
**Information technology — Cloud
computing — Service level agreement
(SLA) framework —**

**Part 1:
Overview and concepts**

*Technologies de l'information — Informatique en nuage — Cadre de
travail de l'accord du niveau de service —*

Partie 1: Aperçu général et concepts

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19086-1:2016



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2016, Published in Switzerland

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Ch. de Blandonnet 8 • CP 401
CH-1214 Vernier, Geneva, Switzerland
Tel. +41 22 749 01 11
Fax +41 22 749 09 47
copyright@iso.org
www.iso.org

Contents

	Page
Foreword	v
Introduction	vi
1 Scope	1
2 Normative references	1
3 Terms and definitions	1
4 Symbols and abbreviated terms	4
5 Overview of SLAs for cloud services	5
6 Relationship between the cloud service agreement and cloud SLAs	6
7 Cloud SLA management best practices	7
7.1 General	7
7.2 Design	7
7.3 Evaluation and acceptance	7
7.4 Implementation and execution	8
7.5 Changes to the cloud SLA	8
8 The role of cloud service level objectives, cloud service qualitative objectives, metrics, remedies and exceptions in the cloud SLA	8
8.1 General	8
8.2 Metrics	8
8.3 SLOs and SQOs	9
8.3.1 Service levels	9
8.3.2 Cloud service level objectives	9
8.3.3 Cloud service qualitative objectives	9
8.4 Remedies and claims	10
8.4.1 Remedies	10
8.4.2 Claims process	10
8.5 Exceptions	10
9 Cloud SLA components	10
9.1 General	10
9.2 Covered services component	10
9.2.1 Description	10
9.2.2 Relevance	11
9.3 Cloud SLA definitions component	11
9.3.1 Description	11
9.3.2 Relevance	11
9.4 Service monitoring component	11
9.4.1 Description	11
9.4.2 Relevance	11
9.4.3 Cloud service qualitative objectives	11
9.5 Roles and responsibilities component	11
9.5.1 Description	11
9.5.2 Relevance	12
10 Cloud SLA content areas and their components	12
10.1 General	12
10.2 Accessibility content area	12
10.2.1 Accessibility component	12
10.3 Availability content area	13
10.3.1 Availability component	13
10.4 Cloud service performance content area	13
10.4.1 General	13
10.4.2 Cloud service response time component	13

10.4.3	Cloud service capacity component	14
10.4.4	Elasticity component	15
10.5	Protection of personally identifiable information (PII) content area	16
10.5.1	Protection of PII component	16
10.6	Information Security content area	17
10.6.1	Information Security component	17
10.7	Termination of service content area	18
10.7.1	Termination of service component	18
10.8	Cloud service support content area	19
10.8.1	Cloud service support component	19
10.9	Governance content area	21
10.9.1	Governance component	21
10.10	Changes to the cloud service features and functionality content area	22
10.10.1	Changes to the cloud service features and functionality component	22
10.11	Service reliability content area	23
10.11.1	General	23
10.11.2	Service resilience/fault tolerance component	23
10.11.3	Customer data backup and restore component	24
10.11.4	Disaster recovery component	25
10.12	Data management content area	26
10.12.1	General	26
10.12.2	Intellectual property rights (IPR) component	27
10.12.3	Cloud service customer data component	27
10.12.4	Cloud service provider data component	28
10.12.5	Account data component	28
10.12.6	Derived Data component	28
10.12.7	Data portability component	29
10.12.8	Data deletion component	29
10.12.9	Data location component	30
10.12.10		
	Data examination component	31
10.12.11		
	Law enforcement access component	31
10.13	Attestations, certifications and audits content area	31
10.13.1	Attestations, certifications and audits component	31
Bibliography		33

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: [Foreword - Supplementary information](#)

The committee responsible for this document is ISO/IEC JTC 1, *Information technology*, Subcommittee SC 38, *Cloud computing and distributed platforms*.

A list of all parts in the ISO/IEC 19086 series can be found on the ISO website.

Introduction

This document provides an overview, foundational concepts, and definitions for the cloud SLA framework. ISO/IEC 19086 builds on the cloud computing concepts defined in ISO/IEC 17788 and ISO/IEC 17789. This document establishes a common framework for helping organizations to understand the purpose of all the parts of ISO/IEC 19086 and the relationships between those parts. It also identifies other documents that have relationships with ISO/IEC 19086 and which are useful in understanding cloud SLAs.

This document can be used by any organization or individual involved in the creation, modification or understanding of a cloud service level agreement which conforms to ISO/IEC 19086. The cloud SLA should account for the key characteristics of a cloud computing service and needs to facilitate a common understanding between cloud service providers and cloud service customers.

In particular, it defines the following fundamental concepts of the cloud SLA framework:

- Cloud Service Agreement (CSA)
- Cloud Service Level Agreement (SLA)
- Cloud Service Level Objectives (SLO)
- Cloud Service Qualitative Objectives (SQO)

This document also describes the content areas and components that consist of a list of SLOs and SQOs.

- ISO/IEC 19086-2 provides the metrics model to be used for creating metrics used in SLOs and SQOs.
- ISO/IEC 19086-3 provides the core conformance requirements derived from the SLOs and SQOs defined in this document.
- ISO/IEC 19086-4 builds upon the foundational concepts and definitions described by this document by describing specific components and the conformance requirements for SLOs and SQOs in the area of Security and Privacy.

More specifically, this document

- a) promotes cohesion between the parts of ISO/IEC 19086 by explaining the concepts and terminology used across all parts,
- b) contributes to the understanding of ISO/IEC 19086 by clarifying the relationships between all the parts, and
- c) provides an overview of other International Standards which can be used in combination with ISO/IEC 19086.

[Figure 1](#) represents an overview of the content of ISO/IEC 19086 and the relationships between the parts of ISO/IEC 19086 and other key International Standards relating to cloud computing.




Figure 1 — Relationship of parts of ISO/IEC 19086 and other cloud computing standards

This document addresses the contents of a cloud SLA in two main groupings: SLA Components, addressed in [Clause 9](#), and SLA Content Areas, addressed in [Clause 10](#), as shown in [Figure 2](#).




Figure 2 — SLA components and SLA content areas

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 19086-1:2016

Information technology — Cloud computing — Service level agreement (SLA) framework —

Part 1: Overview and concepts

1 Scope

This document seeks to establish a set of common cloud SLA building blocks (concepts, terms, definitions, contexts) that can be used to create cloud Service Level Agreements (SLAs).

This document specifies

- a) an overview of cloud SLAs,
- b) identification of the relationship between the cloud service agreement and the cloud SLA,
- c) concepts that can be used to build cloud SLAs, and
- d) terms commonly used in cloud SLAs.

This document is for the benefit and use of both cloud service providers and cloud service customers. The aim is to avoid confusion and facilitate a common understanding between cloud service providers and cloud service customers. Cloud service agreements and their associated cloud SLAs vary between cloud service providers, and in some cases different cloud service customers can negotiate different contract terms with the same cloud service provider for the same cloud service. This document aims to assist cloud service customers when they compare cloud services from different cloud service providers.

This document does not provide a standard structure that can be used for a cloud SLA or a standard set of cloud service level objectives (SLOs) and cloud service qualitative objectives (SQOs) that will apply to all cloud services or all cloud service providers. This approach provides flexibility for cloud service providers in tailoring their cloud SLAs to the particular characteristics of the offered cloud services.

This document does not supersede any legal requirement.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 17788:2014, *Information technology — Cloud computing — Overview and vocabulary*

ISO/IEC 17789, *Information technology — Cloud computing — Reference architecture*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC 17788 and the following apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

— IEC Electropedia: available at <http://www.electropedia.org/>

— ISO Online browsing platform: available at <http://www.iso.org/obp>

3.1 accessibility

usability of a product, service, environment or facility by people within the widest range of capabilities

Note 1 to entry: The concept of accessibility addresses the full range of user capabilities and is not limited to users who are formally recognized as having disability.

Note 2 to entry: The usability-oriented concept of accessibility aims to achieve levels of effectiveness, efficiency and satisfaction that are as high as possible considering the specified context of use, while paying attention to the full range of capabilities within the user population.

Note 3 to entry: It is important in the context of ISO/IEC 19086 to distinguish between the specialized meaning of “accessibility” as defined here and the term “accessible” which is used with its dictionary meaning of “able to be reached or entered.”

[SOURCE: ISO 9241-171:2008, 3.2]

3.2 business continuity

capability of the organization to continue delivery of products or services at acceptable predefined levels following disruptive incident

[SOURCE: ISO/IEC 22301:2012, 3.3]

3.3 cloud service agreement

documented agreement between the cloud service provider and cloud service customer that governs the covered service(s)

Note 1 to entry: A cloud service agreement can consist of one or more parts recorded in one or more documents.

3.4 cloud service level agreement cloud SLA

part of the cloud service agreement (3.3) that includes cloud service level objectives (3.5) and cloud service qualitative objectives (3.6) for the covered cloud service(s)

3.5 cloud service level objective SLO

commitment a cloud service provider makes for a specific, quantitative characteristic of a cloud service, where the value follows the interval scale (3.9) or ratio scale (3.17)

Note 1 to entry: An SLO commitment may be expressed as a range.

3.6 cloud service qualitative objective SQO

commitment a cloud service provider makes for a specific, qualitative characteristic of a cloud service, where the value follows the *nominal scale* (3.11) or *ordinal scale* (3.12)

Note 1 to entry: A cloud service qualitative objective may be expressed as an enumerated list.

Note 2 to entry: Qualitative characteristics typically require human interpretation.

Note 3 to entry: The ordinal scale allows for existence/non-existence.

3.7**disaster recovery**

ability of the ICT elements of an organization to support its critical business functions to an acceptable level within a predetermined period of time following a disaster

[SOURCE: ISO/IEC 27031:2011, 3.7]

3.8**failure notification policy**

policy specifying the processes by which the cloud service customer and cloud service partner can notify the cloud service provider of a service outage and by which the cloud service provider can notify the cloud service customer and cloud service partner that a service outage has occurred.

Note 1 to entry: The policy may also include the process for providing updates on service outages, who receives notifications and updates, the maximum time between the detection of a service outage and the issuance of a notice of service outage, the maximum time interval between service outage updates and how service outage updates are described.

3.9**interval scale**

continuous scale or discrete scale with equal sized scale values and an arbitrary zero

[SOURCE: ISO 3534-2:2006, 1.1.8]

3.10**metric**

standard of measurement that defines the conditions and the rules for performing the measurement and for understanding the results of a measurement

Note 1 to entry: A metric implements a particular abstract metric concept.

Note 2 to entry: A metric is to be applied in practice within a given context that requires specific properties to be measured, at a given time(s) for a specific goal.

3.11**nominal scale**

scale with unordered labelled categories or ordered by convention

[SOURCE: ISO 3534-2:2006, 1.1.6]

3.12**ordinal scale**

scale with ordered labelled categories

[SOURCE: ISO 3534-2:2006, 1.1.7]

3.13**personally identifiable information****PII**

any information that (a) can be used to identify the PII principal to whom such information relates, or (b) is or might be directly or indirectly linked to a PII principal

Note 1 to entry: To determine whether a PII principal is identifiable, account should be taken of all the means which can reasonably be used by the privacy stakeholder holding the data, or by any other party, to identify that natural person.

[SOURCE: ISO/IEC 29100:2011, 2.9]

3.14

PII controller

privacy stakeholder (or privacy stakeholders) that determines the purposes and means for processing personally identifiable information (PII) other than natural persons who use data for personal purposes

[SOURCE: ISO/IEC 29100:2011, 2.10]

3.15

PII principal

natural person to whom the personally identifiable information (PII) relates

[SOURCE: ISO/IEC 29100:2011, 2.11]

3.16

PII processor

privacy stakeholder that processes personally identifiable information (PII) on behalf of and in accordance with the instructions of a PII controller

[SOURCE: ISO/IEC 29100:2011, 2.12]

3.17

ratio scale

continuous scale with equal sized scale values and an absolute or natural zero point

[SOURCE: ISO 3534-2:2006, 1.1.9]

3.18

remedy

compensation available to the cloud service customer in the event the cloud service provider fails to meet a specified cloud service level objective (3.5)

Note 1 to entry: This definition of the term in English is based on the “legal reparation” meaning defined in The Shorter Oxford English Dictionary.

3.19

resilience

ability of a cloud service to recover operational condition quickly after a fault occurs

4 Symbols and abbreviated terms

BLOB	Binary Large Object
CSA	Cloud Service Agreement
CSC	Cloud Service Customer
CSP	Cloud Service Provider
ICT	Information and Communications Technology
IPR	Intellectual Property Rights
IT	Information Technology
PII	Personally Identifiable Information
RPO	Recovery Point Objective
RTO	Recovery Time Objective

SLA	Service Level Agreement
SLO	Cloud Service Level Objective
SQO	Cloud Service Qualitative Objective
VM	Virtual Machine

5 Overview of SLAs for cloud services

A cloud service level agreement (cloud SLA) is a part of the cloud service agreement that includes cloud service level objectives and cloud service qualitative objectives for the covered cloud service(s). The cloud SLA should account for the key characteristics of cloud computing as described in ISO/IEC 17788:2014, 6.2 that include the following.

- **On-demand self-service** — A CSC may gain access to cloud services without human interaction with the CSP. The cloud service agreement (CSA) (see [Clause 6](#)) and the associated cloud SLA may be presented and agreed through software tools and financial arrangements that are automated.
- **Resource pooling** — The public cloud deployment models allow sharing resources across many CSCs that do not have a relationship. The private cloud models allow users to share resources within the same organization. The hybrid cloud models allow users to share some resources within the same organization and some resources across many CSCs that do not necessarily have a relationship with one another. The community cloud deployment models allow sharing resources across CSCs that have some relationship.
- **Multi-tenancy** — Cloud environments are enabled through the use of large-scale virtualization of servers, storage and networks. Overall system usage is typically spread over many CSCs. Multi-tenancy allows sharing of resources in such a way that multiple tenants and their computations and data are isolated from and inaccessible to one another. Cloud environments typically have no persistent relationship between particular physical resources and their use by CSCs. The CSCs are assigned virtual resources, and logging of usage is done at this level of abstraction.
- **Rapid elasticity and scalability** — A characteristic of cloud computing where physical or virtual resources can be rapidly and elastically adjusted, in some cases automatically, to quickly increase or decrease resources.
- **Tradeoff between cost and control** — Large-scale, standardized cloud services may be provided on a low unit cost, utility basis, in conjunction with standardized contracts and cloud SLAs. If a CSC requires more control and customization of cloud services than is available from a standard utility service model, then this may be provided at an additional cost and with a specific cloud SLA.
- **Measured service** — A feature where the metered delivery of cloud services is such that usage can be monitored, controlled, reported and billed. This is an important feature needed to optimize and validate the delivered cloud service. The focus of this key characteristic is that the CSC may only pay for the resources that they use.
- **Broad network access** — The capabilities of cloud services are made available over the network and are typically accessed through standard mechanisms that promote use by heterogeneous client platforms (for example, access through mobile phones, laptops and workstations).

Details of cloud SLAs, SLOs and SQOs can vary for different cloud service categories, cloud capabilities types and different cloud deployment models (see ISO/IEC 17788). Cloud SLAs in this document are intended to be useful for CSCs and CSPs across the variety of cloud service categories and cloud deployment models. As the definitions of SLOs and SQOs are intended to be technology and business model neutral, so not all of these SLOs or SQOs will apply to every cloud service, and those that do apply may be structured and applied in different ways to specific cloud services. For example, service availability can be measured in different ways, some of which depend on the specific cloud service: a

computational cloud service is different from an email cloud service, and service availability for each will be computed differently.

6 Relationship between the cloud service agreement and cloud SLAs

Cloud services, particularly public cloud services, generally involve an agreement between the CSC and the CSP concerning the acquisition and use of the cloud services. For the purposes of this document, the legal agreement is referred to as the “Cloud Service Agreement” or CSA. The CSA has a number of synonyms such as “Master Service Agreement,” “Customer Agreement,” “Terms of Service,” or simply “Agreement.”

A CSA comprises one or more parts recorded in one or more documents. Contents of each part can appear in more than one document. There is no normative relationship between parts and documents, i.e. a part does not have to be in a single document, and a document does not have to contain a whole part. There is neither a standard naming convention for the parts or documents of a CSA, nor a standard structure for the documents or parts.

Examples of common parts of CSAs include the following.

- Cloud Service Level Agreement (cloud SLA)

The cloud SLA ordinarily contains a collection of SLOs and SQOs relating to the cloud service, covering aspects of the service. This might include availability, reliability, performance, security, data protection, compliance and data handling.

- Acceptable Use Policy

The acceptable use policy usually defines boundaries for the CSC’s use of the cloud service. This might include restrictions that prevent the CSC from installing malware on the cloud service or limit the kind of data that can be stored.

- Security Policy

The security policy typically describes responsibilities that apply to the CSC and to the CSP, SLOs and SQOs which the CSP applies to the cloud service in security terms and potentially indicates which security certifications or standards are met by the cloud service.

- Data Protection Policy

The data protection policy typically deals with the handling of personal data or sensitive data by the cloud service, including SQOs for specific data protection measures and privacy certifications or standards that apply to the service.

- Business Continuity Policy

The business continuity policy typically deals with the resilience aspects of the cloud service and can include measures that are implemented by the CSP to avoid data loss and to deal with outages, such as backups and redundant components.

- Upgrade Policy

The upgrade policy usually covers changes to the features and functions for the covered services and related management interface changes. Periodic updates are also usually covered by the upgrade policy.

- Termination Policy

The Termination Policy usually deals with the issues that arise when a CSC terminates their use of one or more cloud services. The termination policy might include SQOs for areas such as notifications, data reversibility and data deletion.

The content of each of the parts and the number of parts in the CSA can vary between different cloud services — any particular item, including the SLOs and SQOs described in this document, can appear in different parts for different services. As an example, security SLOs and SQOs might be described in the Security Policy or they might appear in the cloud SLA. However, it is important for the CSC to know the complete set of documents that govern the cloud service, and the CSA is expected to reference all applicable documents.

7 Cloud SLA management best practices

7.1 General

Cloud SLA management covers the issues related to cloud SLA design, evaluation, negotiation and acceptance, implementation and execution and changes to the cloud SLA. CSCs should ensure that cloud SLAs and other governing documents align with their business cases and overall strategy. CSCs should be aware that there could be several documents that govern a cloud service. See [Clause 6](#) for more details.

7.2 Design

A cloud SLA applies to the covered services identified within the cloud SLA. A single cloud SLA may apply to multiple CSCs or to a single CSC. In cases where the cloud SLA is designed jointly by the CSC and the CSP, both parties should undertake the steps outlined in 7.2 together. A CSP should design the cloud SLAs to meet the needs of their CSCs and to align with the capabilities of the covered services.

The cloud SLA design process should account for the appropriate roles. ISO/IEC 17789 can be used as a reference for determining the appropriate key roles. Key roles for the cloud SLA design process discussed in ISO/IEC 17789 include

- cloud service customer, party which is in a business relationship for the purpose of using cloud services,
- cloud service provider, party which makes cloud services available, and
- cloud service partner, party which is engaged in support of, or auxiliary to, activities of either the CSP or the CSC, or both.

ISO/IEC 17789 contains concepts that can be used in the construction of a cloud SLA. The choice of concepts to be included in a cloud SLA depends on the cloud service and on the business context. The process to change the cloud SLA and notify CSCs of changes should be within the cloud SLA or other governing document.

The design phase should also consider the mechanisms that the CSC and the CSP can use to monitor each service characteristic and report failures to meet SLO and SQO commitments.

7.3 Evaluation and acceptance

CSCs can use this document as a reference when evaluating cloud SLAs. CSCs can review all of the concepts and determine which are critical to their business objectives. CSCs can then consider the CSP's cloud SLAs (and remedies) in evaluating whether the CSP's services meet the CSC's business objectives.

CSCs can review cloud SLAs in the context of their organization's business policies and other requirements, and can identify which SLOs, SQOs and service features are important for each of their use cases. Standards such as the ISO/IEC 20000 series and the ISO/IEC 27000 series may be referenced in cloud SLAs or other documentation. In some cases, CSPs certify their conformance with specific industry standards. CSCs can determine what standards are important to their business objectives or are desirable for their organization's governance and determine whether a cloud service is certified to conform with that standard, or if it is referenced in the cloud SLA or other documentation. CSCs can determine how they monitor cloud service characteristics and report failures to meet SLO and SQO

commitments by familiarizing themselves with a CSP's failure notification policy, if available. Methods for monitoring and reporting include management systems, connected devices apps, web portals, email, text messages, telephone and posts to social media sites.

Acceptance of a cloud SLA may occur by clicking a check box on a web page, by registering for the cloud service, or by a formal signing of the agreement by both parties. Each party should ensure they are ready to undertake the implementation and execution of the cloud SLA, independent of the means of acceptance. In the case of an agreement with unique terms, both parties should ensure they are prepared to support implementation and execution of those terms.

7.4 Implementation and execution

Implementing a cloud SLA involves setting up processes for monitoring and managing cloud service characteristics, reporting any failures to meet SLOs and SQOs and claiming any remedies. In some cases, the CSP may need to collaborate with the CSC when implementing the cloud SLA. CSCs should also include the cloud SLA in their internal governance. The choice of concepts to be included in cloud SLA monitoring and auditing processes depends on the cloud service and the business context.

Execution of the cloud SLA involves the provision and operation of the cloud service by the CSP, including the management and monitoring of service levels. If a CSC believes an SLO or SQO has not been met, the CSC may follow the failure notification policy.

7.5 Changes to the cloud SLA

Change is an inevitable part of any ICT system, and cloud SLAs are no exception, whether due capability change or evolution of CSC requirements. The CSP can include a process to make changes to the cloud SLA and provide notifications to CSCs. The CSP can also include mechanisms that allow the CSCs to request changes to the cloud SLA.

The CSCs can evaluate the cloud SLA change and notification processes in the evaluation and acceptance phase. CSCs can also determine whether the current SLAs or proposed SLA meet their business objectives and if not, request changes to the SLA.

8 The role of cloud service level objectives, cloud service qualitative objectives, metrics, remedies and exceptions in the cloud SLA

8.1 General

It is essential to be able to monitor the cloud service to ensure that the SLOs and SQOs in the cloud SLA are being met. It is important to have remedies described within the cloud SLA or a related document in the event that an objective is not met. Finally, there might be events or incidents that are declared as exceptions. In those cases, even though an SLO or SQO is not met, the related remedy is not triggered.

8.2 Metrics

The definition and usage of appropriate metrics and their underlying measures and measurements are an essential aspect of the cloud SLA. The metrics are used to set the boundaries and margins of error the CSP abides by and their limitations. These metrics may be used at runtime for service monitoring or remediation.

Cloud service metrics address the following needs (list non-exhaustive).

- Determine if SLOs are met.
- Categorize service capabilities.
- Define a purpose for measures and measurements.

- Deliver a consistent representation of measure and measurement information.
- Link properties, measurements and metrics.
- Enable comparison of monitoring between services.
- Determine cloud service effectiveness for business objectives.

Cloud service metrics need to be used in the context of the cloud SLA and a given cloud service. Metrics help define the cloud service properties to be measured. The metric can be defined in terms of its related measures, measurements, relevant parameters, and calculating formulas, as well as measure and measurement rules. The metrics can be used to determine whether a measurement of a cloud service property at a specific point is within stated boundaries of the property. Using a standard set of metrics in the cloud SLA makes it easier and faster to define cloud SLAs and SLOs, and to compare one cloud SLA to another.

Without proper metrics, it is difficult or impossible to enforce a cloud SLA.

8.3 SLOs and SQOs

8.3.1 Service levels

Service levels are measurement results for specific attributes of the cloud service and may change while the service is operating. Service levels are expressed using metrics that may be based on a single measure or may be calculated using several different measures.

Service levels are reported against one or more SLOs (which ordinarily are fixed during service use) and they are often described in the context of the service covered. For example, a “service availability” service level for compute services might be reported against an SLO of “uptime” where uptime would be described in the context of compute as instances being accessible and usable upon demand by an authorized entity, while a “service availability” SLO for storage services might be reported against an SLO of “uptime” described in the context of requests for the stored object returning errors. Specific SLOs, and remedies, are context driven and are clearly defined in the cloud SLA.

Cloud SLOs and cloud SLA content can vary for different cloud deployment models and different cloud service categories.

8.3.2 Cloud service level objectives

SLOs are commitments a CSP makes for specific, quantitative characteristics of a cloud service. Each service level ordinarily has a target commitment that the CSP agrees to meet, typically either a single boundary value or a range of values. For example, for a “service availability” service level, the SLO might be “uptime” and the target might be expressed as a percentage (of uptime to total time) which is a lower bound for the service level, or the SLO might be “downtime” and the target might be a percentage or an amount of time over a time interval which is an upper bound for the service level.

The CSP may offer a range of SLOs and associated remedies.

8.3.3 Cloud service qualitative objectives

SQOs are commitments a CSP makes for specific, qualitative characteristics of a cloud service, where the value follows the nominal or ordinal scale. The ordinal scale allows for cases where the characteristic either exists or does not exist (e.g. “true or false”).

SQO observations require human interpretation and cannot be manipulated algebraically. For example, “security certification” could be an SQO and a related assurance could be “CSP will maintain a current certification for ISO/IEC 27001.” Verification of such an assurance could be the availability of a scanned copy of the certificate. Verification of the performance on an assurance related to an SQO may take many forms including the process of legal discovery.

8.4 Remedies and claims

8.4.1 Remedies

Remedies may be provided by the CSP to the CSC in the event the cloud service fails to meet the SLOs or SQOs defined in the cloud SLA. Remedies for failure to achieve SLOs or SQOs stated in the cloud SLA may take different forms such as refunds on charges, free services, or other forms of compensation.

8.4.2 Claims process

The claims process describes the process for a CSC to claim a remedy if a SLO or SQO is not met. It may be up to the CSC to determine when the cloud service has failed to meet its SLOs or SQOs and report claims to the CSP, while in other cases the CSP will monitor the service levels and automatically initiate claims. If the CSP provides monitoring, the cloud SLA claims (or lack of claim) may require verification.

8.5 Exceptions

Exceptions describe the circumstances under which the SLOs, SQOs and their associated remedies do not apply. These may vary between agreements and be subject to the laws of a particular jurisdiction. Examples of exceptions include scheduled outages, natural disasters and other factors beyond the control of the CSP.

9 Cloud SLA components

9.1 General

The cloud SLA components described in [Clauses 9](#) and [10](#) define concepts commonly used in cloud SLAs. It is important for both the CSC and CSP to have a shared understanding of these concepts. Also, it is important to recognize that these concepts and their associated terms and metrics are dependent on the context established by the cloud service which they cover. Listed SLOs and their associated metrics and listed SQOs in each cloud SLA component in [Clause 9](#) and [Clause 10](#) are not meant to be prescriptive or be an exhaustive list for CSPs to use in their cloud SLAs.

For each component, a description, its relevance and associated SLOs and SQOs are included. Some of the SLOs and SQOs are written as statements describing what would be included in the cloud SLA while other SLOs and SQOs are written describing the associated concepts.

As mentioned in [Clause 6](#), the cloud SLA components described below may be present across more than one document that a CSA comprises.

9.2 Covered services component

9.2.1 Description

The covered services component identifies the cloud services that are covered by the cloud SLA. All other portions of the cloud SLA apply to the services identified in the covered services component.

For example, a cloud SLA may state

“This Service Level Agreement applies to the following service(s) offered by XYZ Inc.:

- XYZ Online Mail
- XYZ BLOB Storage”

If a single cloud SLA covers multiple services, it could be necessary to provide SLOs and SQOs separately for each covered service.

9.2.2 Relevance

A CSP may offer any number of cloud services covered by one or more cloud SLAs. For a given cloud SLA, it is important for the CSC to know exactly which cloud services are covered.

9.3 Cloud SLA definitions component

9.3.1 Description

The definitions component includes terms that are unique to the CSP or that are particularly important to the understanding of the agreement. The definitions component is expected to use definitions from industry standards when possible.

9.3.2 Relevance

It is important for CSCs to understand the definitions of terms important to the cloud SLA as well as terms that have definitions unique to the SLA.

9.4 Service monitoring component

9.4.1 Description

The service monitoring component lists the parameters, for the covered services, that are monitored by the CSP and the data provided to the CSC. These parameters may include those described in this document and other parameters.

9.4.2 Relevance

CSPs may provide reporting capabilities or may provide monitoring tools for CSCs to monitor the performance of the service. Reporting capabilities and monitoring tools may enable CSCs to determine whether an SLO is being met or not.

9.4.3 Cloud service qualitative objectives

Monitoring Parameters

A list of parameters for the covered services that the CSP monitors and the data is provided to the CSC.

Monitoring Mechanisms

A list of mechanisms available to the CSC, such as logs, that includes a description of the monitored parameters and a description of any terms and conditions governing the availability of these mechanisms.

9.5 Roles and responsibilities component

9.5.1 Description

The roles and responsibilities component provides a description of roles and responsibilities for both the CSP and the CSC. Cloud computing involves a number of roles, both on the CSC side and also on the CSP side. Many of these roles are described in ISO/IEC 17789. A clear description of the roles that have relevance to a particular cloud service and the responsibilities of those roles is important for the successful use and operation of the cloud service.

9.5.2 Relevance

A clear description of the division of roles and responsibilities between the CSC and the CSP can help avoid confusion.

10 Cloud SLA content areas and their components

10.1 General

The cloud SLA content areas described below are areas in which a CSP could offer a cloud SLA. A cloud SLA content area is described by one or more SLA components. Cloud SLA content and their associated SLOs, SQOs and metrics are dependent on the context established by the particular cloud services being used.

10.2 Accessibility content area

10.2.1 Accessibility component

10.2.1.1 Description

The accessibility (3.1) component describes the assistive technologies the CSP implements as part of the covered services.

10.2.1.2 Relevance

Around the world, millions of people have disabilities that impede their ability to use Information and Communications Technology (ICT). Assistive technologies such as magnifiers, screen readers, braille readers and alternative input devices are available on client computing platforms to make the use of ICT, including cloud services, easier for people with disabilities.

There are standards available that can be used to provide cloud services to persons with disabilities including

- W3C Web Content Accessibility Guidelines (WCAG) 2.0, also published as ISO/IEC 40500:2012,
- ISO/IEC TR 29138 (all parts) — Accessibility considerations for people with disabilities, and
- ISO/IEC Guide 71 — Guide for addressing accessibility in standards.

Additionally, governments have policies and requirements related to accessible ICT such as Section 508 of the Rehabilitation Act of 1973 in the United States of America and EN 301 549 Accessibility requirements for public procurement of ICT products and services in Europe.

10.2.1.3 Cloud service qualitative objectives

Accessibility Standards

A statement listing accessibility related standards the CSP supports in the covered services.

Accessibility Policies

A statement listing policies and regulations for accessible ICT the CSP supports in the covered services.

10.3 Availability content area

10.3.1 Availability component

10.3.1.1 Description

Availability is the property of being accessible¹⁾ and usable upon demand by an authorized entity (ISO/IEC 17788). The availability component specifies the method for determining that the covered services are accessible and usable.

10.3.1.2 Relevance

Availability provides CSCs with a high-level indicator that the covered services are responding to requests and fulfilling the functions in the service description at a point in time.

Periods when the cloud service is not available are commonly known as “downtime”. There may be cases, such as “scheduled downtime” where the cloud service is not available for reasons other than failures. Periods where the cloud service is not available, but not counted as downtime, are termed “allowable downtime.”

Availability is often expressed as the total time in a defined interval less the downtime during that time interval.

10.3.1.3 Cloud service level objectives

Availability

The amount or percentage of time in a given period that the cloud service is accessible and usable.

Availability may be calculated as the total time over a set of defined intervals less the total downtime during each interval, and may exclude allowable downtime.

For additional details, see ISO/IEC 19086-2.

10.4 Cloud service performance content area

10.4.1 General

Cloud service performance includes individual components that can be used in a cloud SLA to express the performance of a cloud service. Note that definition of performance can vary depending on the cloud service, the CSP and (potentially) the CSC. The components listed here are not meant to be a prescriptive list.

10.4.2 Cloud service response time component

10.4.2.1 Description

There are several properties related to the responsiveness of the cloud service that can be included in the cloud SLA.

Response time is the time between a stimulus to the cloud service and the response of the service to the stimulus.

1) In the context of availability, accessible means that the cloud service can be successfully accessed via its interfaces.

10.4.2.2 Relevance

A stimulus may be initiated by a variety of sources such as CSCs, partners, CSPs and programmatic events.

Response times can vary between different cloud services even when the client system, workloads, and transaction mix are identical. When comparing response times, consider whether (a) the same metric definitions are used and (b) other variables can be held equal (e.g. network latency, network throughput, transmission delay, forwarding latency) in making those comparisons.

For additional details, see ISO/IEC 19086-2.

10.4.2.3 Cloud service level objectives

Cloud Service Maximum Response Time Observation

The maximum time between a defined stimulus or input to the cloud service and a defined point in the response.

Cloud Service Response Time Mean

Statistical mean over a set of cloud service response time observations.

The specification of the method of calculating the mean is defined for metrics associated with the cloud service.

Cloud Service Response Time Variance

Statistical variance describes how far from the mean response times are likely to be within a set of cloud service response time observations.

Variance is often described in terms of the standard deviation from the mean value of a set of measurements.

This SLO can be used to indicate the stability of response time over a time span — it shows how much variation or statistical dispersion exists from the average response time.

The variance of response time is also important when the user wants to measure the system performance under variable load levels and when the user wants to measure the system's performance for a special event.

10.4.3 Cloud service capacity component

10.4.3.1 Description

This subclause explains service properties related to the capacity of the service that can be included in the cloud SLA. These properties related to capacity include not only the capacity of the cloud resources (such as storage space, processing power) but the capacity of the network used to access the resources.

10.4.3.2 Relevance

As the capacity of a cloud system may change, it is important to understand what capacity exists and how it is usable by the CSC. In traditional asset-based systems, the capacity is fixed, while in a cloud system, the capacity is dynamic and needs to be tracked to ensure the system is meeting the CSC requirements. Charging for the cloud service may also be based on capacity and/or capacity limits.

10.4.3.3 Cloud service level objectives

Limit of Simultaneous Cloud Service Connections

Maximum number of simultaneous connections supported by the cloud service.

Limit of Available Cloud Service Resources

Maximum capacity of available resources, i.e. disk space, CPU power, memory size, page view, etc.

Cloud Service Throughput

The number of inputs or the amount of sets of inter-dependent inputs (i.e. a transaction) that can be processed in every unit of time by the cloud service. It is normally measured as web requests per second, page elements per second and transactions per second.

Cloud Service Bandwidth

The amount of data that can be transferred over a period of time.

10.4.4 Elasticity component

10.4.4.1 Description

The elasticity component describes the ability of a cloud service to dynamically adjust the amount of resources that are allocated to an instance of the service. The adjustment is performed on the basis of the current workload of the cloud service instance, i.e. increased workload results in the allocation of more resources, while decreased workload is answered by the de-allocation of resources.

Hence, elasticity depends on the availability of a procedure to monitor variations in the workload and react to this by the allocation of more resources or the de-allocation of resources that are no longer needed. The procedure can either be manual or automatic. In the manual case, it is the CSC's responsibility to assess the amount of resources needed according to the changing workload. The automatic case is executed without direct interaction from the CSC or the CSP. In the automatic case, the procedure can be reactive, i.e. based on monitoring of actual changes the current workload, or proactive, i.e. using some algorithm to predict the future load situation, or a combination of both.

Elasticity is related to one or more resource types. For instance, a virtual machine combines the following resources: number of processors of a given specification, size of memory, number of network interfaces, size of hard drive, etc. Resources for a Web-based application service can be the number of parallel user sessions and/or the number of parallel transactions.

10.4.4.2 Relevance

Elasticity is one of the main characteristics of cloud services. CSCs do not have to allocate a sufficient amount of resources for the cloud service during service configuration, but can rely on the ability of the cloud system to scale up resources over time as needed and scale down resources that are no longer needed. The CSC may add and reduce resources manually, or may use CSP tools that enable the CSC to set rules for scaling resources up and down. Therefore, CSCs do not have to over-provision resources for the cloud service and thus do not have to pay continuously for resources that they use only to deal with peaks in the workload.

10.4.4.3 Cloud service level objectives

Elasticity

Characterization of the elasticity of a cloud service can be evaluated in terms of two objectives:

Elasticity Speed

The elasticity speed quantity describes how fast a cloud service is able to react to a resource request when

- the CSC makes a resource re-allocation request (in the case of manual elasticity), or
- workload changes take place (in the case of automatic elasticity).

The speed quantity can be expressed within thresholds. This quantity can be determined by a measurement process and therefore defines a metric.

Elasticity Precision

The elasticity precision quantity describes how precise the resource allocation meets the actual resource requirements at a given point in time.

- In the manual case, precision depends on the granularity of the resource allocation, i.e. the minimum amount of resources that can be re-allocated. Hence, in the manual case, precision is a technical characteristic of the cloud service that does not require measurements (i.e. no metric is associated with it).
- In the automatic case, precision refers to the difference between the amount of resources that are allocated and the amount of resources that are actually needed (the optimum state) to cope with a given workload. The actual resource allocation may be over-provisioned (i.e. more resources are allocated than are actually needed), or under-provisioned (i.e. the amount of resources that are actually allocated is not sufficient to cope with the actual workload). As opposed to the manual case, in the automatic case, the difference between the allocated and the actually needed amount of resources can be determined by a measurement process and hence imply a metric.

The precision quantity can be expressed within thresholds.

10.5 Protection of personally identifiable information (PII) content area

10.5.1 Protection of PII component

Description

Personally identifiable information (PII) is information that can be used to identify the PII principal to whom the information relates or might be directly or indirectly linked to a PII principal. The PII principal is the person to whom the PII relates. The PII controller determines the purposes and the means for processing PII, while a PII processor processes PII on behalf of and in accordance with the instructions of a PII controller. (See ISO/IEC 29100 and ISO/IEC 27018 for an explanation of the concepts described here.)

Loss, unauthorized disclosure or unauthorized alteration of PII can be damaging to the persons concerned. It is common for laws or regulations to apply to the handling of PII, in addition to the reputational or business damage that can result from the mishandling of PII. Generally, the burden of ensuring the correct handling of PII falls principally on the PII controller.

In the case of cloud computing, it is generally the case that the PII controller is a CSC while the PII processor is a CSP.

As mentioned in [Clause 6](#), some of the components described below may be represented in documents other than the cloud SLA document that make up the CSA, e.g. a privacy policy.

For details, refer to ISO/IEC 19086-4 for cloud SLAs and protection of PII.

Relevance

Principles for the protection of PII are set out in ISO/IEC 29100 and their application for public cloud services is described in ISO/IEC 27018. Implementation of these principles, especially in the context of cloud services, depends on jurisdiction-specific legislation or regulations and also on terms included in the CSA and in the cloud SLA.

There are national and international standards available that may be used by the CSP and/or the CSC to derive terms for the CSA and the cloud SLA. Examples include ISO/IEC 29151 and ISO/IEC 27018 (both based on ISO/IEC 27001), JIS Q15001, NIST/SP 800-53:2013 Appendix J and BS 10012. Some of the standards mentioned above provide vocabulary and advice about the specific PII issues that could be

included as part of the CSA and cloud SLA. See also ISO/IEC 19086-4 that discusses certifications and compliance.

Cloud service level objectives

Specific SLOs relating to the protection of PII component are described in ISO/IEC 19086-4.

10.5.1.1 Cloud service qualitative objectives

Specific SQOs relating to the protection of PII component are described in ISO/IEC 19086-4.

10.6 Information Security content area

10.6.1 Information Security component

10.6.1.1 Description

The information security content area deals with SLOs and SQOs that relate to information security for cloud services. ISO/IEC 27017 and ISO/IEC 27018 address information security for cloud services and data protection for cloud services respectively, and these standards are in turn based on the sets of security objectives and security controls contained in ISO/IEC 27002.

Information security is so general and so pervasive that it touches on many of the cloud SLA components for cloud services, with the result that many SLOs and SQOs that relate to information security are distributed across the components described in this document. As an example, availability is a key information security concern and there is a dedicated component clause for availability in this document.

As mentioned in [Clause 6](#), some of the components described below may be represented in documents other than the cloud SLA document that make up the CSA, e.g. a security policy.

For details, refer to ISO/IEC 19086-4 for cloud SLAs and information security.

10.6.1.2 Relevance

Information security is a key concern for CSCs and the split of responsibilities between the CSC and the CSP puts a focus on the cloud SLA to ensure that all the necessary elements of information security are dealt with appropriately for the covered cloud services.

The information security component deals with those SLOs and SQOs relating to security that are not explicitly dealt with by other components. Other components include the following:

- Service monitoring ([9.4](#));
- Roles and responsibilities ([9.5](#));
- Availability ([10.3](#));
- Protection of PII ([10.5](#));
- Termination of service ([10.7](#));
- Support ([10.8](#));
- Service reliability ([10.11](#));
- Data backup and restore ([10.11.3](#));
- Data management ([10.12](#));
- Attestations, Certifications, Audits ([10.13](#)).

The SLOs and SQOs for information security are described in ISO/IEC 19086-4.

10.7 Termination of service content area

10.7.1 Termination of service component

Description

The termination of service component deals with the exit process, where the use of a cloud service is terminated and there is an orderly process by which the CSC stops using the cloud service.

The exit process involves reversibility, which is the part of the process where the CSC is able to retrieve their cloud service customer data and application artifacts and where the CSP then deletes all cloud service customer data, as well as contractually specified cloud service-derived data after an agreed period. The CSC's expectation is that the CSP will not retain any materials belonging to the CSC after an agreed period.

Various factors can control the process, including requirements for data retention by law, as well as provision for a delay between the user's cessation of use of the service and the actual removal or deletion of data.

The CSC or the CSP can decide to terminate the cloud service contract and bring the CSC's use of the cloud services to an end. While each cloud services contract is unique, the CSA or other governing documents need to address specific termination issues including the exit process and the handling of all classes of data related to the cloud service. (See [10.12.7](#) for Data Portability component.) The exit process includes both business and technical aspects.

The business aspects of the exit process include items such as a notice period, payment of current charges and other payments due such as early termination fees or outbound data transfer fees. Technical aspects of the exit process include retrieval and eventual deletion of the cloud service customer data and the data formats and methods supported for data retrieval.

Relevance

Termination of a cloud services agreement can occur in several scenarios including the following.

- The CSC elects to stop using the service for reasons such as a) failure of the CSP to meet SLOs, SQOs and other parameters, b) a better price is available from another CSP and c) completion of a specific project using the cloud service.

In this scenario, the CSC ordinarily is allowed to retrieve their cloud service customer data from the CSP (data reversibility) provided that the CSC takes responsibility for payment of all charges due for data storage and data transfer. The CSP will delete the cloud service customer data or complete a data sanitization process from the CSP's systems within a specified period of time. However, it is the responsibility of the CSC to ensure data deletion by taking affirmative steps to delete and sanitize its data from the CSP's systems after it has been retrieved. Additionally, the CSP may provide the CSC with other classes of data objects and may delete records (such as log entries) relating to the CSC. In some jurisdictions, some CSC-related data may be retained by the CSP as required by law.

- The CSC goes out of business. The CSC may not be able to pay their outstanding charges and the CSP may not be able to contact the CSC at all.

In this scenario, the CSP may elect to maintain the CSC's data and other related classes of data for some period of time before deleting it.

- The CSP terminates the agreement for causes such as the CSC's failure to adhere to the terms of use or other requirements of the service. In this scenario, the cloud service customer data and other data classes are handled much as they are when the CSC initiates the termination.

- The CSP terminates the agreement for causes such as ceasing to provide the cloud service, e.g. planned shutdown of the CSP. In this scenario, the cloud service customer's data and other data classes are handled much as they are when the CSC initiates the termination.
- The CSP goes out of business. The CSP may enter a state of receivership or initiate a winding-up and issue a notice to CSCs that they have a limited period to recover their data.

The exit process may call for backup data and log data to be maintained by the CSP during the exit process, until a completion point is reached. It is the responsibility of the CSC to ensure its business continuity by taking responsibility for payment of current charges and other payments due such as early termination fees or outbound data transfer fees and retrieval of its data prior to termination. Once the completion point is reached, the CSP is usually obliged to delete all the CSC-related data including things such as log records and user identities. In some jurisdictions, some CSC-related data may be retained by the CSP as required by law. In addition, CSCs may allow CSPs to retain specified data.

At the end of the exit process, the CSP should provide the CSC with notification that the process is complete.

10.7.1.1 Cloud service level objectives

Data Retention Period

The period of time that cloud service customer data is retained after a notification of service termination has been issued.

Log Retention Period

The period of time that cloud service customer-related log files are retained after a notification of service termination has been issued.

10.7.1.2 Cloud service qualitative objectives

Notification of Service Termination

A statement of the process for notifying a CSC that their cloud service agreement is being terminated including the notification period.

Return of Assets

A statement stipulating responsibilities of the CSP and the CSC in relation to the ownership, use, return and disposal of data objects and the disposal of physical artifacts containing data objects as part of the service termination process.

For more information about data portability, see [10.12.7](#).

For more information about data deletion, see [10.12.8](#).

10.8 Cloud service support content area

10.8.1 Cloud service support component

10.8.1.1 Description

The cloud service support component includes SLOs and SQOs relating to support options for the covered services that are available to CSCs.

CSCs require service support for situations ranging from daily operations such as account administration, configuration, billing and the answering of "how to" questions to more serious incidents such as service outages, security breaches and disaster recovery. Responsibility for resolving support incidents might lie with the CSP, the CSC or both in collaboration. CSPs may offer "support plans" or

packages that include a number of varying options that might be free or for a fee. Support for cloud services may be offered through a number of methods such as email, telephone, web forms and APIs, online chat, community forums and social media channels.

The CSP can provide that it will respond to a service failure within a certain timeframe and by following a certain process for keeping the CSC informed in a timely way of service failure. It can also provide that it will keep the CSC informed of measures to repair or fix the service failure.

10.8.1.2 Relevance

Compared to legacy IT, cloud computing makes the CSC far removed from physical access to the computing resources making the CSC dependent on the CSP to resolve many support incidents. Large CSCs ordinarily provide their own first tier support with in-house or contracted personnel, while they will interface with the CSP to resolve more complicated or serious incidents. A clear picture of the support infrastructure, processes and expectations is critical to the successful use of cloud services. Typically, service requests raised by the CSC with the CSP are assigned a severity level based on the impact to the CSC's business. Severity levels might be changed up or down after initial handling and assessment of the service request. For example, the levels could be defined as follows, with associated response times and other SLOs and SQOs specific to each level:

- Level 1: Critical business impact;
- Level 2: Significant business impact;
- Level 3: Minimal business impact;
- Level 4: No business impact.

Changes to the level assigned to a service request are typically done via a service request escalation process as defined in the CSA.

10.8.1.3 Cloud service level objectives

The SLOs for the cloud service support component cover technical aspects such response times and failure notifications. There might be separate SLOs for each incident severity level.

Support Hours

The hours of operation for each support plan.

Service Incident Support Hours

The hours during which CSCs may obtain support specifically for service incidents.

Service Incident Notification Time

The time interval in which the CSP will provide a notification of a service incident to specified contacts at the CSC when provided for in the support plan.

Maximum First Support Response Time

The maximum time between a customer reporting an incident and the cloud service provider's initial response to the report.

Maximum Incident Resolution Time

States the maximum time for resolving an incident.

10.8.1.4 Cloud service qualitative objectives

The SQOs for the cloud service support component cover operational aspects such as points of contact, emergencies and escalations, as well as technical aspects such as disaster recovery and incident reporting.

Support Plans

A list of the service support plans available to CSCs, including any support costs.

The following SQOs may be included under Support Plans in the cloud SLA:

Support Methods

Lists the methods the CSC can use to obtain support.

Support Contacts

Lists specific contacts for service support if available under the support plan.

Service Incident Reporting

Lists the options which the CSC may use to report service incidents to the CSP.

Service Incident Notification

Lists the terms and conditions (severity, timeframe, etc.) under which the CSP will disclose the details of a service outage or condition that affects the operation of the service. The term may also define what constitutes a service incident.

The Service Incident Notification may include

- the cause of the incident,
- the steps the CSP is taking to resolve the incident,
- the time at which the CSP expects to have the incident resolved, and
- any workarounds the CSC may employ while the incident is being resolved.

10.9 Governance content area

10.9.1 Governance component

10.9.1.1 Description

The cloud service governance component defines the metrics, processes and standards that enable the CSP, CSC and other roles defined in ISO/IEC 17789 to support and/or comply with their governance requirements (see ISO/IEC 17998 and ISO/IEC 38500).

Governance indicators (usually metrics) help ensure

- appropriate regulations or standards are supported by the CSP, CSC and services they use,
- governance policies are being complied with by stakeholders, and
- appropriate business or governance processes are occurring.

Generally, indicators or metrics for governance are attained and tracked by operational (runtime) management systems owned by the CSP, CSC or other stakeholders. Governance indicators may include other cloud SLA components such as availability, performance, reliability, information security or protection of personally identifiable information.

A CSC has responsibilities with respect to attaining compliance to regulatory standards for the use of cloud services in certain cases. While a CSP may perform certain audits or gain certification(s) in relation to a cloud service, this may not necessarily ensure end-to-end regulatory compliance for CSCs.

10.9.1.2 Relevance

For cloud services, CSPs and CSCs ordinarily have a customized governance regimen that is set by the business to ensure that cloud services, use of cloud services and solutions using cloud services are meeting and will continue to meet business objectives and goals. It is important that metrics and policies are put in place by CSCs and CSPs that can also detect if the cloud SLA and governance policies are no longer effective and need to be re-evaluated and perhaps changed. For cloud services, the CSC may have to rely on the CSP to supply the metrics and other indicators necessary for governance.

For many industries, there are regulations that need to be adhered to by CSPs, CSCs and other roles defined in ISO/IEC 17789. Examples include security, data protection, financial and health care regulations for data. The Governance component is where the regulations, standards and policies supported by the CSP are noted.

For public cloud in particular, CSCs all share the same features and functionality (depending on the configurations they have purchased), and changes to these cloud services can affect many CSCs at once. CSCs need mechanisms through which they can understand and prepare for any announced changes to the covered services.

10.9.1.3 Cloud service qualitative objectives

This component has the following SQOs:

Regulation Adherence

A list of regulations including name, clause and certification number (if applicable) the CSP attests or has been certified to comply with.

Standards Adherence

A list of industry standards including name, clause and certification number (if applicable) the CSP attests or has been certified to comply with.

Policy Adherence

A statement to the stakeholders that the business or governance policies specific to the service are being adhered to on an ongoing basis.

Audit Schedule

A schedule of audits the CSP undertakes using its own or third-party resources including the schedule for each audit.

10.10 Changes to the cloud service features and functionality content area

10.10.1 Changes to the cloud service features and functionality component

10.10.1.1 Description

The CSP may decide to make changes to the features and functionality of the covered service(s) during the term of the CSA. Changes such as security patches and bug fixes are not covered in this subclause. See [7.5](#) for changes to the cloud SLA itself. Changes may be made for a variety of reasons such as the addition of features and functions requested by CSCs or deprecation of features and functionality that have low adoption.

10.10.1.2 Relevance

For public cloud deployment in particular, cloud service customers all share the same features and functionality depending on the configurations they have purchased. Changes can affect many CSCs at once. CSCs should understand and prepare for any announced changes to the covered service(s).

10.10.1.3 Cloud service level objectives**Minimum Service Change Notification Period**

The minimum period of time between the issuance of a service change notification and the implementation of the change.

Minimum Time Before Feature/Function Deprecation

The minimum time period between the initial availability of a feature or function and deprecation of that feature or function.

10.10.1.4 Cloud service qualitative objectives**Service Change Notification Method**

The method(s) by which the CSP will notify CSCs of changes to the features and functionality of the covered service(s).

10.11 Service reliability content area**10.11.1 General**

Cloud service reliability is an important characteristic of cloud computing systems. This is a complex aspect and is broken down into three components: Service Resilience/Fault Tolerance, Customer Data Backup and Restore, and Disaster Recovery.

10.11.2 Service resilience/fault tolerance component**10.11.2.1 Description**

The availability of a cloud service can be impacted by faults, or failure of hardware and software components that underlie the cloud service. Since cloud services are housed in data centres, potential faults can also occur on the facilities and infrastructure side. Fault tolerance can be defined as the ability of the service to remain in operation in the event one or more components fail, while resilience is the ability of a service to recover after a fault occurs.

10.11.2.2 Relevance

Some CSCs may be able to rely on the CSP's availability SLO while others may need additional SLOs and SQOs for resilience and fault tolerance.

10.11.2.3 Cloud service level objectives**Time to Service Recovery**

The Time to Service Recovery is the time elapsed between a cloud service failing and the service returning back to the normal state of operation.

Mean Time to Service Recovery

The average of a series of Time to Service Restoration calculations.

NOTE Mean time to service recovery is inclusive of the hardware-related *Mean time to repair* service level objective, but in a cloud service environment, hardware is typically virtualized and so the relationship between the time to repair hardware and its impact on service availability is not direct.

Maximum Time to Service Recovery

The largest value of a set of Time to Service Restoration calculations over a defined period of time.

Number of Service Failures

The number of service failures in total or over a defined period of time.

10.11.2.4 Cloud service qualitative objectives

Cloud Service Resiliency/Fault Tolerance methods

A statement of the methods employed by the cloud service provider which afford resilience and fault tolerance for the cloud service(s) and a statement of the methods available to the cloud service customer to afford resilience/fault tolerance for their own workloads.

10.11.3 Customer data backup and restore component

10.11.3.1 Description

The CSC data backup and restore component includes SLOs and SQOs such as backup methods, backup retention periods and the number of backup generations.

10.11.3.2 Relevance

CSPs can provide the CSC with mechanisms to manage their own data backup and in some cases the CSP automatically creates mirrored copies of stored data on different equipment and potentially in a different geographic location.

Some cloud service customer data is transient and may be lost in the event of equipment failure if the data is not regularly committed to permanent storage. The CSP can offer a service specifically to address this scenario. Additionally, it can be appropriate to store backups of virtual machines in the event a server malfunctions and fail-over operations do not restore the VM to full functionality. This is especially important when running applications on a single server instance.

For some cloud services, CSCs might not have the ability to efficiently back up their own data. In this case, CSCs can rely on the CSP's plan for data redundancy and backup.

It is important for CSCs to ensure that their own and the CSP's data backup and recovery plans are sufficient. These plans should include the backup methods, backup interval, backup retention period, number of backup generations retained, the Recovery Point Objectives and Recovery Time Objectives for the selected cloud services.

10.11.3.3 Cloud service level objectives

Backup Interval

The period of time between data backups or the number of data backups made in a defined period of time.

Retention Period for Backup Data

The time period the CSP retains data backups.

Number of Backup Generations

The number of backup generations of cloud service customer data retained by the CSP.

Backup Restoration Testing

The number of restoration tests from backups over a specified time period.

10.11.3.4 Cloud service qualitative objectives**Backup Method**

A list of cloud service customer data backup methods available to the CSC or employed by the CSP.

Backup Verification

A list of methods or technologies to verify the integrity of data backups.

Backup Restoration Test Reporting

A statement describing the content and availability of reports on backup restoration testing.

Alternative Methods for data recovery

A list of methods the CSP can undertake to restore cloud service customer data in the event the primary data restoration method is not successful.

Data Backup Storage Location

List of geographical location(s) where the data backups are stored.

10.11.4 Disaster recovery component**10.11.4.1 Description**

The Disaster Recovery Component covers SLOs and SQOs such as the CSP's disaster recovery plan, Recovery Time Objective, and Recovery Point Objective.

10.11.4.2 Relevance

In addition to faults, availability of a cloud service (see [10.3](#)) can also be impacted by disasters, natural, manmade and accidental. A disaster, compared to a fault, usually results in substantial damage or total loss of the underlying components of the service. Since modern organizations depend on IT services, extended outages of a cloud service can threaten the existence of the organization itself.

In the case of a severe disaster, weeks or months may be required to completely restore the components of an individual data centre. The cloud service might be run on more than one data centre with load shifting and fail-over. Fail-over might be done automatically or manually. However, a CSC's applications and data might not necessarily be run or stored in multiple locations. Such capabilities might be provided with additional costs.

10.11.4.3 Cloud service level objectives**Recovery Time Objective (RTO)**

The maximum time period required to bring the cloud service back from an outage to a correct operational state.

RTO time period starts when the CSP agrees to initiate recovery process in response to a disaster declared by the CSP and ends when the CSC can resume production operation in the standby/secondary environment. If the decision to fail over is made during a planned downtime, the RTO extends to include

the time required to complete the planned maintenance activity by CSP. RTO and RPO do not generally apply to CSC customizations that depend on non-standard components or third-party software.

Recovery Point Objective (RPO)

The maximum time period prior to a failure or disaster during which changes to data may be lost as a consequence of recovery.

NOTE RPO does not specify the amount of acceptable data loss, only the acceptable amount of time. In particular, RPO affects data redundancy and backup.

Data changes preceding the failure or disaster by at least this time period are preserved by recovery. Zero is a valid value and is equivalent to a “zero data loss” requirement.

10.11.4.4 Cloud service qualitative objectives

Cloud Service Provider Disaster Recovery Plan

A plan that includes a documented set of procedures adopted by the CSP for restoring the cloud service as well as the CSC’s applications and data. These procedures can be executed automatically or manually.

NOTE The RTO and RPO SLOs can form part of the cloud service provider disaster recovery plan.

10.12 Data management content area

10.12.1 General

The data management components define how a cloud service deals with data.

Understanding this is an important aspect of an informed agreement between the CSC and CSP since the use of a cloud service changes a CSC’s control of their data. While it is helpful for both the CSP and the CSC to come to a common agreement on how to handle data, it is important to recognize that data rights, roles and obligations are subject to relevant laws.

When using cloud services, the CSC’s data is stored and processed in the CSP’s data centres. Additionally, cloud service customer data from multiple tenants can exist side-by-side within the service. The CSP has access to the data, and in some agreements, the ability to use that data as part of their business. In any case, the CSP has access to cloud service customer data and maintains account data including names and billing information for tenants along with data about the operation of the service and the underlying infrastructure by multiple tenants.

Multiple issues relate to cloud service customer data including confidentiality, portability, deletion, retention, regulation, law enforcement access and geographic location.

These issues mean that gaining an understanding of the handling of data in a cloud service as the basis for an informed agreement requires defining key classes of data objects and agreeing on the intellectual property rights, control and use of these classes. CSCs need to understand the classes of data involved when using cloud services and how the CSP defines each class of data, how the different classes of data may be created and processed and the CSP’s policies for each class of data.

Control of data is defined as the authorization to perform search, create, read, update and delete operations on the data. Additional elements of control of data in cloud computing include the ability to move data from one location to another and the ability to route the flow of data as it traverses the CSP’s systems.

Control is frequently shared. For example, a CSC may create a file using a cloud service which is initially stored in a server’s memory and potentially stored temporarily on a disk drive local to the server. The file then may be placed for longer-term storage in another location. Control also changes over time, for example, the CSC has control of the file when it is introduced to the service; during the processing of that data, the CSP also has control by having the authority to select which server the CSC uses and where the file resides for longer-term storage.

For some CSPs, the data management components (10.12.2 through 10.12.11) could be included in CSP's policy documents and not necessarily as elements of the cloud SLA.

Note that ISO/IEC 19944 elaborates on issues related to different classifications of Cloud Service Customer data, and what rights a Cloud Service customer may or may not have over them.

10.12.2 Intellectual property rights (IPR) component

10.12.2.1 Description

"Ownership" of data is a complex combination of intellectual property rights and control, and separate agreement on each of those issues is key to a meaningful overall agreement. The law, regulation and custom for Intellectual Property Rights for data vary with different locations and the assignment of rights is closely tied to the business arrangement between the CSC and CSP, requiring a clear and comprehensive agreement about IPR.

10.12.2.2 Relevance

Similar to IT outsourcing arrangements, placing or creating data on the CSP's systems does not diminish the CSC's intellectual property rights but does create a multi-party relationship regarding the data that is different from on-premises systems where the CSC has sole access to the data.

10.12.2.3 Cloud service qualitative objectives

Intellectual Property Rights

A statement of any IPR the CSP claims on the cloud service customer data. Alternatively or in addition, it is a statement of any IPR the CSP grants to the CSC on the cloud service provider's data and/or cloud service-derived data.

10.12.3 Cloud service customer data component

10.12.3.1 Description

ISO/IEC 17788 defines cloud service customer data as a class of data objects under the control of the CSC. For example, such objects can include files, BLOBs, tables, database entries, emails and other objects created using the cloud service or transferred to the CSP for temporary or long-term processing or storage.

Cloud service customer data includes data input into the cloud service by the CSC and the results of the CSC's use of the cloud service to process that data.

10.12.3.2 Relevance

Cloud services include different classes of data objects, some under the control of the CSC and some under the control of the CSP. In traditional distributed computing models, all of the data objects are under the control of the CSC.

10.12.3.3 Cloud service qualitative objectives

Cloud Service Customer Data

A statement defining the cloud service customer data, such as the CSC files and database content.

Cloud Service Customer Data Usage

A statement of all uses of cloud service customer data by the cloud service provider.