bundles issues around general availability and possible recovery from failures. *Target service uptime* (1.1) indicates the percentage of minutes the system is available per month. *Downtime* (1.2) specifies what is considered unavailable in terms of billing and *exclusion* (1.3) describes when the provider is exempt from the responsibility of promised uptime. *Service timetable* (1.4) describes the time when the system is up and running, even if no one is working on it. This is relevant, for example, when the system performs scheduled job processing during the night. Last, *recovery time objective (RTO)* (1.5) and *recovery point objective (RPO)* (1.6) are common metrics for the time needed to recover (RTO) and time of maximum data loss (RPO).

Service maintenance (cf. Table 1, Baset) covers the aspects in which the service is planned to be unavailable. This may happen for various reasons. On the one hand, due to necessary *system maintenance* (2.2 - 2.4) on underlying infrastructures or due to *major/release upgrades* of the software to a new release (2.5 - 2.7). The specified parameters are therefore identical for maintenance and upgrades. The *announcement* (2.2, 2.5) indicates the time before the unavailability of the application is announced. The *date* (2.3, 2.6) determines at which time (day and time), according to the stated time zone, the downtime usually takes place. The *duration* (2.4, 2.7) indicates how long the downtime typically lasts from the start time date. The *region* parameter (2.1) is used to specify the location where the service is hosted. This generally has an impact on the scheduled downtimes.

The *service credit* category (cf. Table 1, Guila and Sood, Baset) combines all financial aspects that are relevant once the service fails to fulfill the agreement and the customer receives a fee back. Service *credit calculation* (3.1) indicates how the billing is calculated in relation to the monthly fee for the cloud service. Whereby *maximum credit volume* (3.3) represents the maximum of it. The *credit notification* (3.2) determines the period of time in which the customer is supposed to submit the claim to the provider in order to receive the service credit.

The category *service contract* contains the potential termination of the contract for both parties. The *termination clause* (4.1) defines the number of service level violations after which the customer may terminate the contract for cause. *End of life* (4.2) specifies the period of time in advance for the provider to terminate the cloud service.

5 Evaluation and Discussion

The parameters have been instantiated via the SLA documents of the five providers (see Table 4). The instantiation of the formalization and the capability to be comparable and aggregable (matchmaking) is assessed for all parameters and evaluated across the categories.

With target service uptime it is noticeable that one provider (Salesforce) is not matchable (adjective of matchmaking) because it does not specify a quantitative availability. Downtimes can only be formalized for three vendors and are not matchable there due to the complex wording. The exclusion of availability-reducing factors can be formalized for all vendors. However, the fine details in the SLA are phrased in such a linguistic way that they are not matchable. The service timetable is explicitly described by three out of five providers. However, based on the description of all other parameters, we assume that all providers offer 24/7 service. The formalization of service timetable can therefore be questioned. The formalization of just one provider with regard to the practically highly relevant RTO and RPO documentation is highlighted as potential for improvement in cloud provider practice.

Many cloud providers set planned downtimes depending on the usual business hours per region. For example, these downtimes are usually scheduled for weekends. However, if the cloud service of your choice is not offered in your region, this can lead to scheduled downtimes during the week. Maintenance announcement is only defined for two of the providers. The announcement for upgrades, on the other hand, can be formalized for three providers, but due to vague descriptions it is only matchable for one provider. In this context, matchable means that during the service lifecycle it is possible to check automatically whether maintenance is scheduled by considering the minimum number of days prior notification (as automatically check interval). Maintenance date can be formalized for system maintenance and upgrades in five out of ten potential parameters. All parameters are matchable and give cloud consumers, in combination with maintenance duration, the possibility to compare the potential maintenance windows (which times disrupt their business less) and to aggregate (which maintenance days cover all components of their composite service).

Service credit is almost entirely formalizable across all parameters of four providers. Even the credit calculation of the four different providers is very similar which makes it easily matchable. However, it is noticeable that the maximum credit volume varies significantly (between 10% to 100%). The relevance of this category is considered to be quite high because, in a best-case scenario, the service credit should be able to compensate for the loss caused by business interruption as risk transfer.

Table 4. instantiated formalization and matchmaking of the cloud provider sample

cat.	no.	Adobe	Oracle	Microsoft	Salesforce	SAP
service commitment	1.1	99,9	99,9	99,9	commercially reasonable efforts to make the services available 24/7	99,5
	1.2	service not available to the customer, except any excluded minutes	-	users unable to login (excluded planned downtimes)	-	minutes the system is not available (excluded downtimes)
	1.3	misbehavior of customer	scheduled downtimes from my oracle support	planned downtimes; list of limitations (e.g. networt, inappropriate usage)	planned downtime; circumstances beyond reasonable control	regular maintenance, major upgrades; out of provider control
	1.4	24/7	-	-	24/7	24/7
	1.5	-	12	-	-	-
	1.6	-	1	-	-	-
service maintenance	2.1	America	-	EMEA	EU	Europe
	2.2	-	-	notified at least five business days in advance	ten days prior to the maintenance	-
	2.3	-	-	22:00 (UTC)	SAT, 22:00 (UTC)	SAT, 22:00 (UTC)
	2.4	-	-	8	4	4
	2.5	-	-	choose a specific weekend	approximately one year before the release date	notified at least five business days in advance
	2.6	-	-	-	FRI, 22:00 (UTC)	SAT, 4:00 (UTC)
	2.7	-	-	3 hours	6 hours	24 hours
service credit	3.1	<99,9% -> 5% fee, <99,5% -> 10% fee, <95% -> 15% fee, <90% -> 25% fee	per 1% below availability (99,9) you get 2% credit of your monthly fee; service credit is paid with the second month of missed service availability in a six month period	<99,9% -> 25% credit, <99% -> 50% credit, <95% -> 100% credit	-	per 1% below availability (99,5) you get 2% credit of your monthly fee
	3.2	-	30	within two months of the end of the billing month in which the incident occurred	-	30
	3.3	25	10	100	-	100
service contract	4.1	-	availability violation for three consecutive months	-	-	-
	4.2	-	12	-	-	-

formalizable	matchable

Service contract can be formalized by one provider and is matchable with parameters of other providers. Again, relevant aspects of the SLA are mapped, which could be used for enrichment in the SLA of other providers.

In order to get an overview of the results of the analysis, the assigned labels in Table 4 (formalizable, matchable) were evaluated quantitatively. The result of this analysis can be seen in Table 5 and are discussed in the following.

Service commitment is often not trivial to process by machine due to the lack of formalizability, which is seen as a challenge and a risk for cloud consumers. In addition, even if it can be formalized, it is often difficult to compare or aggregate it due to over-defined terminology and exclusion (see 1.2, 1.3 Table 4).

Service maintenance formalizability is basically enabled by three out of five providers. The formalized parameters allow a suitable processing in general. It remains to be said that both practice and research in the discovery or decision phase must nevertheless reckon with uncertainty in the run-up to the announcement or even the lack of announcement of maintenance. Maintenance must therefore be formalizable, measurable and adaptable in later phases of the service level lifecycle.

Table 5. evaluation of formalization and matchmaking of the cloud provider sample

category	formalization	matchmaking
service commitment	60,0%	30,0%
service maintenance	57,1%	40,0%
service credit	73,3%	73,3%
service contract	20,0%	20,0%

Service credit formalizability and matchmaking has the highest percentage of all categories. This shows that the calculation methods are similar across four depicting providers and are therefore easy to process. Nevertheless, it remains interesting for practice and research to include the maximum loss payment in the risk consideration when deciding on the selection of an SLA.
Service contract formalizability and matchmaking is limited to one provider. For the service contract, analogous to the service credit, the challenge for both practice and research is to consider it in the risk management of the cloud application decision.

6 Threats to Validity

To demonstrate rigor and encourage further research, the threats of validity of this study are discussed. Threats of validity are considered in terms of internal validity and external validity.
Internal validity refers specifically to whether an experimental treatment or condition makes a difference to the outcome or not, and whether there is sufficient evidence to substantiate the claim [35]. The following internal validity threats for this study were identified and should be considered for interpretation or further research.

- Insufficient or improper SLA documents or information were collected by the formalization procedure. Accordingly, it may be that the scores calculated for the categories may be inaccurate.
- The selected sample of leading public cloud SaaS providers is chosen biased or is insufficient. Potentially, adding more providers improves calculated category scores and leads to more underrepresented parameters.
- The interpretation of the descriptions or the order of importance in the SLAs was inaccurate or wrong. Our focus group was set up to evaluate the formalization; another group may possibly come to different results.
- The evaluation of the matchmaking of the parameters was performed with the knowledge of existing methods in SLA management. The actual applicability of the labeled parameters is nevertheless to be verified in each case. The evaluation of SLA

management artifacts with the identified parameters is a promising field for further research.
- The survey is statically time-based, so changes in SLAs over time can lead to other results.

External validity refers specifically to whether the results can be considered in real-world environment [35]. The following external validity threats for this study were identified and should be considered for interpretation or further research.

- It is possible that (1.) agreements besides the SLA between provider and customer are made or (2.) further contractual documents affect the consideration.
- The generalization of our identified formalizability and matchmaking can be challenged due to different requirements of the different application types.

7 Conclusion

In this study, based on the motivation of an as-is analysis, five present-day public cloud SaaS SLAs were analyzed in the context of service level compliance and risk management. The study focus was intentionally set on (I) public cloud, (II) SaaS and (III) business-critical applications in order to address the relevance of downtime-related breakdowns in business processes.
With the help of four derived SLA categories and 18 underlying SLA parameters, a general formalizability (concerning at least one provider) was determined. Across the four different categories, formalizability was found to range from an average 20% to 73.3% across the entire sample (concerning RQ1). The high variance in formalizability confirms the common assumption of the lack of SLA standards in practice.
To enable semi-automated SLA management, all parameters were evaluated for matchmaking (comparable and aggregable). Across the four different categories, matchmaking was found to range from an average of 20% to 73.3% across the entire sample (concerning RQ2). Matchmaking has high importance in the context of IT-supported SLA management, and is threatened especially due to low rates (20%, 30% and 40%) in three out of four categories. The emerging deficit can be closed on the one hand (I) by further analysis of matchmaking or (II) by an additional manual evaluation step of potential cloud customers.
An extract of contributions to research and practice elaborated in section 5 are finally summarized (concerning RQ3).

- Practitioners get an understanding of common and uncommon public cloud SaaS SLA parameters and categories to analyze risk and service level compliance prior to an adoption.
- It is also possible for practitioners to identify SLA aspects that may have a high economic importance (e.g. RPO, RTO, planned downtimes) but may not be offered by all providers.

- Researchers get an up-to-date view of SLA parameters that SLA management methods and concepts must take into consideration in order to be applicable to present-day clouds (e.g. temporal logic for downtime aggregation).
- Researchers should consider how business-critical SLA parameters (e.g., service credit calculation and downtime exclusion) can be reflected in terms of risk assessment extending traditional QoS aggregation (e.g., availability multiplication).

References

1. Gartner Inc.: Proportion of Enterprise IT Spending on Public Cloud Computing Continues to Increase, https://www.gartner.com/en/newsroom/press-releases/2020-11-17-gartner-forecasts-worldwide-public-cloud-end-user-spending-to-grow-18-percent-in-2021 (Accessed: 12.01.2021)
2. Synergy Research Group: Quarterly SaaS Spending Reaches $20 billion as Microsoft Extends its Market Leadership, https://www.srgresearch.com/articles/quarterly-saas-spending-reaches-20-billion-microsoft-extends-its-market-leadership (Accessed: 12.01.2021)
3. B. Link: Considering the Company's Characteristics in Choosing between SaaS vs. On-Premise-ERPs. In: Wirtschaftsinformatik (2013)
4. Bennett, K., Munro, M., Gold, N., Layzell, P., Budgen, D., Brereton, P.: An Architectural model for service-based software with ultra rapid evolution. In: Proceedings, IEEE International Conference on Software Maintenance. Systems and software evolution in the era of the Internet : Florence, Italy, 7-9 November, 2001. IEEE Computer Society, Los Alamitos, Calif. (2001)
5. Benlian, A., Hess, T., Buxmann, P.: Drivers of SaaS-Adoption – An Empirical Study of Different Application Types. Bus. Inf. Syst. Eng. 1, 357–369 (2009)
6. M. Janssen, Anton Joha: Challenges for adopting cloud-based software as a service (saas) in the public sector. In: ECIS (2011)
7. Baset, S.A.: Cloud SLAs: Present and Future. SIGOPS Oper. Syst. Rev. 46, 57–66 (2012)
8. Gulia, P., Sood, S.: Comparative Analysis of Present Day Clouds using Service Level Agreements. IJCA 71, 1–8 (2013)
9. Paquette, S., Jaeger, P.T., Wilson, S.C.: Identifying the security risks associated with governmental use of cloud computing. Government Information Quarterly 27, 245–253 (2010)
10. Patel, P., Ranabahu, A.H., Sheth, A.P.: Service Level Agreement in Cloud Computing. Kno.e.sis Publications, THE OHIO CENTER OF EXCELLENCE IN KNOWLEDGE-ENABLED COMPUTING, 1–10 (2009)
11. P. Hoberg, J. Wollersheim, H. Krcmar: The Business Perspective on Cloud Computing - A Literature Review of Research on Cloud Computing. In: AMCIS (2012)
12. Seifert, M., Kühnel, S.: "HySLAC"– A Conceptual Model for Service Level Agreement Compliance in Hybrid Cloud Architectures. Proceedings of the 50th

Annual Conference of the German Informatics Society, Lecture Notes in Informatics (LNI), 195–208 (2021)
13. Sun, W., Zhang, X., Guo, C.J., Sun, P., Su, H.: Software as a Service: Configuration and Customization Perspectives. In: 2008 IEEE Congress on Services Part II (services-2 2008), pp. 18–25. IEEE (2008 - 2008)
14. Andreas Jede, Frank Teuteberg: Understanding Socio-Technical Impacts Arising from Software-as-a-Service Usage in Companies. Bus Inf Syst Eng 58, 161–176 (2016)
15. Kritikos, K., Plexousakis, D.: Requirements for QoS-Based Web Service Description and Discovery. IEEE Trans. Serv. Comput. 2, 320–337 (2009)
16. Vakili, A., Navimipour, N.J.: Comprehensive and systematic review of the service composition mechanisms in the cloud environments. Journal of Network and Computer Applications 81, 24–36 (2017)
17. L. Qi, W. Dou, X. Zhang, J. Chen: A QoS-aware composition method supporting cross-platform service invocation in cloud environment. J. Comput. Syst. Sci. 78, 1316–1329 (2012)
18. Xin Zhao, Liwei Shen, Xin Peng, Wenyun Zhao: Toward SLA-constrained service composition: An approach based on a fuzzy linguistic preference model and an evolutionary algorithm. Information Sciences 316, 370–396 (2015)
19. Kritikos, K., Plexousakis, D., Plebani, P.: Semantic SLAs for Services with Q-SLA. Procedia Computer Science 97, 24–33 (2016)
20. Labidi, T., Mtibaa, A., Brabra, H.: CSLAOnto: A Comprehensive Ontological SLA Model in Cloud Computing. J Data Semant 5, 179–193 (2016)
21. Morgan, D.: Focus Groups as Qualitative Research. SAGE Publications, Inc, 2455 Teller Road, Thousand Oaks California 91320 United States of America (1997)
22. Mell, P.M., Grance, T.: The NIST definition of cloud computing. National Institute of Standards and Technology, Gaithersburg, MD (2011)
23. Yangui, S., Ravindran, P., Bibani, O., Glitho, R.H., Ben Hadj-Alouane, N., Morrow, M.J., Polakos, P.A.: A platform as-a-service for hybrid cloud/fog environments. In: IEEE LANMAN 2016. The 22nd IEEE International Symposium on Local and Metropolitan Area Networks, June 13-15, 2016, Rome, Italy, pp. 1–7. IEEE, Piscataway, NJ (2016)
24. Goyal, S.: Public vs Private vs Hybrid vs Community - Cloud Computing: A Critical Review. IJCNIS 6, 20–29 (2014)
25. Breiter, G., Naik, V.K.: A Framework for Controlling and Managing Hybrid Cloud Service Integration. In: Campbell, R. (ed.) 2013 IEEE International Conference on Cloud Engineering (IC2E). 25 - 27 March 2013, San Francisco Bay, California, pp. 217–224. IEEE, Piscataway, NJ (2013)
26. Suakanto, S., Supangkat, S.H., Suhardi, Saragih, R.: Performance Measurement of Cloud Computing Services (2012)
27. Kritikos, K., Plexousakis, D.: Semantic QoS Metric Matching. In: Bernstein, A., Gschwind, T., Zimmermann, W. (eds.) 4th European Conference on Web Services, 2006. ECOWS '06 ; 4 - 6 December 2006, Zurich, Switzerland. IEEE Computer Society, Los Alamitos, Calif. [u.a.] (2006)

28. Adobe Inc.: Service Level Agreement, https://www.adobe.com/content/dam/cc/en/legal/terms/enterprise/pdfs/MasterSLA-2016DEC5.pdf
29. Adobe Inc.: Service Level Exhibit - AEM as a Cloud Service, https://www.adobe.com/content/dam/cc/en/legal/terms/enterprise/pdfs/SLAExhibit-AEMCloudService-2019DEC12.pdf (Accessed: 12.01.2021)
30. Microsoft Corporation: Service Level Agreement for Microsoft Online Services, https://www.microsoftvolumelicensing.com/Downloader.aspx?documenttype=OSCS&lang=English (Accessed: 12.01.2021)
31. Oracle Corporation: Oracle SaaS Public Cloud Services-Pillar Document, https://www.oracle.com/assets/saas-public-cloud-services-pillar-3610529.pdf (Accessed: 12.01.2021)
32. Salesforce Inc.: Master Subscription Agreement, https://c1.sfdcstatic.com/content/dam/web/en_us/www/documents/legal/salesforce_MSA.pdf (Accessed: 12.01.2021)
33. Salesforce Inc.: Preferred Salesforce Maintenance Schedule, https://help.salesforce.com/articleView?id=000331027&type=1&mode=1 (Accessed: 12.01.2021)
34. SAP SE: Service Level Agreement for SAP Cloud Services, https://assets.cdn.sap.com/agreements/product-use-and-support-terms/cls/en/service-level-agreement-for-sap-cloud-services-english-v7-2020.pdf (Accessed: 12.01.2021)
35. J. Siegmund, N. Siegmund, S. Apel: Views on Internal and External Validity in Empirical Software Engineering. In: 2015 IEEE/ACM 37th IEEE International Conference on Software Engineering, pp. 9–19 (2015)