

**FEDERAL AVIATION ADMINISTRATION  
HANDBOOK**

**Preparation of SWIM Cloud Service Level Agreement  
(SLA-C)**

**June 11, 2024**

## Foreword

Handbooks are defined by the FAA as non-mandatory guidance documents. This handbook is for Service Level Agreement (SLA) developers and could be used in developing Cloud SLAs. This handbook describes the elements that can be used to compose a Cloud SLA. For the purposes of this handbook, a Cloud SLA is the SLA between an application on the Cloud and an application user. It assumes that the Cloud technical architecture is being offered as a Platform as a Service (PaaS). This handbook does not describe the SLA between the application and the Cloud Provider, but it does explore the relationship between this SLA and the application/user SLA.

This handbook has been developed in accordance with FAA-STD-069 *Preparation of Handbooks* (December 2009).

Comments, suggestions, or questions on this document should be addressed to System Wide Information Management (SWIM) Governance Operations, 800 Independence Avenue, S.W., Washington DC, 20591

Change Record

Version	Date	Description of Changes
1.0.0	06/11/2024	Initial Draft for CCB review

## Table of Contents

1	Scope.....	1
1.1	Scope .....	1
1.2	Intended Audience .....	1
1.3	Background .....	1
1.4	Document Organization.....	1
2	Applicable References .....	1
2.1	FAA Documents .....	1
2.2	Non-FAA Government Documents .....	2
2.3	Non-Government Documents .....	2
3	Definitions .....	3
3.1	Terms and Definition .....	3
3.2	Document Organization.....	5
4	General Guidance.....	5
4.1	General Description .....	5
5	Detailed Guidance.....	6
5.1	Service Level Management.....	6
5.2	The FAA Cloud Application and the Government Sponsored Cloud Environment (GSCE) .....	8
5.3	Service Level Agreement-Cloud (SLA-C) .....	10
5.3.1	Service Level Targets .....	11
5.3.2	Support and Maintenance.....	12
5.3.3	Terms and Conditions.....	13
	Appendix A – SLA-C SLT and Cloud Provider SLT Relationships.....	14
	Appendix B – SLA Laws, Standards and Policies.....	16
	Appendix C – Examples of a Limitation of Liability and Indemnification Clauses .....	22
	Appendix D - Acronyms.....	23

<b>FIGURE 1: VENN DIAGRAM OF SLM COMPONENTS.....</b>	<b>7</b>
<b>FIGURE 2 RELATIONSHIP BETWEEN ACTORS AND THE SLA.....</b>	<b>7</b>
<b>FIGURE 3: RELATIONSHIP BETWEEN CONTRACTS AND SYSTEM ACTORS.....</b>	<b>8</b>
<b>FIGURE 4: IAAS, PAAS, AND SAAS CLOUD MODELS .....</b>	<b>9</b>
<b>FIGURE 5: RELATIONSHIP BETWEEN THE CLOUD APPLICATION, AWS AND GSCE .....</b>	<b>10</b>
<b>FIGURE 6: RELATIONSHIP BETWEEN ACTORS AND CONTRACTS.....</b>	<b>11</b>

# 1 Scope

## 1.1 Scope

This handbook provides a set of guidelines for the preparation of a SLA between an application on the Cloud and a user of that application (hereafter referred to in this document as SLA-C). It also describes the relationship between the SLA-C measurable targets and the targets established in the SLA between the application and the Cloud Provider (hereafter referred to in this document as an Underpinning Contract (UC)). It also describes the SLA-C non-measurable targets, which include the laws and standards to which the SLA-C must conform. The handbook assumes that the application is located on a Platform as a Service technical architecture.

This handbook is for guidance only and cannot be cited as a requirement.

## 1.2 Intended Audience

This handbook is intended to be used by those governance individuals who will be responsible for the creation and management of a SLA-C. No prior knowledge of Service Level Agreements or of Cloud technology is needed to use this handbook.

## 1.3 Background

The FAA maintains a “Cloud-first” strategy, resulting in FAA systems and applications moving to the Cloud. In addition, new systems and applications are being developed on the Cloud itself. In recent years, the System Wide Information Management (SWIM) program has defined a conceptual vision for establishing and preparing SLAs within SWIM. This vision is outlined in the *Concept of Operations SWIM Service Level Management*, Version 1.0.0, dated October 26, 2022. However, at that time, the subject of SLAs in the Cloud environment was not explored and was not included in the ConOps. This handbook builds on previous works but focuses on realizing the concept of SLAs in SWIM's Cloud environment.

## 1.4 Document Organization

This handbook first gives a list of reference documents and then a glossary. It then provides a brief description of the SLM framework and then discusses the extension of adding the Cloud to this framework. It then describes the elements included in the SLA-C: the Service Level Targets (SLT) and the non-SLT SLA-C elements. The appendices describe the relationship between the SLA-C SLT and the Cloud provider SLT, the laws that must be followed and the legal clauses that must be included in the SLA-C, along with an example of two of these clauses.

# 2 Applicable References

## 2.1 FAA Documents

United States federal Aviation Administration System Wide Information Management, *Concept of Operations SWIM Service Level Management*, Version 1.0.0, October 26, 2022

## 2.2 Non-FAA Government Documents

United States, Government Accountability Office, *Cloud Computing Agencies Need to Incorporate Key Practices to Ensure Effective Performance*, GAO-16-325, April 2016

## 2.3 Non-Government Documents

Wieder, Philipp, et al. *Service Level Agreements for Cloud Computing*. Springer, 2011

Erl, Thomas, Zaigham Mahmood, and Ricardo Puttini. *Cloud Computing Concepts, Technology & Architecture*. Prentice Hall Pearson Education, 2013

Bakalos, Nikos, et al. *Slalom Legal & Open Terms for Cloud SLA and Contracts, SLA specification and reference model-a*, European Commission through the H2020 Programme under Grant Agreement 644720

Baudoin, Claude, et al. *Practical Guide to Cloud Service Agreements Version 2.0*, Cloud Standards Customer Council, 2015

*Nvidia Cloud End User License Agreement, Amazon Web Services User*, May 7, 2019

Buck, Kevin, Diane Hanf and Daniel Harper, *Cloud SLA Considerations for the Government Consumer*, MITRE, July 2015

*Practical Guide to Cloud Service Agreements – Version 3.0*, Object Management Group, February 2019

Simmon, Eric *Cloud Service Level Agreements Meeting Customer and Provider needs*, NIST Software and System Division, January 28<sup>th</sup> 2014,

<https://www.youtube.com/watch?v=2rRqgZaXkdM>

Hartley, Cara. *Legal Requirements for SaaS*, TermsFeed, <https://www.termsfeed.com/blog/legal-requirements-saas/>

Cytowski & Partners, *SAAS LEGAL CHECKLIST*, Medium, November 20,2020, <https://cytlaw.medium.com/saas-legal-checklist-94fd069aba80>

*AWS Customer Agreement*, Amazon Web Services, December 22, 2023, <https://aws.amazon.com/agreement/>

*What's the Difference Between Throughput and Latency?*, Amazon Web Services, <https://aws.amazon.com/compare/the-difference-between-throughput-and-latency/>

Bonta, Rob *California Consumer Privacy Act (CCPA)*, State of California Department of Justice, <https://oag.ca.gov/privacy/ccpa>

*The Virginia Consumer Data Protection Act*, <https://www.oag.state.va.us/consumer-protection/files/tips-and-info/Virginia-Consumer-Data-Protection-Act-Summary-2-2-23.pdf>

Wolford, Ben *What is GDPR, the EU's new data protection law*, GDPR EU, <https://gdpr.eu/what-is-gdpr/>

ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection, Information security management systems Requirements, ISO, <https://www.iso.org/standard/27001>

Wikipedia, The Free Encyclopedia, ISO/IEC 27001, [https://en.wikipedia.org/wiki/ISO/IEC\\_27001](https://en.wikipedia.org/wiki/ISO/IEC_27001)

Wikipedia, The Free Encyclopedia, FedRAMP, <https://en.wikipedia.org/wiki/FedRAMP>

Wikipedia, The Free Encyclopedia, NIST Special Publication 800-53, [https://en.wikipedia.org/wiki/NIST\\_Special\\_Publication\\_800-53](https://en.wikipedia.org/wiki/NIST_Special_Publication_800-53)

A-LIGN, SOC 2 Compliance The Definitive Guide, <https://www.a-lign.com/resources/soc-2-the-definitive-guide>

Barney, Nick, TechTarget, SOC 3 (System and Organization Controls 3), <https://www.techtarget.com/searchsecurity/definition/Soc-3-Service-Organization-Control-3>

Whitaker, Scott, Contract Sent, What is limitation of liability in a SaaS agreement?, July 26, 2023, <https://www.contractsent.com/what-is-limitation-of-liability-in-a-saas-agreement/>

Prakash, Sameer, LinkedIn, Unlocking the Secrets of SaaS agreements: Navigating LOL, Indemnification, and License Clauses, June 29, 2023, <https://www.linkedin.com/pulse/unlocking-secrets-saas-agreements-navigating-lol-license-prakash>

Whitaker,, Scott, Contract Sent, Contract indemnification in SaaS contracts, December 27, 2022 <https://www.contractsent.com/contract-indemnification-in-saas-contracts/>

Google, Concepts in service monitoring, <https://cloud.google.com/stackdriver/docs/solutions/slo-monitoring>

Bhardwaja, Utkarsha, Jesse Gordon, Oracle Linux Blog: When will my instance be ready?— understanding cloud launch time performance metrics, June 15, 2021 <https://blogs.oracle.com/linux/post/when-will-my-instance-be-ready-understanding-cloud-launch-time-performance-metrics>

## 3 Definitions

### 3.1 Terms and Definition

**Amazon Machine Image (AMI):** An Amazon resource that is a pre-configured virtual machine (VM) image that contains the operating system (OS) that stores all the configuration, metadata, permissions, and data from multiple disks of a virtual machine instance.

**Artificial Intelligence (AI):** The theory and development of computer systems able to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.

**Central Processing Unit (CPU):** The primary component of a computer that acts as its control center. It is a complex set of electronic circuitries that runs the machine's operating system and applications.

**Cloud:** Servers that are accessed over the Internet and the software and databases that run on those servers.



**Graphical Processing Unit (GPU):** This is an electronic circuit that can perform mathematical calculations at high speed. It can perform tasks like graphics rendering, machine learning and video editing.

**Horizontal Provisioning:** A way to increase the capacity of a system by adding more Virtual Machines that are identical to the present Virtual Machines.

**Indemnification:** Security against legal liability for one's actions.

**Infrastructure as a Service (IaaS):** The division of the cloud technology architecture where the Cloud user is responsible for the application, application data, runtime, middleware, and the operating system and the Cloud Provider is responsible for the rest of the technology stack.

**Liability:** the state of being legally responsible for something.

**Measurement, Monitoring and Reporting (MMR):** A set of processes that support continuous measuring, monitoring, reporting, and analyzing service performance metrics to ensure they meet the agreed-upon service level targets.

**Middleware:** Software that lies between the operating system and the applications running on that operating system. It functions as a translation layer and enables communication and data management for distributed applications.

**National Airspace System (NAS):** The airspace, navigation facilities, and airports of the United States along with their associated information services, rules, regulations, policies, procedures, personnel, and equipment.

**Operating System (OS):** System software that manages computer hardware and software resources and provides common services for computer programs.

**Operational Level Agreement (OLA):** A documented agreement between a service provider and another unit of the same organization who both participate in delivering the service specified in the SLA.

**Platform as a Service (PaaS):** The division of the cloud technology architecture where the Cloud user is responsible for the application and the application data and the Cloud Provider is responsible for the rest of the technology stack.

**Random Access Memory (RAM):** A computer's short-term memory, where the data that the processor is currently using is stored.

**Runtime:** A computer sub-system where the program is created and run.

**Service Consumer (SC):** An organizational entity that uses the service and maintains a business relationship with the service provider.

**Service Integrator (SI):** An organizational entity responsible for managing and integrating interdependent services from various internal and external service providers into an end-to-end solution that meets business objectives.

**Service Level Agreement (SLA):** A documented agreement between a service provider and consumer that identifies services and their agreed-upon performance.

**Service Level Management (SLM):** A framework by which services are defined, service levels required to support business processes are agreed upon, Service Level Agreements (SLAs) and Operational Level Agreements (OLAs) are developed to satisfy the agreements, and the costs of service are developed.

**Service Level Target (SLT):** A service level that an organization commits to.

**Service Provider (SP):** An organizational entity responsible for provisioning the service for a service consumer.

**Software as a Service (SaaS):** The cloud architecture where the Cloud Provider is responsible for the whole Cloud Technology Stack.

**System Wide Information Management (SWIM):** FAA infrastructure that is responsible for passing information between NAS clients and between Non-NAS and NAS clients.

**Underpinning Contract (UC):** A contract between a service provider and a third party. The third party provides supporting services that enable the service provider to deliver a service to a consumer.

**Vertical Provisioning:** The act of adding additional capabilities (such as processing power, memory or storage) to a single machine to manage higher loads without adding multiple machines to the system.

**Virtual Machine (VM):** A computer system created using software on one physical machine to emulate the functionality of another separate physical computer.

**Virtualization:** Technology used to create virtual representations of servers, storage, networks, and other physical machines. Virtual software mimics the functions of physical hardware to run multiple virtual machines simultaneously on a single physical machine.

## 3.2 Document Organization

This document has a forward that is a quick summary of the handbook, table of contents and figures and then a Scope section that describes the handbook's purpose and the intended audience, a listing of applicable government and non-government documents and websites and then a section where common terms used in this document are defined. It then has two sections that define the Service Level Management model and how the SLA-C is an extension of that model. It then talks about measurable Service Level Targets that belong in a SLA as well as the non-measurable targets and laws that need to be addressed by the SLA-C. It contains four appendices: the first contains acronym definitions used in this document, the second describes the relationship between the SLA-C Service Level Targets and the Cloud Service Level Targets, the third a listing of laws and liabilities that the SLA-C should address and the last examples of Limited Liabilities and Indemnification Clauses.

# 4 General Guidance

## 4.1 General Description

A SLA is an agreement between a provider of services and the consumer of these services. The SLA and other contract documents associated with an SLA are defined in the Service Level Management (SLM) ConOps. This ConOps, using the SWIM system as an example, defines three different types of contracts: the SLA, which is the contract between a provider and a consumer or is the contract between a program and a consumer; the Operational Level Agreement (OLA), which is the contract between the different services; and the Underpinning Contract (UC), which is a contract between the service and a contractor.

At the present time, the SLM ConOps does not apply to the Cloud. The purpose of this handbook is to define a SLA-C as an extension of the Service Level Management Model.

## 5 Detailed Guidance

### 5.1 Service Level Management

The Cloud Service Level Agreement is a component of the Service Level Management framework. The Service Level Management framework defines a solution for establishing and maintaining dependable and sustainable relationships between service providers and consumers in a non-cloud environment.

Figure 1 below is a Venn Diagram that shows the relationship between the SLM and its components. Inside the circle labeled “Agreement” are the three different contracts that are part of the SLM framework: the SLA which is a contract between either the system integrator or service provider and the service consumer; the OLA which is a contract between two service providers or between a service provider and a service integrator; and the UC which is a contract between a contractor and the service provider. Inside the circle labeled **Measurement, Monitoring and Reporting (MMR)** are the four different methods used to ensure that the SLM’s contracts achieve their objectives. These methods are: measuring which is the act of determining whether the Service Target Level is within acceptable bounds; monitoring which is the act of periodic measuring of the Service Target Level; reporting which is the act of revealing the measured value for the monitored Service Level Targets to a user or analysis; and analyzing which is the act of determining if action needs to be taken based on the measured, monitored Service Level Targets. The agreements (contracts) and MMR methods are both inside the SLM circle which indicates that they are part of the Service Level Management framework. The intersection of the agreement and MMR circle is the space where the Cloud Application (CA) is in compliance with the SLM Framework, i.e., using both the contracts and Measuring, Measurement and Reporting methods in daily operations.




Figure 1: Venn Diagram of SLM Components

.Figure 2 shows the relationship between system actors and their associated SLAs. Note the SLA can be between a provider and a consumer or it can be between the program and the consumer as shown in the figure. Also, the provider system is composed of providers and contractors. The yellow box below each actor describes the service provided by the actor.

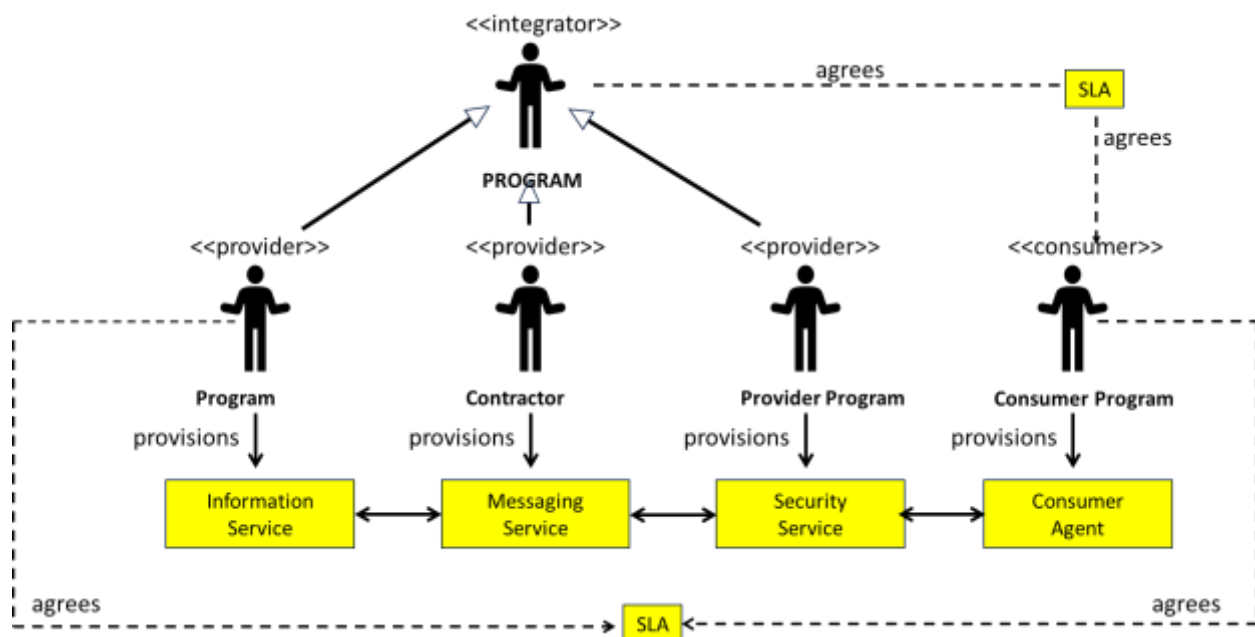
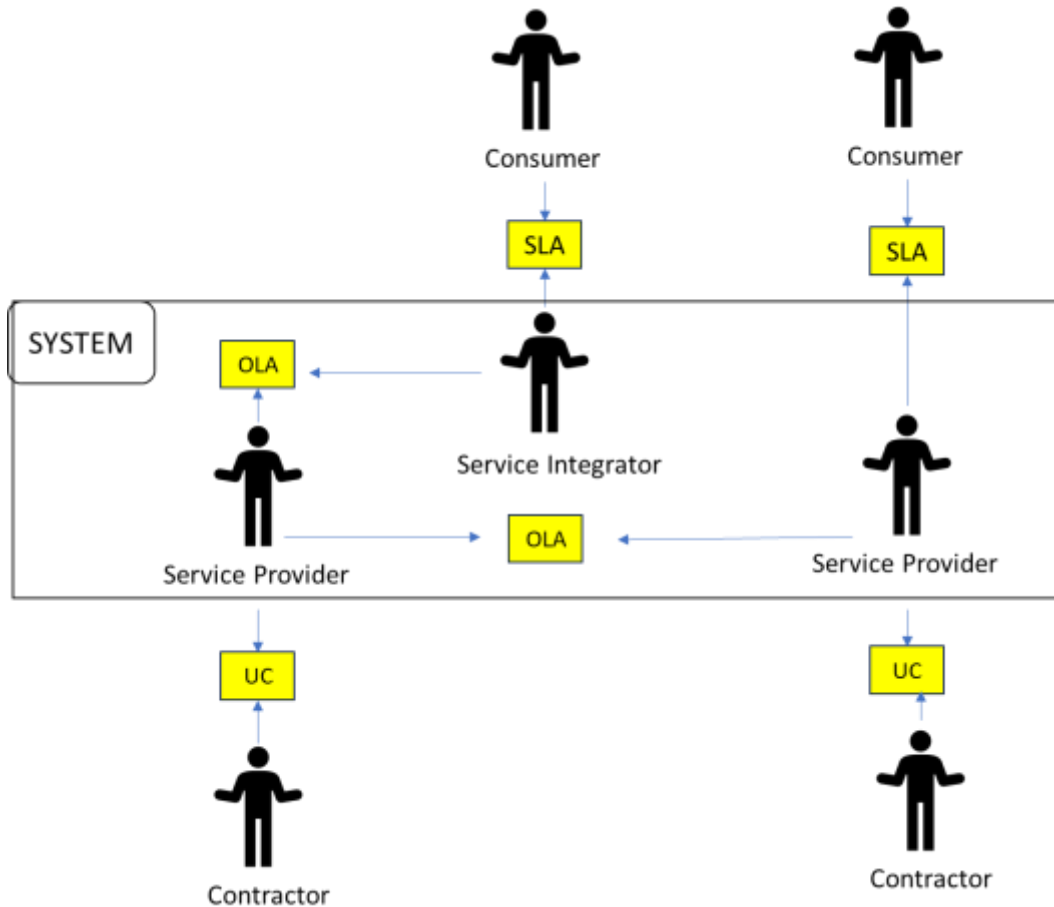


Figure 2 Relationship between Actors and the SLA

Figure 3 shows the relationship of the different contracts among the system actors. An SLA is a contract between either the system integrator or service provider and the service consumer. An OLA is a contract between two service providers or between a service provider and a service integrator. A UC is a contract between a contractor and the service provider. Note that with respect to the system, the SLA is the contract between the consumer and the program, the OLA is a contract between the different actors which make up the system ecosystem and the UC is a contract between a contractor that is providing a service to the system and the system itself

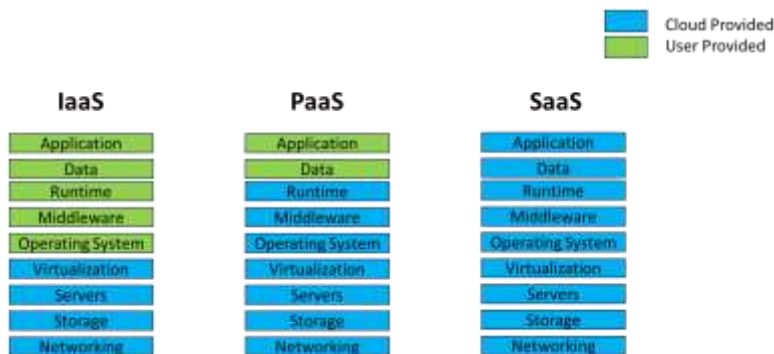


**Figure 3: Relationship Between Contracts and System Actors**

## 5.2 The FAA Cloud Application and the Government Sponsored Cloud Environment (GSCE)

The Cloud application and its data will be hosted on the Government Sponsored Cloud Environment (GSCE). The GSCE is a private cloud that is hosted on the Amazon Web Services (AWS) infrastructure. There are three types of Cloud models, Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). The difference between these models is who has responsibility for which

portions of the Cloud technology stack. Figure 4 shows the division of responsibility for each of these models.




**Figure 4: IaaS, PaaS, and SaaS cloud models**

In this handbook, the Cloud model used is Platform as a Service (PaaS) which means that the program is responsible for the application and the application data on the Cloud technology stack. AWS will be responsible for the following on the Cloud technology stack: Runtime, Middleware, Operating System (O/S), Virtualization, Servers, Storage and Networking.

Another definition of the PaaS model is provided by Wikipedia:

“Platform as a service or application platform as a service or platform-based service is a category of cloud computing services that allows customers to provision, instantiate, run, and manage a modular bundle comprising a computing platform and one or more applications, without the complexity of building and maintaining the infrastructure typically associated with developing and launching the application, and to allow developers to create, develop, and package such software bundles.”

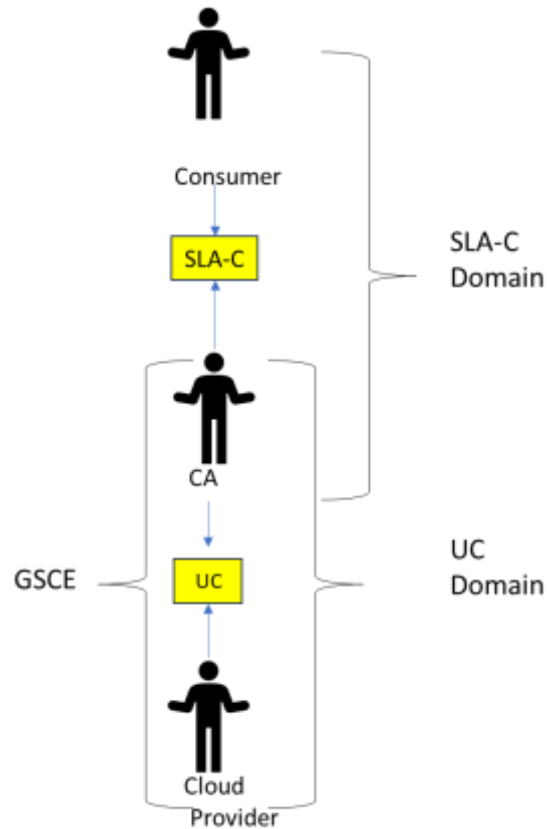
Figure 5 shows the division of the GSCE technology stack (note that the portion of the Cloud Technology Stack that the FAA is responsible for is called the Cloud Application (CA) in the figure).



**Figure 5: Relationship Between the Cloud Application, AWS and GSCE**

### 5.3 Service Level Agreement-Cloud (SLA-C)

Figure 6 shows the relationship between the Consumer and the Cloud Application (CA) and that between the SLA-C and the UC. The SLA-C domain is the service-level targets presented to the consumer derived from the UC service-level targets. The UC domain refers to the Service-level targets associated with the Cloud Provider. Note that the SLA-C Service Level Targets depend on the UC Service Level Targets. This relationship is explored further in Appendix A.



**Figure 6: Relationship between Actors and Contracts**

### 5.3.1 Service Level Targets

The Service Levels Targets used to build an SLA include the following: Availability, Throughput, Latency, Response Time, Mean Time to Restore, Mean Time to Failure, Horizontal Elasticity, Vertical Elasticity, Capacity, Instance Cold Starting Time, and Instance Warm Starting Time as defined below. The SLA-C also contains obligations that are not measurable. These obligations are organized as support elements, data elements, legal obligations and guidelines, and terms and conditions.

The following SLA-C Service Level Targets may be used in the definition of the SLA-C:

**Availability:** The property that the CA is accessible and usable upon demand from an authorized user.

**Throughput:** The minimum number of specific requests that the CA can process in a stated time.

**Latency:** The round-trip time from the moment a request is sent by the customer to the CA to when the response is received by the customer.

**Response Time:** The exact time between the CA receiving a request and the CA's response to the request.

**Incident Response Time:** The time it takes for the CA to respond to a troubleshooting incident.

**Mean Time to Restore:** The average time that it takes the CA to recover from a total failure.



**Mean Time to Failure:** The average time between CA failures.

**Horizontal Elasticity:** The average time needed for the CA to dynamically provision or de-provision horizontal resources as necessary.

**Vertical Elasticity:** The average time needed for the CA to dynamically provision or de-provision vertical resources as necessary.

**Capacity:** The maximum number of CA instances that can run at the same time on the same system.

**Instance Cold Starting Time:** Average length of time required to initialize a new CA instance on a new Virtual Machine.

**Instance Warm Starting Time:** Average length of time required to initialize a new CA instance on a running Virtual Machine.

The SLA-C also contains obligations that are not Service Level Targets (measurable). These obligations are organized as support elements, data elements, legal obligations and guidelines, and terms and conditions.

### 5.3.2 Support and Maintenance

Since the CA is offered as a PaaS, any maintenance of the physical/virtual infrastructure is the responsibility of the Cloud provider. The Cloud Application will still need software updates and security patches, as well as periodic security scans, which could take the CA offline for a period of time. Also, the CA application owners may offer support to help new users to come online with the application as well as support to troubleshoot any problems. The support/maintenance elements included in the SLA-C are:

**Support:** What type of support the CA Provider will supply. For instance, the CA provider may offer support in how to develop software to access/use the CA and it could offer troubleshooting support. These types of support could eliminate connection/use issues as well as detect bugs in the CA itself.

**Support Notifications:** The methods that the CA customer can use to obtain support from the CA Provider.

**Planned Maintenance:** Maintenance schedule whose date and time is fixed. Usually, maintenance is not the result of an unexpected issue, such as a troubleshooting event.

**Planned Maintenance Notification:** The methods used by the CA Provider to notify customers of any Planned Maintenance.

**Information Security:** The information security methods provided by the CA Provider to ensure that the transmitted and stored customer data is kept safe.

**Cybersecurity:** The cybersecurity methods implemented in the CA.

**Incident Reporting:** The list of options that the CA provider (or its customers) can use to report service incidents.

**Pricing and Billing:** Pricing structure, models, and applicable fees

Data elements are important because there are usually legal ramifications for violating or not providing them. The data elements included in the SLA-C are:

**Data Portability:** The ability for the CA provider to transfer data between instances of the CA without having to reenter the data.

**Data Privacy:** The methods used by the CA provider to ensure data privacy.

**Data Location:** The methods used by the CA provider to ensure that data is only located at locations determined by law.

**Data Rights:** The CA provider and the CA consumer shall agree on who owns the data and on any specifications on how the data can be used (e.g. Can the CA Consumer sell the data).

The Legal elements and obligations included in the SLA-C are:

**Laws, Standards and Policies:** The laws, standards and policies that the CA provider must follow.

**Liabilities:** A description of the CA liabilities and liabilities that the CA provider is not responsible for.

**Limitation of Liabilities Clause:** A clause which restricts the amount a party can claim in the event of a breach, error, or overstep during the execution of the CA as a PaaS.

**Indemnification Clause:** The Indemnification Clause defines the responsibilities of each party in case of claims, damages, or losses resulting from the use of the CA.

**License:** A license defines the ways that clients may use the CA and its associated data.

### 5.3.3 Terms and Conditions

The terms and conditions in the SLA-C are not legal requirements but are conventions that need to be established for SLA-C changes and termination events. The terms and conditions included in the SLA-C are:

**Change Management:** The management of changes to the content of the SLA-C.

**Change Management Notification:** The methods used by the CA provider to notify consumers of an SLA-C change.

**Termination:** The conditions under which the CA provider will terminate a customer's usage.

**Termination Notification:** The process used to notify the CA customers that their service is being terminated.

**SLA-C Termination:** The termination of the SLA-C.

**SLA-C Termination Notification:** The method used to notify SLA-C users that the SLA-C is being terminated.

## Appendix A – SLA-C SLT and Cloud Provider SLT Relationships

This appendix describes the relationship between a SLA-C Service Level target and a Cloud Provider Service Level target.

**Availability:** The relationship between the SLA-C availability and the Underlying Cloud Provider availability is that the SLA-C availability is the minimum availability of the virtual machines that compose the CA infrastructure. The Virtual Machine availability is defined as Virtual Machine Run Time/ Virtual Machine Planned Run time. The SLA-C availability will change as VMs are added and subtracted to accommodate the load on the CA.

**Throughput:** CA throughput is related to the processing power associated with the Virtual Machines that make up the CA infrastructure. The VM processing power can be represented by the sum of the number of physical cores per physical machine times the number of physical machines used in CA processing times the average CPU utilization.

**Latency:** CA Latency depends on the latency associated with the latency of the Virtual Network that the CA PaaS uses. CA latency equals the latency associated with the physical path from the customer to the GSCE (location) plus the latency associated with the network protocol (protocol efficiency) plus latency associated with network congestion plus the latency associated with the Cloud virtual network (network infrastructure).

**Response Time:** CA Response time is equal to the maximum response time associated with the VMs that make up the CA PaaS infrastructure. The VM response time equals the average response time for the virtual machine times the Call Percentage times the Calls per request. The average VM response time equals the Average time the VM is active divided by the Consuming API Average response time. The Calls per request equals the Request divided by the consuming API requests.

**Mean Time to Restore:** The CA MTTR is equal to the maximum MTTR among the Virtual Machines that make up the CA PaaS infrastructure.

**Mean Time to Failure:** CA MTTF is equal to the minimum MTTF associated with the Virtual Machines that make up the CA PaaS infrastructure.

**Horizontal Elasticity:** CA Horizontal Elasticity is equal to the maximum time that it takes to spin up or down a virtual machine with the same capacity as the other virtual machine in the CA infrastructure.

**Vertical Elasticity:** CA Vertical Elasticity is equal to the maximum time that it takes to spin up or down a virtual machine with a larger/smaller capacity as the other virtual machines in the CA infrastructure.

**Capacity:** CA Capacity is the amount of server hardware resources required to provide the desired levels of service for a given workload mix for the least cost. It is a function of the Network Capacity, Storage Device Capacity, and the Server Capacity. Network Capacity is the maximum amount of traffic that the network can transfer in a given time. Storage Device Capacity is the maximum number of bytes that can be stored on the device. Server Capacity is the maximum number of independent processes that the server can process simultaneously (it is a function of the number of CPU cores, CPU frequency, RAM size and storage size).

**Instance Cold Starting Time:** CA Instance Cold Starting Time is equal to the time required to acquire all of the resources (e.g., CPUs, memory, disks, networking capabilities) needed for the instance plus time to create the boot volume plus the time to copy over the selected operating system image (for VM instances).

**Instance Warm Starting Time:** CA Instance Warm Starting Time is equal to the time required to acquire all of the resources (e.g., CPUs, memory, disks, networking capabilities) needed for the instance plus time to create the boot volume plus the time to copy over the selected operating system image (for VM instances) on a running VM.

## Appendix B – SLA Laws, Standards and Policies

The SLA-C should contain the following: a Privacy Policy, a Terms and Conditions clause, an End User License agreement, a Subscription Agreement, an Acceptable Use Policy, an Intellectual Property Rights clause, and an Indemnification Provision among other legal agreements. The CA Provider should make sure that these clauses are in compliance with both US and global legislation. Some of the laws that the SLA-C should be in compliance with are the:

- California Consumer Privacy Act of 2018 (CCPA)
- Virginia Consumer Data Protection Act (CDPA)
- New York Shield Act
- General Data Protection Regulation (GDPR) (Europe)
- ISO/IEC 27001 – Establishes an Information Security Management System
- Federal Risk and Authorization Management Program (FedRAMP)
- NIST 800-53
- Service Organization Control 2 and 3 (SOC 2, SOC 3) (Manages customer information)

The California Consumer Privacy Act (CCPA) gives consumers the following rights:

- The right to know about personal information that a business collects about them and how it is used and shared.
- The right to delete personal information collected from them.
- The right to opt-out of the sale or sharing of their personal information.
- The right to non-discrimination for exercising their CCPA rights.
- The right to correct inaccurate personal information that a business has about them.
- The right to limit the use and disclosure of sensitive personal information collected about them.

The Virginia Consumer Data Protection Act (CDPA) defines personal data as any information that is linked or reasonably linkable to a Virginia resident. It does not include any publicly available data or Health Information. The CDPA allows consumers to:

- Confirm if the data controller is processing their personal data.
- Correct inaccuracies in the consumer's personal data collected by the controller.
- Delete personal data provided by or obtained about the consumer.
- Obtain copies of personal data collected by the controller.
- Opt out of the processing of personal data for purposes of targeted advertising, the sale of personal data or further profiling.

The New York SHIELD Act imposes the following requirements on companies that collect information on New York residents:

- It broadens the definition of private information to include account numbers, biometric information, credit/debit card information, access codes, usernames, email addresses, passwords and security questions and answers.

- It broadens the definition of a breach to refer to unauthorized access of computerized data that compromises the security, confidentiality, or integrity of private information.
- It expands the territorial scope that the act applies to so that it includes any person or business that owns or licenses private information of a New York resident that operates in the U.S.
- It imposes new data security requirements to include safeguards to protect the security, confidentiality and integrity of private information, implementation of a security program and the designation of an employee to oversee cybersecurity operations.

The General Data Protection Regulation (GDPR) applies to any organization that targets or collects data related to people in the European Union. It contains the following principals:

- Lawfulness, fairness, and transparency – processing must be lawful, fair, and transparent to the data subject.
- Purpose limitation – organizations must process data for the legitimate purposes specified explicitly to the data subject when organizations collect it.
- Data minimization – Organizations should collect and process only as much data as necessary for the purposes specified.
- Accuracy – Organizations must keep personal data accurate and up to date.
- Storage limitation – organizations may only store personally identifying data for as long as necessary for the specified purpose.
- Integrity and confidentiality – Processing must be done in such a way to ensure appropriate security, integrity, and confidentiality (e.g. by using encryption)
- Accountability – The data controller is responsible for being able to demonstrate GDPR compliance with all of these principles.

ISO/IEC 27001 is an international standard to manage information security. It:

- Defines the requirements that an Information Security Management System (ISMS) must meet.
- Follows the three principals of information security:
  - Confidentiality – Only the right people can access the information held by an organization.
  - Information integrity – Data that the organization uses to pursue its business or keeps safe for others is reliably stored and not erased or damaged.
  - Availability of data – The organization and its clients can access the information whenever it is necessary so that business purposes and customer expectations are satisfied.
- Keeps security controls organized.
- Links business continuity planning, physical security, and information security
- Requires that the organization management:
  - Systematically examine the organization's information security risks, taking account of the threats, vulnerabilities, and impacts.

- Design and implement a coherent and comprehensive suite of information security controls and/or other forms of risk treatment to address those risks that are deemed unacceptable.
- Adopts an overarching management process to ensure that the information security controls continue to meet the organization's information security needs on an ongoing basis.
- States that organizations are required to implement risk management processes to identify potential threats, evaluate their impact and develop appropriate mitigation strategies.

The Federal Risk and Authorization Management Program (FedRAMP):

- Provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services.
- Provides accreditation for cloud services for IaaS, PaaS, and SaaS
- Is FISMA (Federal Information Security Management Act) for the Cloud
- Security baselines are derived from NIST SP 800-53 with a set of control enhancements that pertain to the unique security requirements of cloud computing.

NIST 800-53 is an information security standard that provides a catalog of security and privacy controls for all U.S. federal information systems except those related to national security. Federal Systems Programs select and implement a subset of controls listed in NIST 800-53 Appendix F. NIST 800-53 covers the steps in the Risk Management Framework that address security control selection for federal information systems in accordance with the security requirements in the Federal Information Processing Standard (FIPS) 200. The controls are classified into the following 18 control families:

- Access Control
- Audit and Accountability
- Awareness and Training
- Configuration Management
- Contingency Planning
- Identification and Authentication
- Incident Response
- Maintenance
- Media Protection
- Personnel Security
- Physical and Environmental Protection
- Planning
- Program Management
- Risk Assessment
- Security Assessment and Authorization
- System and Communications Protection
- System and Information Integrity
- System and Services Acquisition

SOC 3 is a cybersecurity audit that:

- Examines controls in place that protect and secure a system or services.
- Is necessary if you process, manage or store data in the Cloud or data centers.
- The report is valid for 12 months.

Soc 3 is a controls audit that:

- Creates a report that outlines the information related to a service organization's internal controls for security, availability, processing integrity, confidentiality, and privacy.
- Provide a high-level overview of an organization's controls and security risks.
- Is applicable for SaaS, cloud computing or data center storage.
- Is a way of proving compliance with industry standards (e.g. GDPR)

The CA owner could be held responsible for the following liabilities:

- Shutdown Time (measured by MTTR)
- System Crashes (measured by MTTF)
- Data Breaches
- Data Corruption
- Data Location (certain locations taboo due to regulations and laws)
- CA Client loss of (due to Shut down or Crash)
  - Profits
  - Revenue
  - Data
  - Data Use

The SLA-C should include a Limitation of Liability clause which restricts the amount a CA client can claim in the event of a breach, error, or overstep during the execution of the CA as a PaaS. The Limitation of Liabilities include the following information:

- Monetary Cap – A cap on monetary claims by a CA Client.
- Exclusions of certain types of damages – Indirect, Consequential, Incidental, Special and Exemplary damages (e.g. lost profits or data breaches).
- Exceptions for Gross Negligence and Willful Misconduct – Limitation of Liabilities does not apply in cases of gross negligence, willful misconduct, or fraud on the part of the CA Provider
- Carve-outs for Certain Obligations – Carve-out for indemnification obligations or confidentiality breaches.

The SLA-C should include an Indemnification Clause which defines the responsibilities of each party in case of claims, damages, or losses resulting from the use of the CA. Examples of indemnification that should be included in the Indemnification Clause are:

- Indemnification from third-party claims arising from the customer's use of the CA.



- Indemnification for any claims arising from a customer's use of the CA (e.g., customer's failure to comply with applicable laws, infringement of intellectual property rights, or unauthorized use of service).
- Indemnification of customer arising from infringement of third-party intellectual properties rights (e.g., arises from CA interaction with third-party).

The SLA-C should include a License that defines the ways that clients may use the CA. The elements of a license should include:

- Non-exclusivity – many customers can use the product.
- Worldwide – any customer can access your service wherever they are sitting.
- Revocable – clarifies that you reserve the right to terminate a particular customer's use of the service.
- Use and Copy – this language clarifies that customers are not infringing your rights by using the service in the way you want them to.
- Prohibited Uses – listing behavior that causes a customer to forfeit his license.

The CA owner and CA clients, assuming that AWS is the Cloud Provider, must abide by the AWS Contract that states that the End User:

- Will consent to the use of content by AWS to provide AWS Services.
- Will not violate the Acceptable Use Policy.
- Will not attempt to reverse engineer, disassemble, or decompile AWS Services and AWS Content or use any other process to derive the source code of any AWS Services and Content.
- Will not access or use AWS Services or Content in a way to avoid incurring fees or exceeding using limits or quotas.
- Will not resell AWS Services or Content.
- Will not hold AWS or any of its associates responsible related to any third-party claim that arises out of:
  - Use of the AWS Services.
  - Breach of this agreement or violation of applicable law by CA owner or CA clients.
  - Dispute between CA owner and CA clients.
- Will not hold AWS to any obligations or liability if CA client uses AWS after being notified to discontinue use.
- Does not have any third-party beneficiary rights.

Nvidia is a company that makes Graphical Processing Units (GPU) which are used as the hardware for different applications of Artificial Intelligence (AI) (e.g. Deep Learning). Nvidia has a Cloud End User License Agreement for an AWS User that applies to the CA owner and the CA client. The License states the End User:

- Has the right to access and use the Nvidia Software as part of an Amazon Machine Image (AMI) or for use with an AMI that runs on NVIDIA GPUs.

- May not reverse engineer, decompile, disassemble, or separate parts of the Software as packaged.
- May not modify or create derivative works of the Software.
- May not remove copyright or other proprietary notices from any portion of the Software or copies of Software.
- Except for running or developing AMIs you may not use, distribute, or provision the Software or authorize others to do so.
- May not bypass, disable, or circumvent any encryption, security, digital rights management, or authentication mechanism in the Software.
- May not use the Software in any manner that would cause it to become subject to an open software license.
- Unless you have an agreement with NVIDIA for this purpose, you may not use the Software with any system or application where the use or failure of the system or application can reasonably be expected to threaten or result in personal injury, death, or catastrophic loss. Examples include use in nuclear, avionics, navigation, military, medical, life support or other life-critical applications. NVIDIA does not design, test, or manufacture the Software for these critical uses and NVIDIA shall not be liable to you or any third party, in whole or in part, for any claims or damages arising for such uses.

## Appendix C – Examples of a Limitation of Liability and Indemnification Clauses

### An example Limitation of Liability Clause:

“To the fullest extent permitted by applicable law, in no event shall the DocuSign parties be liable to you for any direct, indirect, incidental, special, punitive, or consequential damages whatsoever resulting from any (I) use of the site; (II) errors, mistakes or inaccuracies of content, (III) personal injury or property damage, of any nature whatsoever, resulting from your access to and use of the site, (IV) any unauthorized access to or use of our servers and/or any and all personal information and/or financial information stored therein, (V) any interruption or cessation of transmission to or from our servers, (VI) any bugs, viruses, trojan horses, or the like which may be transmitted to or through the site by any third party, (VII) any loss of your data or content from the site, (VIII) any errors or omissions in any content or for any loss or damage of any kind incurred as a result of your use of any content posted, transmitted or otherwise made available via the site, whether based on warranty, contract, tort or any other legal theory and whether or not the DocuSign parties are advised of the possibility of such damages, and/or (IX) the disclosure of information pursuant to these TOU or our privacy policy despite the above, or total liability to you for any cause of action you take against us will at all times be limited to no more than one hundred dollars (\$100.00).”

### An example of an Indemnification Clause:

“Provider shall indemnify, defend, and hold harmless Client and its affiliates, Officers, directors, employees, and agents from and against any claims, demands, losses, damages, judgments, settlements, costs, and expenses (including reasonable attorneys’ fees) that arise from or relate to any breach of this Agreement by Provider or any errors or malfunctions in the Software. Provider shall have no liability for any claims, demands, losses, damages, judgments, settlements, costs, and expenses that arise from or relate to the unauthorized use of the Software, or from any third-party claims against Client related to the use of the Software.”

## Appendix D- Acronyms

AI: Artificial Intelligence

AMI: Amazon Machine Image

AWS: Amazon Web Services

CCPA: California Consumer Privacy Act

CDPA: Virginia Consumer Data Protection Act

CPU: Central Processing Unit

SLA-C: Cloud Service Level Agreement

FedRAMP: Federal Risk and Authorization Management Program

FIPS: Federal Information Processing Standard

FISMA: Federal Information Security Management Act

GDPR: General Data Protection Regulation

GPU: Graphical Processing Unit

GSCE: Government Sponsored Cloud Environment

IEC: International Electrotechnical Commission

ISMS: Information Security Management System

ISO: International Organization for Standardization

MMR: Measurement, Monitoring and Reporting

MMTF: Mean Time to Failure

MTTR: Mean Time to Repair

NAS: National Airspace System

NIST: National Institute of Standards and Technology

OLA: Operational Level Agreement

O/S: Operating System

PaaS: Platform as a Service

RAM: Random Access Memory

SC: Service Consumer

SHIELD: Stop Hacks and Improve electronic Data Security Act

SI: Service Integrator

SLA: Service Level Agreement

SLM: Service Level Management

SLT: Service Level Target

SOC: Service Organization Control

SP: Service Provider

SWIM: System Wide Information Management

UC: Underpinning Contract

US: United States

VM: Virtual Machine