GxP Systems on AWS

July 2026

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Abstract

This whitepaper provides information on how AWS approaches GxP-related compliance and security and provides customers guidance on using AWS products in the context of GxP. The content has been developed based on experience with and feedback from AWS pharmaceutical and medical device customers, in addition to software partners, who are currently using AWS products in their validated GxP systems.

Introduction

Prioritization of cloud technologies in the life sciences sector is steadily increasing as customers seek out highly reliable, scalable, and secure solutions to operate their regulated IT systems. Amazon Web Services (AWS) provides cloud services designed to help customers run their most sensitive workloads in the cloud, including the computerized systems that support Good Manufacturing Practice, Good Laboratory Practice, and Good Clinical Practice (GxP). GxP guidelines are established by the US Food and Drug Administration (FDA) and exist to ensure safe development and manufacturing of medical devices, pharmaceuticals, biologics, and other food and medical product industries.

Since the last version of this whitepaper, there have been many changes in the industry, including revised guidance from the FDA on computer systems assurance which clarified many of the misconceptions around computer systems validation, and from ISPE with the second edition of their GAMP 5 Guide: A Risk-Based Approach to Compliant GxP Computerized Systems, which acknowledged the increased importance of cloud services and the increased use of tools and automation to achieve greater control, higher quality, and lower risk. There is also increased acknowledgement of the value of cloud services from regulatory agencies like the FDA as we discuss in Policymakers: Cloud can Accelerate Life Sciences Innovation. There is even clearer guidance emerging about the use of cloud technologies as part of a medical device though the upcoming AAMI TIR 115.

The first section of this whitepaper assists customers with a supplier assessment by outlining the AWS services and the organizational approach to security and compliance that support GxP requirements as part of the Shared Responsibility Model and AWS Quality System. After establishing this information, the whitepaper provides information to assist you in using AWS services to implement GxP-compliant environments, that is, to assist with your side of the shared responsibility model. Many customers already use industry guidance to influence their regulatory interpretation of GxP requirements. Therefore, the primary industry guidances used to form the basis of this whitepaper are the GAMP (Good Automated Manufacturing Practice) guides from ISPE (International Society for Pharmaceutical Engineering). This paper is in effect a type of good cloud computing practice guide.

While the following content provides information on use of AWS services in GxP environments, you should ultimately consult with your own counsel to ensure that your GxP policies and procedures satisfy regulatory compliance requirements.

Whitepapers containing more specific information about AWS products, privacy, and data protection considerations are available at AWS Compliance.

Assessing AWS as a supplier

This section of the whitepaper provides information and insights into AWS to assist with a supplier assessment.

About AWS

In 2006, AWS began offering on-demand IT infrastructure services to businesses in the form of web services with pay-as-you-go pricing. Today, AWS provides highly reliable, scalable, low-cost cloud that power hundreds of thousands of businesses in countries around the world. Using AWS, businesses no longer need to plan for and procure servers and other IT infrastructure weeks or months in advance. Instead, they can instantly spin up hundreds or thousands of servers in minutes and deliver results faster. Offering over 200 fully featured services from data centers globally, AWS gives you the ability to take advantage of a broad set of global cloud-based products including compute, storage, databases, networking, security, analytics, machine learning, generative AI, mobile, developer tools, management tools, Internet of Things (IoT), media services, and enterprise applications. The rapid pace of innovation by AWS allows you to focus on what's most important to you and your end users without the undifferentiated heavy lifting.

AWS Healthcare and Life Sciences

AWS started its dedicated Healthcare and Life Sciences (HCLS) Practice in 2014 in response to the growing demand for an experienced and reliable life sciences cloud industry leader. Today, the AWS Healthcare and Life Sciences Practice team consists of members that have been in the industry on average for over 20 years and had previous titles such as Chief Medical Officer, Chief Digital Officer, Physician, Radiologist, and Researcher among many others. The AWS Healthcare and Life Sciences practice serves a large ecosystem of HCLS customers, including pharmaceutical, biotechnology, medical device, genomics, health tech, start-ups, university and government institutions, in addition to healthcare payers and providers. A full list of customer case studies can be found at Healthcare & Life Sciences Case Studies.

In addition to the resources available within the HCLS practice at AWS, you can also work with AWS Life Sciences Competency Partners to drive innovation and improve efficiency across the life sciences value chain including cost-effective storage and compute capabilities, advanced analytics, and patient personalization mechanisms. AWS Life Sciences Competency Partners have demonstrated technical expertise and customer success in building life science solutions on AWS. A full list of AWS Life Sciences Competency Partners can be found at AWS Life Sciences Competency Partners.

AWS services

AWS delivers a scalable cloud computing platform with high availability and dependability, providing the tools that enable you to run a wide range of applications. Helping to protect the confidentiality, integrity, and availability of our customers' systems and data is of the utmost importance to AWS, as is maintaining customer trust and confidence.

Similar to other general-purpose IT products such as operating systems and database engines, AWS offers commercial off-the-shelf (COTS) IT services according to IT quality and security standards such as ISO, NIST, SOC and many others. For the purposes of this paper, we will use the definition of COTS in accordance with the definition established by FedRAMP, a United States government-wide program for procurement and security assessment. FedRAMP references the US Federal Acquisition Regulation (FAR) for its definition of COTS, which outlines COTS items as:

- Products or services that are offered and sold competitively in substantial quantities in the commercial marketplace based on an established catalog.
- Offered without modification or customization.
- Offered under standard commercial terms and conditions.

Under GAMP guidelines (such as GAMP 5: A Risk-Based Approach to Compliant GxP Computerized Systems), organizations implementing GxP-compliant environments will need to categorize AWS services using respective GAMP software and hardware categories (for example, Software Category 1 for infrastructure software, including operating systems, database managers, and security software or Category 5 for custom or bespoke software). Most often, organizations using AWS services for validated applications will categorize them under Software Category 1.

AWS offers products falling into several categories. The following is a subset of those AWS offerings spanning Compute, Storage, Database, Networking and Content Delivery, Security and

Compliance, and Machine Learning. A later section of this whitepaper, AWS Products in GxP Systems, will provide information to assist you in using AWS services to implement your GxP-compliant environments.

Table 1: Subset of AWS offerings by group

Group	AWS products
Compute	Amazon Elastic Compute Cloud (Amazon EC2), Amazon EC2 Auto Scaling, Amazon Elastic Container Registry (Amazon ECR), Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), Amazon Lightsail, AWS Batch, AWS Elastic Beanstalk, AWS Fargate, AWS Lambda, AWS Outposts, AWS Serverless Application Repository (AWS SAM), AWS Wavelength, AWS App Runner
Storage	Amazon Simple Storage Service (Amazon S3), Amazon Elastic Block Store (Amazon EBS), Amazon Elastic File System (Amazon EFS), Amazon FSx for Lustre, Amazon FSx for Windows File Server, Amazon S3 Glacier, AWS Backup, AWS Snow Family, AWS Elastic Disaster Recovery, AWS Storage Gateway, CloudEndure Disaster Recovery
Database	Amazon Aurora, Amazon DynamoDB, Amazon DocumentDB, Amazon ElastiCache, Amazon Keyspaces, Amazon Neptune, Amazon Relational Database Service (Amazon RDS), Amazon Redshift, Amazon Timestream, AWS Database Migration Service (AWS DMS)
Networking and Content Delivery	Amazon VPC, Amazon API Gateway, Amazon CloudFront, Amazon Route 53, AWS PrivateLink, AWS App Mesh, AWS Cloud Map, AWS Direct Connect, AWS Global Accelerator, AWS Transit Gateway, Elastic Load Balancing, AWS Private 5G
Security, Identity, and Compliance	AWS Identity and Access Management (IAM), Amazon Cognito, Amazon Detective, Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Artifact, AWS Certificate Manager (ACM), AWS CloudHSM, AWS Directory Service, AWS Firewall Manager, AWS Key Management Service (AWS KMS), AWS Resource Access Manager (AWS RAM), AWS Secrets Manager, AWS Security Hub, AWS Shield, AWS IAM Identity Center (IAM), AWS WAF, Amazon Verified Permissions, Amazon Security Lake
Machine Learning	Amazon SageMaker, Amazon Augmented AI, Amazon CodeGuru, Amazon Comprehend, Amazon Personalize, Amazon Rekognition, Amazon Transcribe, Amazon Kendra, Amazon Lex, Amazon Bedrock, Amazon Q

In addition to these offerings, AWS also empowers the HCLS industry by providing purpose-built health services and solutions. Some of the AWS health care and life sciences services are listed in the following table.

Table 2: AWS health and life sciences services

AWS product	Brief description
AWS HealthLake	Provide a complete view of individual or patient population health data using the HealthLake console
AWS HealthImaging	Store, transform, and analyze medical images in the cloud at petabyte scale.
AWS HealthOmics	Transform genomic, transcriptomic, and other omics data into insights.
Amazon Transcribe Medical	Automatically convert medical speech to text.
Amazon Comprehend Medical	Extract information from unstructured medical text accurately and quickly.

Details and specifications for the full portfolio of AWS products are available online at https://aws.amazon.com/.

AWS cloud security

AWS infrastructure has been architected to be one of the most flexible and secure cloud computing environments available today. It's designed to provide an extremely scalable, highly reliable platform that enables customers to deploy applications and data quickly and securely. This infrastructure is built and managed not only according to security best practices and standards, but also with the unique needs of the cloud in mind. AWS uses redundant and layered controls, continuous validation and testing, and a substantial amount of automation to ensure that the underlying infrastructure is continuously monitored and protected.

We have many customer testimonials that highlight the security benefits of using the AWS cloud, in that the security capabilities provided by AWS far exceed the customer's own on-premises capabilities.

> "We had heard urban legends about 'security issues in the cloud,' but the more we looked into AWS, the more it was obvious to us that AWS is a secure environment and we would be able to use it with peace of mind."
>
> **– Yoshihiro Moriya, Certified Information System Auditor at Hoya**

> "There was no way we could achieve the security certification levels that AWS has. We have great confidence in the logical separation of customers in the AWS Cloud, particularly through Amazon VPC, which allows us to customize our virtual networking environment to meet our specific requirements."
>
> **– Michael Lockhart, IT Infrastructure Manager at The GPT Group**

> "When you're in telehealth and you touch protected health information, security is paramount. AWS is absolutely critical to do what we do today. Security and compliance are table stakes. If you don't have those, the rest doesn't matter."
>
> **– Cory Costley, Chief Product Officer, Avizia**

Many more customer testimonials, including those from health and life science companies, can be found in Customer Success Stories.

IT security is often not the core business of our customers. IT departments operate on limited budgets and do a good job of securing their data centers and software given limited resources. In the case of AWS, security is foundational to our core business and so significant resources are applied to ensuring the security of the cloud and helping our customers ensure security in the cloud, as described further in the following section.

Shared Responsibility Model

Security and Compliance is a shared responsibility between AWS and the customer. This shared model can help relieve your operational burden because AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. Customers assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, and the configuration of the AWS-provided security group firewall. You should carefully consider the services you choose because your responsibilities vary depending on the services used, the integration of those services into your IT environment, and applicable laws and regulations.

The following figure provides an overview of the Shared Responsibility Model. This differentiation of responsibility is commonly referred to as Security *of* the Cloud compared to Security *in* the Cloud, which will be explained in more detail in the following paragraphs.

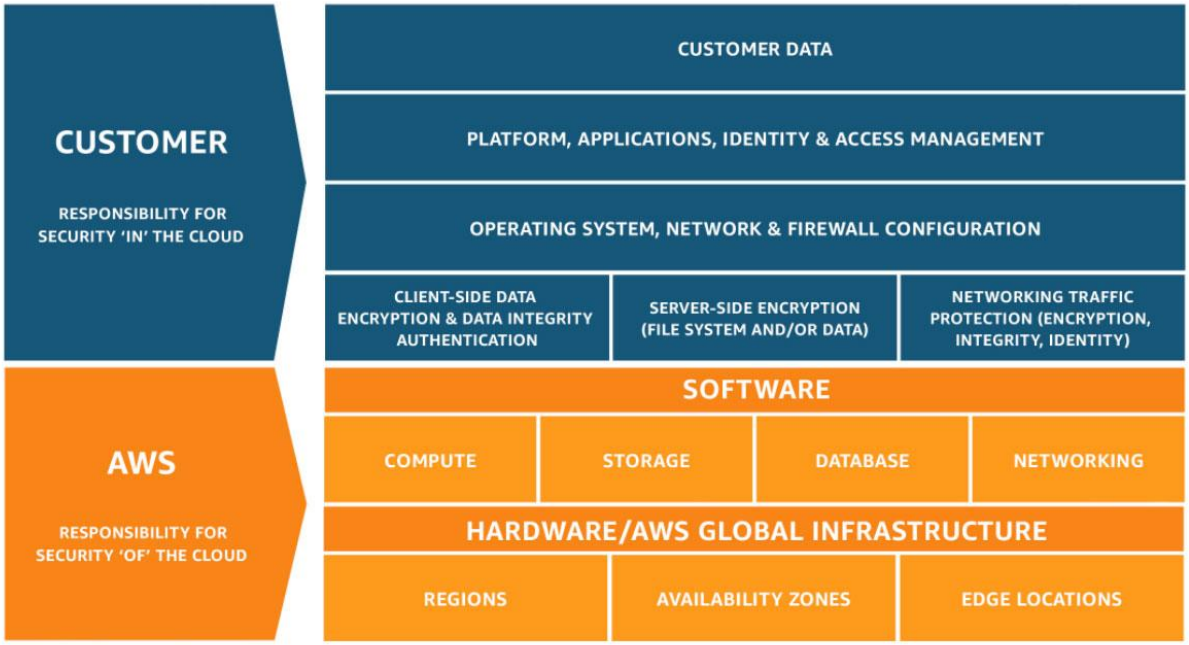

Figure 1: AWS Shared Responsibility Model

AWS is responsible for the security and compliance *of* the cloud, the infrastructure that runs all the services offered in the AWS Cloud. Cloud security at AWS is the highest priority. AWS customers benefit from a data center and network architecture that are built to meet the

requirements of the most security-sensitive organizations. This infrastructure consists of the hardware, software, networking, and facilities that run AWS cloud services.

Customers are responsible for the security and compliance *in* the cloud, which consists of customer-configured systems and services provisioned on AWS. Responsibility within the AWS Cloud is determined by the AWS cloud services that you select and ultimately the amount of configuration work you must perform as part of your security responsibilities. For example, a service such as Amazon Elastic Compute Cloud (Amazon EC2) is categorized as infrastructure as a service (IaaS) and, as such, requires you to perform all the necessary security configuration and management tasks. Customers that deploy an Amazon EC2 instance are responsible for management of the guest operating system (including updates and security patches), any application software or utilities installed by you on the instances, and the configuration of the AWS-provided firewall (called a security group) on each instance. For abstracted services, such as Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB, AWS operates the infrastructure layer, the operating system, and platforms, while customers access the endpoints to store and retrieve data. You're responsible for managing your data and component configuration (including encryption options), classifying your assets, and using AWS Identity and Access Management (IAM) tools to apply the appropriate permissions.

The customer and AWS Shared Responsibility model also extends to IT controls. The same as the responsibility to operate the IT environment is shared between you and AWS, so is the management, operation, and verification of IT controls shared. AWS can help relieve your burden of operating controls by managing those controls associated with the physical infrastructure deployed in the AWS environment that might previously have been managed by you. Because every customer is deployed differently in AWS, you can take advantage of shifting management of certain IT controls to AWS which results in a (new) distributed control environment. You can then use the AWS control and compliance documentation available to you, in addition to techniques discussed later in this whitepaper, to perform your control evaluation and verification procedures as required. The following are examples of controls that are managed by AWS, AWS customers, or both.

Inherited controls – Controls that you fully inherit from AWS.

- Physical and environmental controls

Shared controls – Controls that apply to both the infrastructure layer and customer layers, but in completely separate contexts or perspectives. In a shared control, AWS provides the requirements for the infrastructure and you must provide your own control implementation within your use of AWS services. Examples include:

- **Patch management** – AWS is responsible for patching and fixing flaws within the infrastructure, but you are responsible for patching your guest OS and applications.
- **Configuration management** – AWS maintains the configuration of its infrastructure devices, but you are responsible for configuring your own guest operating systems, databases, and applications.
- **Awareness and training** – AWS trains AWS employees, but you must train your own employees.

Customer specific controls – Controls that are ultimately your responsibility based on the application you're deploying within AWS services. Examples include:

- **Data management** – For instance, placement of data on Amazon S3 where you activate encryption.

While certain controls are customer specific, AWS strives to provide you with the tools and resources to make implementation easier.

For further information about AWS physical and operational security processes for the network and server infrastructure under the management of AWS see AWS Cloud Security.

For customers who are designing the security infrastructure and configuration for applications running in AWS, see the Best Practices for Security, Identity, & Compliance.

AWS certifications and attestations

The AWS global infrastructure is designed and managed according to security best practices in addition to a variety of security compliance standards. With AWS, you can be assured that you're building web architectures on top of some of the most secure computing infrastructure in the world. The IT infrastructure that AWS provides is designed and managed in alignment with security best practices and a variety of IT security standards including the following that life science customers might find most relevant:

- SOC 1, 2, 3
- ISO 9001 / ISO 27001 / ISO 27017 / ISO 27018
- HITRUST
- FedRAMP
- CSA Security, Trust & Assurance Registry (STAR)

- National Institute of Standards and Technology (NIST)

There are no specific certifications for GxP compliance for cloud services to date, however the controls and guidance described in this whitepaper, in conjunction with additional resources supplied by AWS, provide information on AWS services GxP-compatibility, which will assist you in designing and building your own GxP-compliant solutions.

AWS provides on-demand access to security and compliance reports and select online agreements through AWS Artifact, with reports accessible through AWS customer accounts under NDA. AWS Artifact is a central resource for compliance related information and is where you can find additional information on the AWS compliance programs described in this section.

SOC 1, 2, and 3

AWS System and Organization Controls (SOC) Reports are independent third-party examination reports that demonstrate how AWS achieves key compliance controls and objectives. The purpose of these reports is to help you and your auditors understand the AWS controls established to support operations and compliance.

The SOC 1 reports are designed to focus on controls in a service organization that are likely to be relevant to an audit of a user entity's financial statements. The AWS SOC 1 report is designed to cover specific key controls that are likely to be required during a financial audit, in addition to covering a broad range of IT general controls to accommodate a wide range of usage and audit scenarios. The AWS SOC1 control objectives include security organization, employee user access, logical security, secure data handling, physical security and environmental protection, change management, data integrity, availability and redundancy, and incident handling.

The SOC 2 report is an attestation report that expands the evaluation of controls to the criteria set forth by the American Institute of Certified Public Accountants (AICPA) Trust Services Principles. These principles define leading practice controls relevant to security, availability, processing integrity, confidentiality, and privacy applicable to service organizations such as AWS. The AWS SOC 2 is an evaluation of the design and operating effectiveness of controls that meet the criteria for the security and availability principles set forth in the AICPA's Trust Services Principles criteria. This report provides additional transparency into AWS security and availability based on a pre-defined industry standard of leading practices and further demonstrates the commitment of AWS to protecting customer data. The SOC2 report includes outlining AWS controls, a description of AWS services relevant to security, availability and

confidentiality, and test results against controls. You will probably find the SOC 2 report to be the most detailed and relevant SOC report as it relates to GxP compliance.

AWS also publishes a SOC 3 report. The SOC 3 report is a publicly available summary of the AWS SOC 2 report. The report includes the external auditor's assessment of the operation of controls (based on the AICPA's Security Trust Principles included in the SOC 2 report), the assertion from AWS management regarding the effectiveness of controls, and an overview of AWS infrastructure and services.

FedRAMP

The Federal Risk and Authorization Management Program (FedRAMP) is a US government-wide program that delivers a standard approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMP uses the NIST Special Publication 800 series and requires cloud service providers to receive an independent security assessment conducted by a third-party assessment organization (3PAO) to ensure that authorizations are compliant with the Federal Information Security Management Act (FISMA).

For AWS services in scope for FedRAMP assessment and authorization, see AWS Services in Scope by Compliance Program.

ISO 9001

ISO 9001:2015 outlines a process-oriented approach to documenting and reviewing the structure, responsibilities, and procedures required to achieve effective quality management within an organization. Specific sections of the standard contain information on topics such as:

- Requirements for a quality management system (QMS), including quality documentation, document control, and determining process interactions
- Responsibilities of management
- Management of resources, including human resources and an organization's work environment
- Service development, including the steps from design to delivery
- Customer satisfaction
- Measurement, analysis, and improvement of the QMS through activities like internal audits and corrective and preventive actions

The AWS ISO 9001:2015 certification directly supports customers who develop, migrate and operate their quality-controlled IT systems in the AWS Cloud. You can use AWS compliance reports as evidence for your own ISO 9001:2015 programs and industry-specific quality programs, such as GxP in life sciences and ISO 131485 in medical devices.

ISO/IEC 27001

ISO/IEC 27001:2013 is a widely adopted global security standard that sets out requirements and best practices for a systematic approach to managing company and customer information that's based on periodic risk assessments appropriate to ever-changing threat scenarios. To achieve the certification, a company must show it has a systematic and ongoing approach to managing information security risks that affect the confidentiality, integrity, and availability of company and customer information.

This widely recognized international security standard specifies that AWS does the following:

- We systematically evaluate AWS information security risks, considering the impact of threats and vulnerabilities.
- We design and implement a comprehensive suite of information security controls and other forms of risk management to address customer and architecture security risks.
- We have an overarching management process to ensure that the information security controls meet our needs on an ongoing basis.

AWS has achieved ISO 27001 certification of the Information Security Management System (ISMS) covering AWS infrastructure, data centers, and services.

ISO/IEC 27017

ISO/IEC 27017:2015 provides guidance on the information security aspects of cloud computing, recommending the implementation of cloud-specific information security controls that supplement the guidance of the ISO/IEC 27002 and ISO/IEC 27001 standards. This code of practice provides additional information security controls implementation guidance specific to cloud service providers.

The AWS attestation to the ISO/IEC 27017:2015 standard not only demonstrates an ongoing commitment to align with globally recognized best practices but also verifies that AWS has a system of highly precise controls in place that are specific to cloud services.

ISO/IEC 27018

ISO 27018 is the first international code of practice that focuses on protection of personal data in the cloud. It's based on ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls applicable to public cloud personally identifiable information (PII). It also provides a set of additional controls and associated guidance intended to address public cloud PII protection requirements not addressed by the existing ISO 27002 control set.

AWS has achieved ISO 27018 certification, an internationally recognized code of practice, which demonstrates the commitment of AWS to the privacy and protection of your content.

HITRUST

The Health Information Trust Alliance Common Security Framework (HITRUST CSF) uses nationally and internationally accepted standards and regulations such as GDPR, ISO, NIST, PCI, and HIPAA to create a comprehensive set of baseline security and privacy controls.

HITRUST has developed the HITRUST CSF Assurance Program, which incorporates the common requirements, methodology, and tools that enable an organization and its business partners to take a consistent and incremental approach to managing compliance. Further, it allows business partners and vendors to assess and report against multiple sets of requirements.

Certain AWS services have been assessed under the HITRUST CSF Assurance Program by an approved HITRUST CSF Assessor as meeting the HITRUST CSF Certification Criteria. The certification is valid for 2 years, describes the AWS services that have been validated, and can be accessed at Health Information Trust Alliance Common Security Framework. You might consider using the AWS HITRUST CSF certification of AWS services to support your own HITRUST CSF certification, to complement your GxP compliance programs.

CSA Security, Trust & Assurance Registry

In 2011, the Cloud Security Alliance (CSA) launched the Security, Trust & Assurance Registry (STAR), an initiative to encourage transparency of security practices within cloud providers. STAR is a free, publicly accessible registry that documents the security controls provided by various cloud computing offerings, thereby helping users assess the security of cloud providers they currently use or are considering.

AWS participates in the voluntary STAR self-assessment to document AWS compliance with CSA-published best practices. AWS publishes the completed CSA Consensus Assessments Initiative Questionnaire (CAIQ) on the AWS website.

National Institute of Standards and Technology

The National Institute of Standards and Technology (NIST) 800-53 security controls are generally applicable to US Federal Information Systems. Federal Information Systems typically must go through a formal assessment and authorization process to ensure sufficient protection of confidentiality, integrity, and availability of information and information systems.

AWS cloud infrastructure and services have been validated by third-party testing performed against the NIST 800-53 Revision 4 controls and additional FedRAMP requirements. AWS has received FedRAMP Authorizations to Operate (ATO) from multiple authorizing agencies for the AWS GovCloud (US) and the AWS US East and AWS US West Regions.

Infrastructure controls

Cloud models (nature of the cloud)

Cloud computing is the on-demand delivery of compute power, database storage, applications, and other IT resources through a cloud services platform through the internet with pay-as-you-go pricing. As cloud computing has grown in popularity, several different models and deployment strategies have emerged to help meet specific needs of different users. Each type of cloud service and deployment method provides you with different levels of control, flexibility, and management.

Cloud computing models

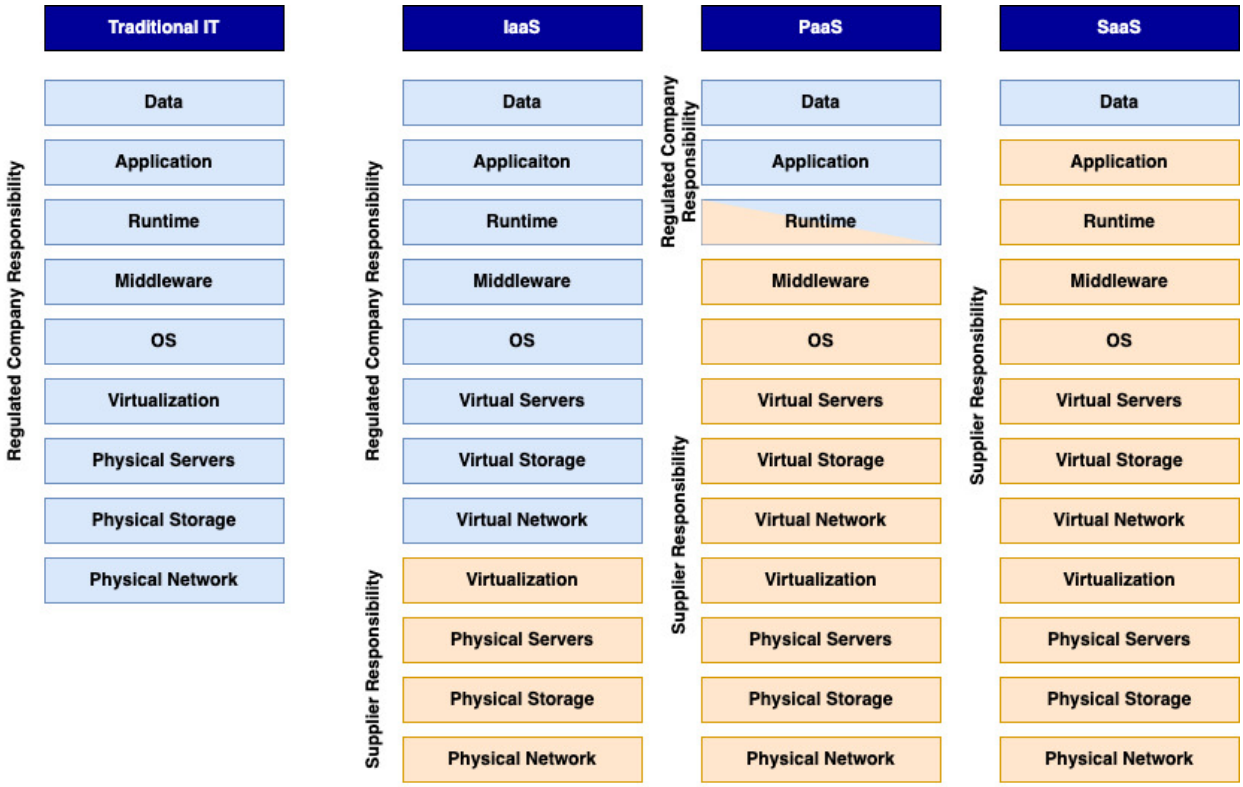

Figure 2 – Comparing shared responsibility of traditional IT, IaaS, PaaS, and SaaS

In a traditional on-premises model, the organization is responsible for every layer of the technology stack including the physical network, servers, applications, and data. As organizations move from IaaS to platform as a service (PaaS) to software as a service (SaaS), more responsibility shifts to the supplier. IaaS transfers physical infrastructure and virtualization. PaaS additionally transfers the operating system, middleware, and runtime. SaaS transfers responsibility for the application itself, while the organization retains responsibility for the data and how it is used.

The preceding diagram is an adaptation of existing models which took the same layers depicted for *Traditional IT* and applied them to the cloud. However, this gives the impression that all server, storage and network management is handled by the cloud service provider. Although this is true for the physical infrastructure, a regulated company will still have a virtual network, use storage services and need to manage a fleet of virtual servers. This means control domains, such as patch management, still apply to the regulated company. For that reason, the Virtual Network, Storage and Server layers have been added.

The following diagram applies the preceding model to AWS services:

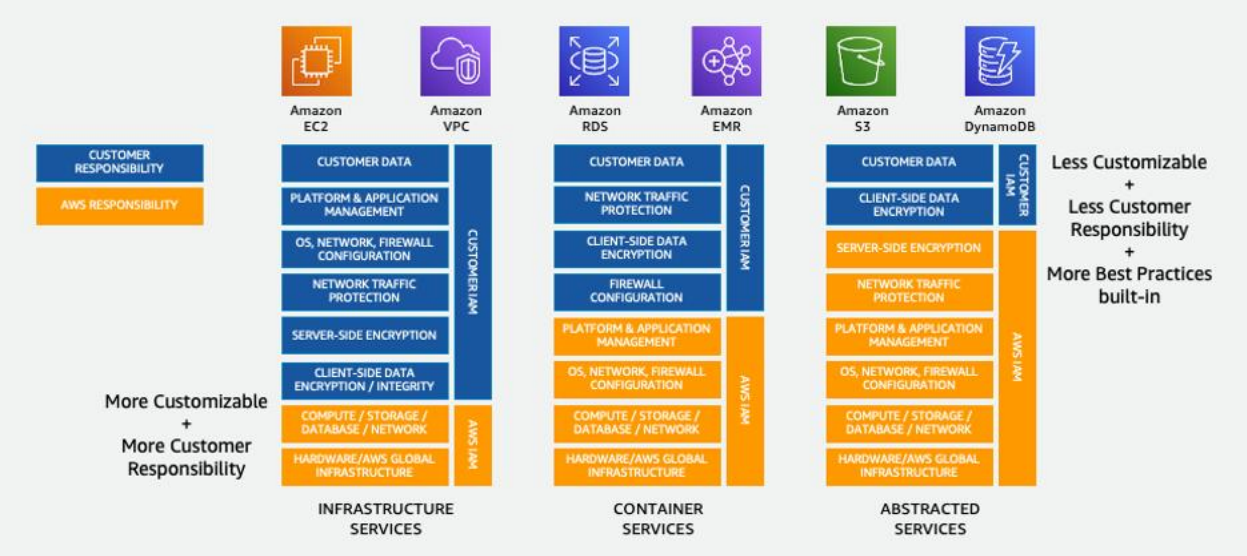

Figure 3 – Shared responsibility spectrum across AWS services

Different AWS services provide different customer responsibility boundaries. For most AWS services, AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. The customer assumes responsibility and management of the guest operating system (including updates and security patches), other associated application software, and the configuration of the AWS-provided security group firewall. Core infrastructure services like Amazon EC2 offer more customization but require more customer responsibility, while managed services like Amazon S3 shift more responsibility to AWS with security best practices built in by default. Container services provide a balance between customization and customer responsibility. The customer always retains responsibility for their data and client-side encryption, while AWS manages progressively more of the underlying infrastructure, services, and network controls, and so on. Because the boundary between customer and AWS responsibilities differs by AWS service, it's recommended that customers assess where the shared responsibility boundary falls for each service to implement proper controls and configurations.

Infrastructure as a service

Infrastructure as a service (IaaS) contains the basic building blocks for cloud IT and typically provides access to networking features, computers (virtual or on dedicated hardware), and data storage space. IaaS provides you with the highest level of flexibility and management control

over your IT resources and is most similar to the existing IT resources that many IT departments and developers are familiar with today (for example, Amazon EC2).

While cloud providers maintain the physical hardware infrastructure, organizations using IaaS retain substantial control over their operating systems, storage configurations, networking components, security protocols, and application deployment. The specific balance of management responsibilities between the service provider and the customer organization varies according to the chosen service offering.

Platform as a service

Platform as a service (PaaS) removes the need for organizations to manage the underlying infrastructure (usually hardware and operating systems) and so you can focus on the deployment and management of your applications (for example, AWS Elastic Beanstalk). This helps you be more efficient because you don't need to worry about resource procurement, capacity planning, software maintenance, patching, or any of the other undifferentiated heavy lifting involved in running your application.

Software as a service

Software as a service (SaaS) provides you with a completed product that is run and managed by the service provider. In most cases, people referring to Software as a Service are referring to end-user applications (for example, Amazon Connect). With a SaaS offering, you don't have to think about how the service is maintained or how the underlying infrastructure is managed; you only need to think about how you will use that particular piece of software. A common example of a SaaS application is web-based email, which can be used to send and receive email without having to manage feature additions to the email product or maintain the servers and operating systems on which the email program is running.

Cloud computing deployment models

Cloud

A cloud-based application is fully deployed in the cloud and all parts of the application run in the cloud. Applications in the cloud have either been created in the cloud or have been migrated from an existing infrastructure to take advantage of the benefits of cloud computing (see What is cloud computing?). Cloud-based applications can be built on low-level infrastructure pieces or can use higher level services that provide abstraction from the management, architecting, and scaling requirements of core infrastructure.

Hybrid

A hybrid deployment is a way to connect infrastructure and applications between cloud-based resources and existing resources that aren't located in the cloud. The most common method of hybrid deployment is between the cloud and existing on-premises infrastructure to extend—and grow—an organization's infrastructure into the cloud while connecting cloud resources to the internal system. For more information about how AWS can help you with hybrid deployment, see (Hybrid Cloud with AWS).

On-premises

The deployment of resources on-premises, using virtualization and resource management tools, is sometimes sought for its ability to provide dedicated resources. In most cases this deployment model is the same as legacy IT infrastructure while using application management and virtualization technologies to try and increase resource utilization. For more information, see Hybrid Cloud with AWS.

Security

Physical security

Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in facilities that aren't branded as AWS facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff using video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. All visitors are required to present identification and are signed in and continually escorted by authorized staff. AWS only provides data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or AWS. All physical access to data centers by AWS employees is logged and audited routinely.

For more information about infrastructure security, see AWS Data Center Controls.

Single- or multi-tenant environments

As cloud technology has rapidly evolved over the past decade, one fundamental technique used to maximize physical resources and lower customer costs has been to offer multi-tenant services to cloud customers. To facilitate this architecture, AWS has developed and

implemented powerful and flexible logical security controls to create strong isolation boundaries between customers. At AWS, Security is our top priority, and you will find a rich history of AWS steadily enhancing its features and controls to help customers achieve their security posture requirements such as GxP. Coming from operating an on-premises environment, you will often find that Cloud Service Providers like AWS enable you to effectively optimize your security configurations in the cloud compared to your on-premises solutions.

The AWS logical security capabilities and security controls address the concerns driving physical separation to protect your data. The provided isolation combined with the added automation and flexibility offers a security posture that matches or bests the security controls seen in traditional, physically separated environments.

Additional detailed information on logical separation on AWS can be found in the Logical Separation on AWS whitepaper.

Global infrastructure

Geography

AWS serves over a million active customers in more than 200 countries. As customers grow their businesses, AWS will continue to provide infrastructure that meets their global requirements.

The AWS Cloud is built around AWS Regions and Availability Zones. An AWS Region is a physical location in the world and has multiple AZs. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities. These AZs offer you the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center. The AWS Cloud operates in over 70 AZs in over 20 geographic Regions around the world, with announced plans for more AZs and Regions. For more information about AWS AZs and Regions, see AWS Global Infrastructure.

Each AWS Region is designed to be completely isolated from the other Regions. This achieves the greatest possible fault tolerance and stability. Each Availability Zone is isolated, but the AZs in a Region are connected through low-latency links. AWS provides customers with the flexibility to place instances and store data within multiple geographic regions and across multiple AZs within each AWS Region. Each AZ is designed as an independent failure zone. This means that AZs are physically separated within a typical metropolitan region and are located in lower risk flood plains (specific flood zone categorization varies by AWS Region). In addition to discrete uninterruptable power supplies (UPS) and onsite backup generation facilities, they are each fed

through different grids from independent utilities to further reduce single points of failure. Availability Zones are all redundantly connected to multiple tier-1 transit providers.

Data locations

Where geographic limitations apply, unlike other cloud providers, who often define a region as a single data center, the multiple Availability Zone (AZ) design of every AWS Region offers you advantages. If you're focused on high availability, you can design your applications to run in multiple AZs to achieve even greater fault-tolerance. AWS Regions meet the highest levels of security, compliance, and data protection. If you have data residency requirements, you can choose the AWS Region that is in close proximity to your desired location. You retain complete control and ownership over the Region in which your data is physically located, making it straightforward to meet regional compliance and data residency requirements.

In addition, for moving on-premises data to AWS for migrations or ongoing workflows, Cloud Data Migration describes the various tools and services that you can use to ensure data onshoring compliance, including:

- Hybrid cloud storage (AWS Storage Gateway, AWS Direct Connect)
- Online data transfer (AWS DataSync, AWS Transfer Family, Amazon S3 Transfer Acceleration, AWS Snowcone, Amazon Kinesis Data Firehose, and APN Partner Products)
- Offline data transfer (AWS Snowcone, AWS Snowball, and AWS Snowmobile)

Capacity

When it comes to capacity planning, AWS examines capacity at both a service and rack usage level. The AWS capacity planning process also automatically triggers the procurement process for approval so that AWS doesn't have additional lag time to account for, and AWS relies on capacity planning models—which are informed in part by customer demand—to trigger new data center builds. AWS enables you to reserve instances so that space is guaranteed in the AWS Regions of your choice. AWS uses the number of reserved instances to inform planning for FOOB (future out of bound).

Uptime and change management

To ensure accountability, AWS maintains documented service level agreements (SLAs) for each service, with core services guaranteed at minimum 99.9% availability. Regular SLA performance reporting, service credits for any breaches, and transparent incident reporting with root cause analysis demonstrate the commitment of AWS to service reliability. A full list of AWS SLAs can be found at AWS Service Level Agreements.

AWS actively manages service uptime through a comprehensive system of monitoring and change management. AWS also maintains continuous real-time service health monitoring through the AWS Service Health Dashboard, employing automated detection and response mechanisms for potential availability issues while strategically scheduling proactive maintenance to minimize customer impact.

Change management follows a sophisticated approach where service updates are deployed through controlled rollout waves and blue/green deployments to ensure zero-downtime updates. AWS communicates these changes through advance notifications on the AWS Health Dashboard and implements a regional deployment strategy to contain any potential impact of changes.

Supporting these technical measures, AWS provides extensive guidance to customers on architecting for high availability, offering technical documentation for resilient design patterns and supporting business continuity planning through the Well-Architected Framework. This comprehensive approach ensures service reliability while enabling continuous improvement and feature deployment, all while maintaining transparent communication with customers about service health and changes.

AWS Quality Management System

Life Science customers with obligations under GxP requirements need to ensure that quality is part of manufacturing and controls during the design, development and deployment of their GxP-regulated product. This quality assurance includes an appropriate assessment of cloud service suppliers, such as AWS, to meet the obligations of your quality system.

For a deeper description of the AWS Quality Management System, you can use AWS Artifact to access additional documents under NDA. In this section, AWS provides information on some of the concepts and components of the AWS Quality System of most interest to GxP customers like you.

Quality infrastructure and support processes

Quality management system certification

Customers need to evaluate and select their potential suppliers, contractors, and consultants on the basis of their ability to meet specified requirements. To ensure the quality and security of AWS products, AWS operates an industry-leading management control framework that conforms to current quality, security, and trust standards for commercial IT organizations. AWS

has undergone a systematic, independent examination of our quality system to determine whether the activities and activity outputs comply with ISO 9001:2015 requirements. A certifying agent found our quality management system (QMS) to comply with the requirements of ISO 9001:2015 for the activities described in the scope of registration.

The AWS quality management system has been certified to ISO 9001 since 2014. The reports cover 6-month periods each year (April–September and October–March). New reports are released in mid-May and mid-November. To see the AWS ISO 9001 registration certification, certification body information in addition to date of issuance and renewal, please see the information on the ISO 9001 AWS compliance program webpage.

AWS recognizes that customers might conduct their own periodic supplier reassessments according to their internal quality requirements and timelines, which might differ from the AWS compliance program assessment schedules.

The certification covers the QMS over a specified scope of AWS services and Regions of operations. If you are pursuing ISO 9001:2015 certification while operating all or part of your IT systems in the AWS cloud, you are not automatically certified by association, however, using an ISO 9001:2015 certified provider like AWS can make your certification process easier. AWS Security Assurance Services are a set of professional services where customers partner with AWS auditors and engineers to help meet their security and compliance requirements. These services help customers with compliance tasks, risk control measures, and security governance by providing expertise, tools, and guidance to align with cloud-native technologies and regulatory needs.

AWS provides additional detailed information on the quality management system accessible within AWS Artifact through customer accounts in the AWS Management Console.

Quality procedures

In addition to the software, hardware, human resource and real estate assets that are encompassed in the scope of the AWS Quality Management System supporting the development and operations of AWS services, it also includes documented information including, but not limited to source code, system documentation and operational policies and procedures.

AWS implements formal, documented policies and procedures that provide guidance for operations and information security within the organization and the supporting AWS environments. Policies address purpose, scope, roles, responsibilities and management commitment. All policies are maintained in a centralized location that is accessible by employees.

Quality organization roles

AWS Security Assurance is responsible for familiarizing employees with the AWS security policies. AWS has established information security functions that are aligned with defined structure, reporting lines, and responsibilities. Leadership involvement provides clear direction and visible support for security initiatives. AWS provides additional detailed information on the quality management system accessible within AWS Artifact through customer accounts in the AWS console.

Quality project planning and reporting

The AWS planning process defines service requirements, requirements for projects and contracts, and ensures customer needs and expectations are met or exceeded. Planning is achieved through a combination of business and service planning, project teams, quality improvement plans, review of service-related metrics and documentation, self-assessments and supplier audits, and employee training. The AWS quality system is documented to ensure that planning is consistent with all other requirements.

AWS continuously monitors service usage to project infrastructure needs to support availability commitments and requirements. AWS maintains a capacity planning model to assess infrastructure usage and demands at least monthly, and usually more frequently. In addition, the AWS capacity planning model supports the planning of future demands to acquire and implement additional resources based upon current resources and forecasted requirements.

Electronic records and electronic signatures

In the United States (US), GxP regulations are enforced by the US Food and Drug Administration (FDA) and are contained in Title 21 of the Code of Federal Regulations (21 CFR). Within 21 CFR, Part 11 contains the requirements for computer systems that create, modify, maintain, archive, retrieve, or distribute electronic records and electronic signatures in support of GxP-regulated activities (and in the EU, EudraLex - Volume 4 - Good Manufacturing Practice (GMP) guidelines – Annex 11 Computerized Systems). Part 11 was created to permit the adoption of new information technologies by FDA-regulated life sciences organizations, while simultaneously providing a framework to ensure that the electronic GxP data is trustworthy and reliable.

There is no GxP certification for a commercial cloud provider such as AWS. AWS offers commercial off-the-shelf (COTS) IT services according to IT quality and security standards such as ISO 27001, ISO 27017, ISO 27018, ISO 9001, NIST 800-53 and many others. GxP-regulated life sciences customers, like you, are responsible for purchasing and using AWS services to develop and operate your GxP systems, and to verify your own GxP compliance, and compliance with 21 CFR 11.

This document, used in conjunction with other AWS resources noted throughout, may be used to support your electronic records and electronic signatures requirements. A further description of the shared responsibility model as it relates to your use of AWS services in alignment with 21 CFR 11 can be found in the Appendix.

Company self-assessments

AWS Security Assurance monitors the implementation and maintenance of the quality management system by performing verification activities through the AWS audit program to ensure compliance, suitability, and effectiveness of the quality management system. The AWS audit program includes self-assessments, third party accreditation audits, and supplier audits. The objective of these audits are to evaluate the operating effectiveness of the AWS quality management system. Self-assessments are performed periodically. Audits by third parties for accreditation are conducted to review the continued performance of AWS against standards-based criteria and to identify general improvement opportunities. Supplier audits are performed to assess the supplier's potential for providing services or material that conform to AWS supply requirements. AWS maintains a documented schedule of all assessments to ensure implementation and operating effectiveness of the AWS control environment to meet various objectives.

Contract reviews

AWS offers services for sale under a standardized customer agreement that has been reviewed to ensure the services are accurately represented, properly promoted, and fairly priced. Please contact your account team if you have questions about AWS service terms.

Corrective and preventive actions

AWS takes action to eliminate the cause of nonconformities within the scope of the quality management system to prevent recurrence. The following procedure is followed when taking corrective and preventive actions:

1. Identify the specific nonconformities
2. Determine the causes of nonconformities
3. Evaluate the need for actions to ensure that nonconformities do not recur
4. Determine and implement the corrective actions needed
5. Record results of actions taken
6. Review of the corrective actions taken
7. Determine and implement preventive action needed
8. Record results of action taken

9. Review of preventive action

The records of corrective actions can be reviewed during regularly scheduled AWS management meetings.

Customer complaints
AWS relies on procedures and specific metrics to support you. Customer reports and complaints are investigated and, where required, actions are taken to resolve them. You can contact AWS or speak directly with your account team for support. AWS provides additional detailed information accessible within AWS Artifact through customer accounts in the AWS console.

Third-party management
AWS maintains a supplier management team to foster third-party relationships and monitor third party performance. SLAs and SLOs are implemented to monitor performance.

AWS creates and maintains written agreements with third parties (for example, contractors or vendors) in accordance with the work or service to be provided (for example, network services, service delivery, or information exchange) and implements appropriate relationship management mechanisms in line with their relationship to the business. AWS monitors the performance of third parties through periodic reviews using a risk-based approach, which evaluate performance against contractual obligations.

Infrastructure management
The infrastructure team maintains and operates a configuration management framework to address hardware scalability, availability, auditing, and security management. By centrally managing hosts by using automated processes that manage change, Amazon is able to achieve its goals of high availability, repeatability, scalability, security, and disaster recovery. Systems and network engineers monitor the status of these automated tools on a continuous basis, reviewing reports to respond to hosts that fail to obtain or update their configuration and software. AWS provides additional detailed information accessible within AWS Artifact through customer accounts in the AWS console.

Software development

Software development processes
The Project and Operation stages of the life cycle approach in GAMP, for instance, are reflected in the AWS information and activities surrounding organizational mechanisms to guide the

development and configuration of the information system, including software development lifecycles and software change management. Elements of the organizational mechanisms include policies and standards, the code pathway, deployment, a change management tool, ongoing monitoring, security reviews, emergency changes, management of outsourced and unauthorized development and communication of changes to customers.

The software development lifecycle activities at AWS include the code development and change management processes at AWS which are centralized across AWS teams developing externally facing and internally facing code with processes applying to both internal and external service teams. Code deployed at AWS is developed and managed in a consistent process, regardless of its ultimate destination. There are several systems used in this process, including:

- A code management system used to assemble a code package as part of development.
- Internal source code repository.
- The hosting system in which AWS code pipelines are staged.
- The tool used for automating the testing, approval, deployment, and ongoing monitoring of code.
- A change management tool which breaks change workflows down into discrete, easy to manage steps and tracks change details.
- A monitoring service to detect unapproved changes to code or configurations in production systems. Any variances are escalated to the service owner or team.

Deployment and testing

A pipeline represents the path approved code packages take from initial check-in through a series of automated (and potentially manual) steps to execution in production. The pipeline is where automation, testing, and approvals happen.

At AWS, the deployment tool is used to create, view, and enforce code pipelines. This tool is utilized to promote the latest approved revision of built code to the production environment.

A major factor in ensuring safe code deployment is deploying in controlled stages and requiring continuous approvals prior to pushing code to production. As part of the deployment process, pipelines are configured to release to test environments (for example, *beta*, *gamma*, and others, as defined by the team) prior to pushing the code to the production environment. Automated quality testing (for example, integration testing, black box testing, and white box testing) is performed in these environments to ensure code is performing as anticipated. If code

is found to deviate from standards, the release is halted and the team is notified of the need to review.

These development and test environments emulate the production environment and are used to properly assess and prepare for the impact of a change to the production environment. To reduce the risks of unauthorized access or change to the production environment, the development, test, and production environments are all logically separated.

The tool additionally enforces phased deployment, if the code is to be deployed across multiple Regions. Should a package include deployment for more than one AWS Region, the pipeline will enforce deployment on a single-Region basis. If the package were to fail integration tests at any Region, the pipeline is halted and the team is notified for need to review. Within each Region, deployments use a rolling deployment strategy that replaces no more than 33% of capacity at once. AWS employs automated monitoring and rollback mechanisms during deployments, including continuous health checks, anomaly detection, and automatic rollback triggers when performance degradation is detected. This automation enables safe, hands-off deployments while maintaining service reliability.

Configuration and change management

Configuration management is performed during information system design, development, implementation, and operation through the use of the AWS Change Management process.

Routine, emergency, and configuration changes to existing AWS infrastructure are authorized, logged, tested, approved, and documented in accordance with industry norms for similar systems. Updates to the AWS infrastructure are done to minimize any impact on you and your use of the services. This rigorous change management process enables AWS to safely deploy changes multiple times per day across services while maintaining the high standards required for GxP compliance.

Change management

AWS applies a systematic approach to managing change so that changes to customer-impacting services are thoroughly reviewed, tested, approved, and well-communicated. The AWS change management process is designed to avoid unintended service disruptions and to maintain the integrity of service to you.

Types of changes and customer communication

AWS categorizes service changes into three primary types based on customer impact:

Deprecations: These involve the retirement of existing features, APIs, or services. Customers receive advanced notification (typically more than 12 months) through multiple channels including the AWS Health Dashboard, service-specific notifications, and direct customer outreach. Clear migration paths and documentation are provided to support the transition.

New features and services*:* These represent additions to the AWS platform, including new service releases, significant feature additions, and new API operations. Customers are notified through AWS What's New announcements, service-specific release notes, the AWS Health Dashboard, developer documentation, and blog posts.

Non-customer impacting changes: These include internal improvements, infrastructure updates, and backend optimizations that do not affect customer-facing functionality or require customer action. While these changes undergo the same rigorous testing and deployment processes described in this document, they are not communicated to customers. The volume of such changes is substantial and communicating each would result in notification fatigue without providing actionable information. AWS focuses customer communications on changes that require awareness or action, ensuring that notifications remain meaningful and relevant.

Changes deployed into production environments are:

- **Prepared**: this includes scheduling, determining resources, creating notification lists, scoping dependencies, minimizing concurrent changes, and a special process for emergent or long-running changes.
- **Submitted**: this includes using a change management tool to document and request the change, determine potential impact, conduct a code review, create a detailed timeline and activity plan, and develop a detailed rollback procedure.
- **Reviewed and approved**: Peer reviews of the technical aspects of a change are required. Changes must be authorized in order to provide appropriate oversight and understanding of business and security impact. The configuration management process includes key organizational personnel that are responsible for reviewing and approving proposed changes to the information system.
- **Tested**: Changes being applied are tested to help ensure they will behave as expected and not adversely impact performance.
- **Performed**: This includes pre- and post-change notification, managing the timeline, monitoring service health and metrics, and closing out the change.

AWS service teams maintain a current authoritative baseline configuration for systems and devices. Change management tickets are submitted before changes are deployed (unless it is an emergency change) and include impact analysis, security considerations, description, timeframe and approvals. Changes are pushed into production in a phased deployment starting with lowest impact areas. Deployments are tested on a single system and closely monitored so impacts can be evaluated. Service owners have a number of configurable metrics that measure the health of the service's upstream dependencies. These metrics are closely monitored with thresholds and alarming in place. Rollback procedures are documented in the change management (CM) ticket. AWS service teams retain older versions of AWS baseline packages and configurations necessary to support rollback and previous versions are stored in the repository systems. Integration testing and the validation process is performed before rollbacks are implemented. When possible, changes are scheduled during regular change windows.

In addition to the preventative controls that are part of the pipeline (for example, code review verifications, test environments), AWS also uses detective controls configured to alert and notify personnel when a change is detected that may have been made without standard procedure. AWS checks deployments to ensure that they have the appropriate reviews and approvals to be applied before the code is committed to production. Exceptions for reviews and approvals for production lead to automatic ticketing and notification of the service team.

After code is deployed to the production environment, AWS performs ongoing monitoring of performance through a variety of monitoring processes. AWS host configuration settings are also monitored as part of vulnerability monitoring to validate compliance with AWS security standards. Audit trails of the changes are maintained.

Emergency changes to production systems that require deviations from standard change management procedures are associated with an incident and are logged and approved as appropriate. Periodically, AWS performs self-audits of changes to key services to monitor quality, maintain high standards, and facilitate continuous improvement of the change management process. Any exceptions are analyzed to determine the root cause, and appropriate actions are taken to bring the change into compliance or roll back the change if necessary. Actions are then taken to address and remediate the process or people issue.

Reviews

AWS performs internal security reviews against Amazon security standards of externally launched products, services, and significant feature additions prior to launch to ensure security risks are identified and mitigated before deployment to a customer environment. AWS security reviews include evaluating the service's design, threat model, and impact to the risk profile of

AWS. A typical security review starts with a service team initiating a review request to the dedicated team and submitting detailed information about the artifacts being reviewed. Based on this information, AWS reviews the design and identifies security considerations; these considerations include but are not limited to: appropriate use of encryption, analysis of data handling, regulatory considerations, and adherence to secure coding practices. Hardware, firmware and virtualization software also undergo security reviews, including a security review of the hardware design, actual implementation and final hardware samples.

Code package changes are subject to the following security activities:

- Full security assessment
- Threat modeling
- Security design reviews
- Secure code reviews (manual and automated methods)
- Security testing
- Vulnerability/penetration testing

Successful completion of the preceding activities are pre-requisites for service launch. Development teams are responsible for the security of the features they develop that meet the security engineering principles. Infrastructure teams incorporate security principles into the configuration of servers and network devices with least privilege enforced throughout. Findings identified by AWS are categorized in terms of risk and are tracked in an automated workflow tool.

Customer training

AWS has implemented various methods of external communication to support its customer base and the community. Mechanisms are in place to allow the customer support team to be notified of operational issues that impact your experience. A Service Health Dashboard is available and maintained by the customer support team to alert you to any issues that may be of broad impact. The AWS Cloud Security Center and Healthcare and Life Sciences Center are available to provide you with security and compliance details and Life Sciences related enablement information about AWS. You can also subscribe to AWS Support offerings that include direct communication with the customer support team and proactive alerts to any customer impacting issues.

AWS also has a series of training and certification programs on a number of cloud-related topics in addition to a series of service and support offerings available through your AWS account team.

AWS products in GxP systems

With limited technical guidance from regulatory and industry bodies, this section aims to describe some of the best practices we've seen customers adopting when using cloud services to meet their regulatory compliance needs.

The Final FDA Guidance Document, Data Integrity and Compliance With Drug CGMP explicitly brings cloud infrastructure into scope through the revised definition of "computer or related systems":

> "The American National Standards Institute (ANSI) defines systems as people, machines, and methods organized to accomplish a set of specific functions. Computer or related systems can refer to computer hardware, software, peripheral devices, networks, **cloud infrastructure**, personnel and associated documents (for example, user manuals and standard operating procedures)."

Industry organizations like ISPE are increasingly dedicating publications on cloud usage in the life sciences (Getting Ready For Pharma 4.0: Data integrity in cloud and big data applications).

Further, this recent publication of FDA document provides guidance for considerations when using cloud computing services.

As described throughout this whitepaper, there is no unique certification for GxP regulations, so each customer defines their own risk profile. Therefore, it is important to note that although this whitepaper is based on AWS experience with life science customers, you must take final accountability and determine your own regulatory obligations.

To begin with, even when deployed in the cloud, GxP applications still need to be validated and their underlying infrastructure still needs qualifying. The basic principles governing on-premise infrastructure qualification still apply to virtualized cloud infrastructure. Therefore, the current industry guidance should still be used.

Traditionally, a regulated company was accountable and responsible for all aspects of their infrastructure qualification and application validation. With the introduction of public cloud providers, part of that responsibility has been shifted to a cloud supplier. The regulated

company is still accountable, but the cloud supplier is now responsible for the qualification of the physical infrastructure, virtualization, and service layers and to completely manage the services they provide, that is, the big difference now is that there is a shared compliance responsibility model which is similar to the shared security responsibility model described earlier in this whitepaper.

Previous sections of this whitepaper described how AWS takes care of their part of the shared responsibility model. This section provides recommended strategies on how to cover your part of the shared responsibility model for GxP environments.

Involving AWS

Achieving GxP compliance when adopting cloud technology is a journey. AWS has helped many customers along this journey, and there is no compression algorithm for experience.

For example, Core Informatics states:

> "Using AWS we can help organizations accelerate discovery while maintaining GxP compliance. It's transforming our business and, more importantly, helping our customers transform their businesses."
>
> **– Richard Duffy Vice President of Engineering, Core Informatics**

For the complete case study, see Core Informatics Case Study. For a selection of other customer case studies, see AWS Customer Success.

Industry guidance recommends that companies should try and maximize supplier involvement and use our knowledge, experience and even our documentation as much as possible, as we provide in the following sections and throughout this whitepaper. Please contact us to discuss starting your journey to the cloud.

Cloud strategy and planning for life science organizations

Life sciences organizations face a critical decision point when adopting cloud technology: how to innovate rapidly while maintaining compliance with stringent GxP regulations. The fundamental challenge lies in transforming from traditional, manual, and document-centric IT processes into what can be termed *Regulated Cloud Natives*—organizations that operate at the speed of consumer-focused companies while maintaining the quality and compliance standards required by regulatory bodies.

Three approaches to cloud adoption

Option 1: Deferring regulatory compliance

Some organizations attempt to avoid dealing with GxP regulations for as long as possible. While this might be a valid tactic for quickly demonstrating cloud value through small-scale projects, it becomes problematic when extended too long, because retrospectively addressing GxP compliance requires significantly more effort. This approach fails to use the full potential of cloud transformation and can create technical debt that becomes increasingly difficult to remediate.

Option 2: Extending traditional practices

Many organizations attempt to simply extend their current ways of working to the cloud, primarily because their existing processes are documented in their quality management system (QMS) and changing them requires significant effort. However, this approach often perpetuates old and inefficient processes that were originally created with very manual controls and haven't evolved as technology and practices have progressed. Organizations taking this route frequently fail to demonstrate the value of cloud and may stall or block further cloud adoption initiatives.

The same basic principles that govern on-premises infrastructure qualification still apply to cloud-based systems. However, replicating on-premises approaches in the cloud without adaptation fails to capture the benefits of cloud-native capabilities and automation.

Option 3: Becoming regulated cloud natives

The optimal approach involves significant transformation to adopt cloud-centered operating practices within a regulated environment where highly regulated customers can operate at a cadence similar to unregulated customers. This approach recognizes that many life sciences organizations constrain themselves with hypercautious interpretations of compliance demands, creating elaborate approval processes that regulatory bodies don't mandate[1]. Tools that aren't directly part of GxP-regulated processes don't require the same level of validation as systems that directly impact product quality or patient safety.

One of the concerns for regulated enterprise customers becomes how to qualify and demonstrate control over a system when so much of the responsibility is now shared with a supplier. Your cloud strategy should directly address this concern. The strategy will employ various tactics to address your regulatory needs.

To better scope the strategy the architecture should be viewed in its entirety. Enterprise scale customers typically define the architecture similar to the following:

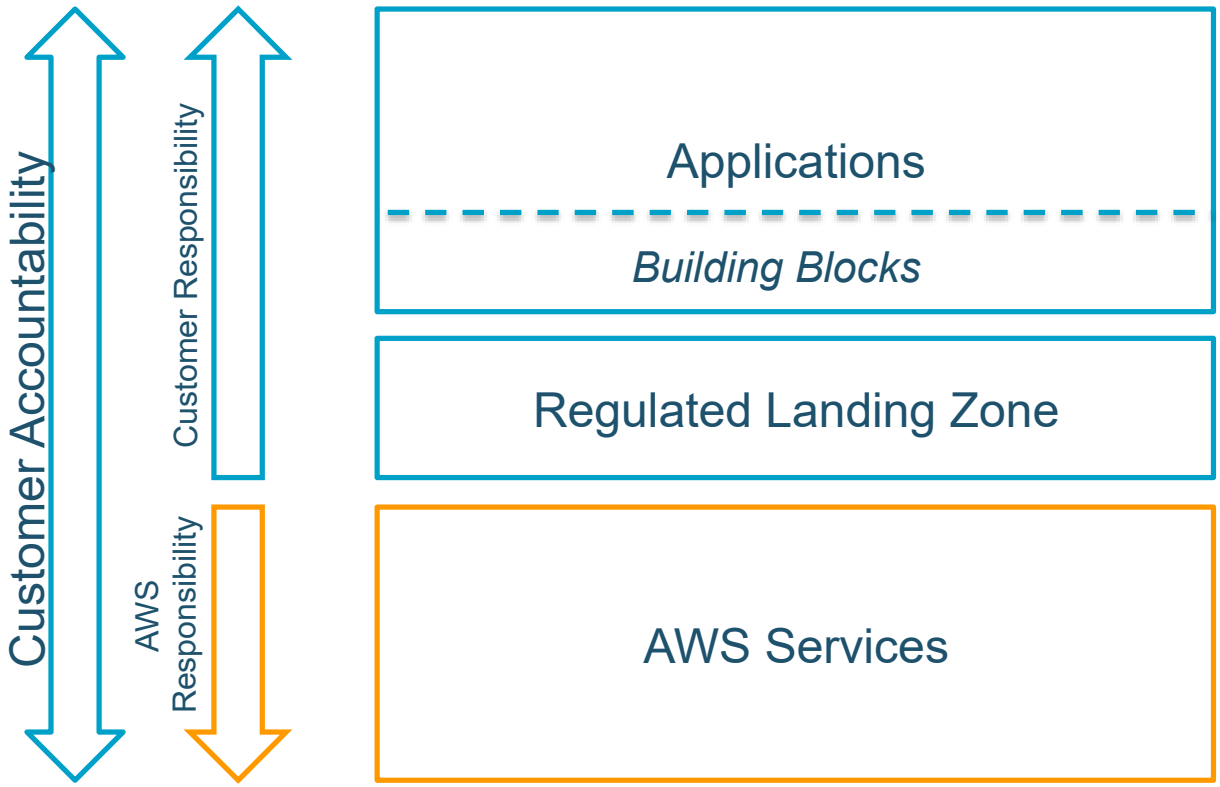

Figure 4: Layered architecture

The preceding diagram illustrates a layered architecture where a large part is delegated to AWS. From this approach, a strategy can be defined to address four main areas:

1. How to work with AWS as a supplier of services.
2. The qualification of the regulated landing zone.
3. The qualification of building blocks.
4. Supporting the development of GxP applications.

The situation also changes slightly if the customer uses a service provider, like AWS Managed Services, where the build, operation, and maintenance of the landing zone is done by the service provider. Conversely, for workloads that must remain on premises, AWS Outposts extends AWS services including compute, storage and networking to customer sites. Data can be configured to be stored locally, and customers are responsible for controlling access around Outposts equipment. Data that is processed and stored on premises is accessible over the customer's local network. In this case, customer responsibility extends into the *AWS services* box (*Figure* 5*5*).

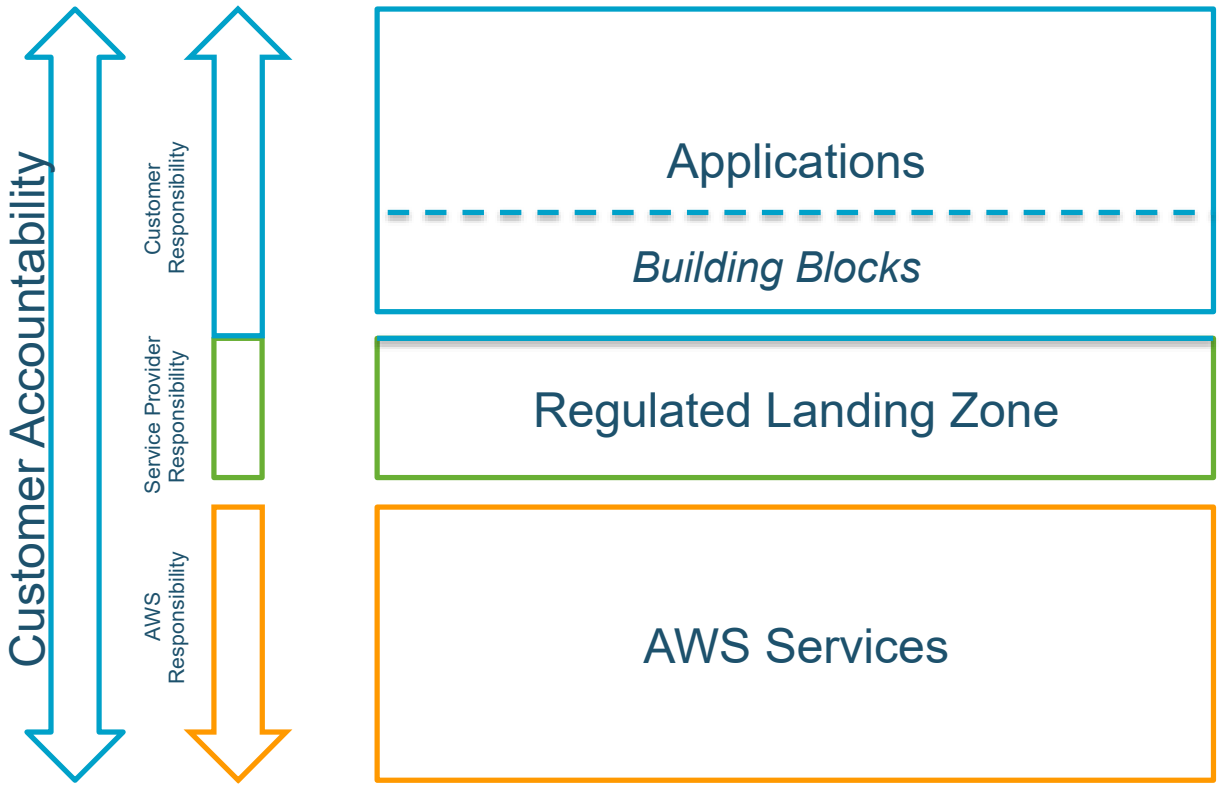

Figure 5: Layered architecture with service provider

In this situation, even more responsibility is delegated by the customer and so the controls that are typically to be put in place by the customer to control their own operations, now need adaptations to check that similar controls are implemented by the service provider. The controls that are inherited from AWS, are shared or that remain with the customer were covered previously in the Shared Security Responsibility Model section of this whitepaper.

This section describes these layers at a high level. These layers are expanded upon in later sections of this whitepaper.

Industry guidance

The following guidance is at a minimum, a best practice for your environment. You should still work with your professionals to ensure you comply with applicable regulatory requirements.

The same basic principles that govern on-premises infrastructure qualification also apply to cloud-based systems. Therefore, this strategy uses a tactic of using and building upon that same industry guidance, using a cloud perspective, based on the following ISPE GAMP Good Practice Guides (*Figure* 6*6*):

- GAMP Good Practice Guide: IT Infrastructure Control and Compliance, Second Edition
- GAMP 5: A Risk-Based Approach to Compliant GxP Computerized Systems

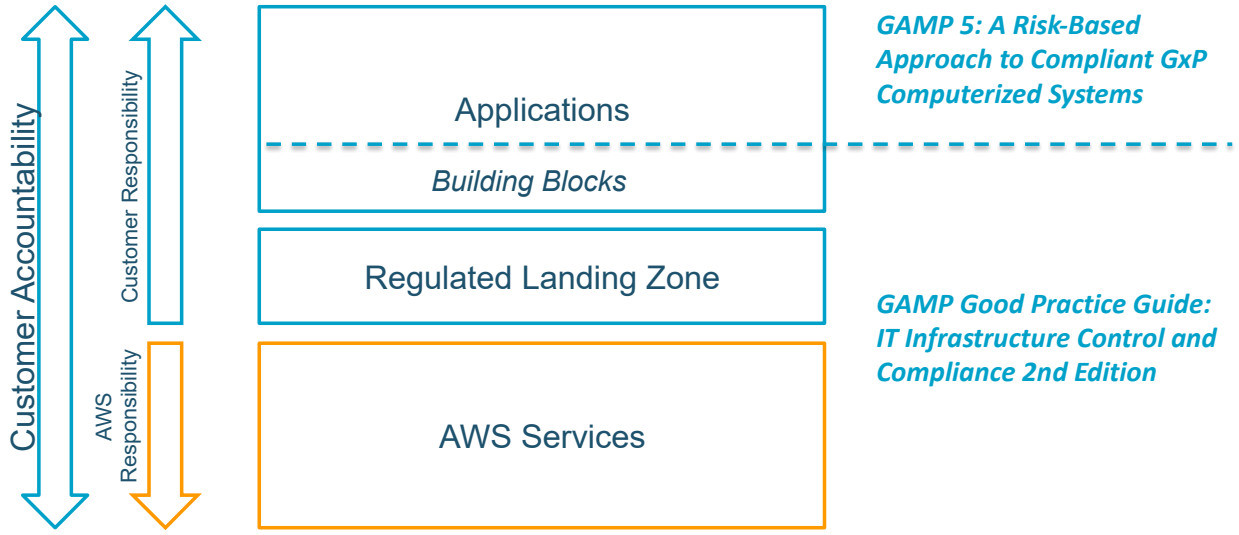

Figure 6: Mapping industry guidance to architecture layers

GxP systems assurance – The layered approach

GxP regulations require you to demonstrate control over the environment within which GxP workloads will operate. A top-level assurance model is a good starting point to explain how each layer of the entire cloud environment will be kept under a state of control. The layered model that follows provides a structure to perhaps guide the creation of new quality management system (QMS) documents or highlight existing QMS documents in need of revision.

The layered approach

One of the concerns for regulated enterprise customers becomes how to qualify and demonstrate control over a system when so much of the responsibility is now shared with a supplier. The purpose of the layered approach is to help answer this question.

To better scope the effort, the architecture should be viewed in its entirety. Enterprise scale customers typically define the high-level architecture similar to the following.

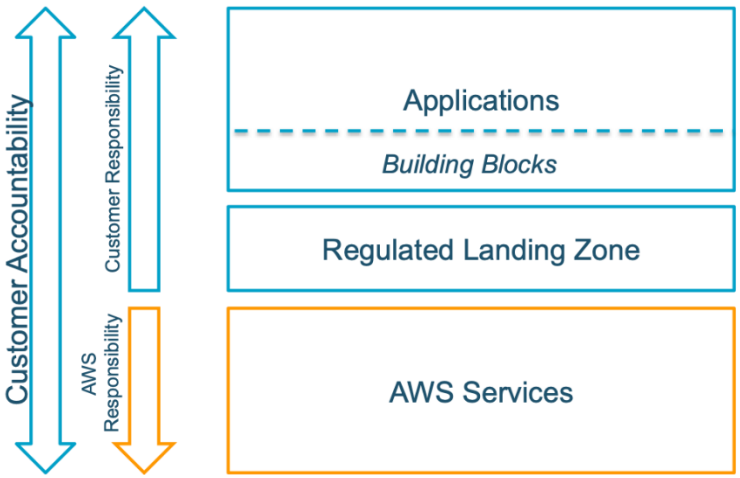

Figure 7: The layered approach

Figure 7 can be further broken down into finer grained layers as in Figure 8.

Figure 8: Detailed layer model

The layered model shows the same shared responsibility model with AWS responsible for the first foundational layers. The customer is then responsible for layers 2 and above but these layers are often owned by different internal teams, such as a platform team and the application development teams.

Layers 0 and 1 – AWS

Demonstrating control starts at the lowest level of the model with a *supplier assessment*. This assessment should consider the AWS Global Infrastructure, the AWS services that will be used,

where the shared responsibility boundary falls for each service, and use the documentation available from AWS Artifact. For GxP compliance purposes the most popular documents are often the SOC 2 and C5 reports, certifications such as ISO 9001, ISO 27001, ISO 27017, ISO 27018 and the AWS Quality Management System Overview document.

Layer 2 – Landing zone

This layer focuses on providing secure and compliant access to the underlying cloud services through accounts. A *landing zone* (LZ) is a well-architected, multi-account environment, which is a starting point from which you can deploy workloads and applications.

You can create an LZ through AWS Organizations or to expedite the creation of a complex multi-account environment AWS has developed the *Landing Zone Accelerator* (LZA), which you can use to describe your LZ in a template and use automation to deploy it. This has the advantage of simplifying change management of the LZ itself and demonstrating control. Any change to the LZ would involve a change to the template which is version controlled, approved and then deployed through the automated pipeline, increasing control and consistency.

Organizational units (OUs) can be used to group accounts together to administer as a single unit. This greatly simplifies the management of your accounts. For example, you can attach policy-based control to an OU and all accounts within the OU automatically inherit the policy. This enables the grouping of accounts subject to certain compliance regimes like GxP and apply appropriate policy controls.

Policies in AWS Organizations enable you to apply additional types of management to the AWS accounts in your organization. These are very useful when demonstrating control over the environment and can be linked to GxP requirements. There are 2 policy types, management and authorization.

Management policies enable you to centrally configure and manage AWS services and their features. One example is *backup policies*, which help you centrally manage and apply backup plans to the AWS resources across your AWS organization's accounts. This type of policy can be attached to an OU containing all GxP accounts to ensure backups are taken so ALCOA+ requirements can be satisfied. Another example is *tag policies*, which can be used to standardize the tags attached to the AWS resource in your accounts. This will help identify resources used as part of GxP workloads and help with configuration management and an application inventory.

Authorization policies help you centrally manage the security of AWS accounts in your organization. One common GxP related request is to only allow approved services to be used as

part of GxP workloads. For a service to be approved it must have been reviewed and there must be evidence available that the services is being maintained under a state of control by the cloud service provider. For example, your policy might state that only services within the scope of specific compliance programs, like SOC2, ISO 9001, ISO 27001, will be approved for general use and the audit reports and certifications will form the evidence. This policy can be implemented through *service control policies (SCPs)*, which can control which services are allowed. Other workloads with specific regulatory constraints may require additional controls. For example, only allowing HIPAA-eligible services for the processing, storage or transmission of PHI data. This variance in compliance controls can be accommodated through the use of policies applied to different OUs.

Guardrails are governance rules for security, operations, and compliance that you can define and apply either across your AWS environment or to specific groups of accounts. Guardrails protect users from making choices that aren't aligned with your overall requirements and another great way of demonstrating control over your environment. Guardrails are classified as either preventive or detective.

Preventive guardrails establish intent and prevent deployment of resources that don't conform to your policies. For example, require AWS CloudTrail to be enabled in all accounts.

Detective guardrails continuously monitor deployed resources for non-conformance and generate alerts when detected. You can automate responses to alerts to take action. For example, disallow public read access to Amazon S3 buckets. These capabilities move away from static point in time compliance checks and closer to achieving a continuous compliance stance.

When vending new accounts for application teams, having a secure account baseline is a great way of demonstrating control. For example, drop default VPC, install approved VPC, deploy security roles and stacks, hook up CloudTrail for logging, and so on. When using the previously mentioned Landing Zone Accelerator (LZA), vending new baselined accounts is simply a matter of updating the LZA template and triggering a deployment.

Deployment will also typically include verification tests to demonstrate the LZ is functioning as expected. To accomplish this the deployment pipeline can include automated tests to provide the evidence that the LZ has been deployed, configured and is working as expected.

The *delivery and operations* of the regulated landing zone and the associated capabilities will include a formal qualification effort and handover. This will include *appropriate documentation* such as designs and runbooks. This documentation needs to be maintained and updated as part

of change management. See the **Good Technical Documentation Mechanisms** section for further information.

A *cloud center of excellence (CCoE) or cloud platform team* is the team responsible for design, build, deployment and ongoing governance and operation of the landing zone. Part of the CCoE is a Cloud Business Office (CBO), which often includes people that liaise with other parts of IT, like IT Finance. A role should exist for liaising with IT Quality and Compliance to keep track of regulatory changes, feed those as requirements to the cloud platform engineering team and define the control objectives to be automated as part of the LZ.

Demonstrating control over the IT operating environment is often a focus of an inspection. This traditionally meant showing extracts from a quality management system, standard operating procedures (SOPs), and a lot of reports and evidence showing compliance to those SOPs. With the cloud, we're making this much easier with tools like Amazon Security Lake, CloudTrail Lake, AWS Config, and AWS Audit Manager. One key takeaway is that the CCoE should be looking at ways to be inspection-ready at all times.

Layer 3 – Cloud foundations and automation

This layer focuses on implementing foundational cloud capabilities and automation to support GxP workloads. It includes developing common capabilities like centralized logging and core networking, in addition to GxP required capabilities such as backup and restore, all properly documented for design, delivery, and operations.

The cloud foundations and automation layer includes identifying and templatizing common service configurations or *building blocks* that can be reused across multiple workloads, potentially undergoing formal service qualification for GxP compliance. There's a shift towards using infrastructure as code (IaC) templates for describing and deploying infrastructure, replacing traditional installation qualification protocols. Automation streamlines processes, reduces manual errors, and enables consistent deployments, including automated tests to verify correct configuration of services.

Any common tooling requirements can also be included in the landing zone and centrally managed. For example, security tooling required by a security operations center (SOC). Similarly, the migration of workloads will involve the use of some common tooling. This tooling could be installed in a shared services account.

Data integrity is core to GxP regulations and so will be an area of focus during the migration of any GxP application. The tools used for the migration of data should include the ability to verify the data has been migrated without any impact to integrity. Plan how data integrity will be

verified post migration. This can be a common approach defined for the entire migration effort and so not directly linked to an individual application. However, evidence will need to be linked to each application's databases.

Layer 4 – Workload and application infrastructure

Layer 4 focuses on the unique infrastructure requirements for each workload, which are captured in infrastructure as code (IaC) templates, such as CloudFormation templates. These IaC templates are kept under version control in a source code repository, allowing for controlled modifications and versioning when infrastructure changes are needed.

Infrastructure deployment is automated using continuous integration and deployment (CI/CD) pipelines, with AWS CodePipeline coordinating the deployment process. This shifts the approach to an exception based process where manual action is only taken if CloudFormation reports an error. To maintain synchronization between the IaC template and the actual infrastructure, all changes should only be made through automation, with CloudFormation drift detection helping to identify any discrepancies.

The traditional manual installation qualification (IQ) process is replaced by automated deployments, with CloudFormation logs serving as deployment evidence. Services like AWS Config are used to continuously monitor and verify that deployed services remain compliant with configured rules, moving towards a model of continuous compliance monitoring. For example, as part of adhering to ALCOA+ principles, Amazon S3 could be configured with versioning turned on. AWS Config can monitor this setting and alert if versioning is ever turned off. This use of IaC, automation and cloud services like AWS Config enables us to move towards continuous compliance monitoring over static verification activities.

Application-specific documentation, including design specifications and operational runbooks, should be updated to reflect the move to cloud infrastructure and services.

Layer 5 – Business applications

Layer 5 focuses on the initial validation and ongoing maintenance of the GxP applications themselves, integrating cloud activities into existing change management processes and maintaining thorough documentation throughout. For migrated workloads it involves possibly re-validating applications to ensure functionality remains intact post-migration, with the extent of testing based on the migration strategy used. Risk management is crucial, meaning a lift-and-shift approach to migrating GxP workloads is often preferred to minimize risk. Ultimately, this layer ensures that business applications remain compliant and functional in a cloud environment.

Supplier assessment

Industry guidance suggests you use a supplier's experience, knowledge and documentation as much as possible. With so much responsibility now delegated to a supplier, the supplier assessment becomes even more important. A regulated company is still ultimately accountable for demonstrating that a GxP system is compliant, even if a supplier is responsible for parts of that system, so the regulated customer needs to establish enough trust in their supplier.

The cloud service provider must be assessed to determine if they can deliver the services offered, but also to determine the suitability of their quality system and that it is systematically followed. The supplier needs to show that they have a QMS and follow a documented set of procedures and standards governing activities such as:

- Infrastructure verification and operation
- Software development
- Change management
- Release management
- Configuration management
- Supplier management
- Training
- System security

Details of the AWS QMS are covered in the AWS Quality Management System section of this whitepaper. The capabilities of AWS to satisfy these areas may be reassessed on a periodic basis, typically by reviewing the latest materials available through AWS Artifact (for example, AWS certifications and audit reports).

The approach described by industry guidance involves several steps which we will cover here.

Basic supplier assessment

The first step is to perform a basic supplier assessment to check the supplier's market reputation, knowledge, and experience working in regulated industries and prior experience working with other regulated companies and what certifications they hold.

You can use industry assessments such as Gartner's analyst report which recognizes AWS as a Leader and has been placed highest on Ability to Execute, customer testimonials and customer case studies.

The earlier sections of this whitepaper provide useful information for this initial assessment.

Documentation review

A supplier assessment often includes a deep dive into the assets available from the supplier describing their QMS and operations. This includes reviewing certifications, audit reports and whitepapers. For more information, see the AWS Risk and Compliance whitepaper.

AWS and its customers share control over the IT environment, and both parties have responsibility for managing the IT environment. The AWS part in this shared responsibility includes providing services on a highly secure and controlled platform and providing a wide array of security features customers can use. The customers' responsibility includes configuring their IT environments in a secure and controlled manner for their purposes. While customers don't communicate their use and configurations to AWS, AWS does communicate its security and control environment relevant to customers. AWS does this by doing the following:

- Obtaining industry certifications and independent third-party attestations
- Publishing information about the AWS security and control practices in whitepapers and web site content
- Providing certificates, reports, and other documentation directly to AWS customers under NDA (as required)

For a more detailed description of AWS Security, see AWS Cloud Security.

AWS Artifact provides on-demand access to AWS security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls. Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA). For a more detailed description of AWS Compliance, see AWS Compliance.

If you have additional questions about the AWS certifications or the compliance documentation AWS makes available, please bring those questions to your account team.

Review service level agreements

AWS offers service level agreements for certain AWS services. Further information can be found under Service Level Agreements (SLAs).

Audit

Mail audit – To supplement the AWS documentation you have gathered, a mail audit questionnaire (sometimes referred to as a supplier questionnaire) can be submitted to AWS to gather additional information or to ask clarifying questions. A questionnaire tailored to cloud service providers rather than a generic questionnaire will expedite the processes. You should work with your account team to request a mail audit.

Onsite audit – AWS regularly undergoes independent third-party attestation audits to provide assurance that control activities are operating as intended. Currently, AWS participates in over 50 different audit programs. The results of these audits are documented by the assessing body and made available for all AWS customers through AWS Artifact. These third-party attestations and certifications of AWS provide you with visibility and independent validation of the control environment, eliminating the need for customers to perform individual onsite audits. Such attestations and certifications may also help relieve you of the requirement to perform certain validation work yourself for your IT environment in the AWS Cloud. For details, see the AWS Quality Management System section of this whitepaper.

AWS has also been assessed by the Ingelheim Kreis Initiative Joint Audits group which represents quality and compliance professionals from some of our largest pharmaceutical and life sciences customers. Details can be found in this blog post.

Contractual agreement

After you have completed a supplier assessment of AWS, the next step is to set up a contractual agreement for using AWS services. The AWS Customer Agreement is available at: https://aws.amazon.com/agreement/. You are responsible for interpreting regulations and determining whether the appropriate requirements are included in a contract with standard terms. If you have any questions about entering into a service agreement with AWS, please contact your account team.

Cloud management

It is important to consider and plan how operational processes that span the shared responsibility model will operate. For example, how to manage changes made by AWS to services used as part of your landing zone or applications, incident response management in cases of outages, or portability requirements should there be a need to change cloud service provider.

Customer quality management system

It is important for you to analyze your existing quality management system (QMS) and the impact adopting cloud technology might have on it. It was probably written based on the use of on-premises infrastructure and some policies and procedures will need updating or exceptions raised. For example, a policy on IT asset management might imposes processes that no longer fit with cloud best practices, such as the need to record every IT asset in a CMDB. In the cloud this is done automatically with AWS Config. Using a CMDB also doesn't fit with scenarios such as auto-scaling groups where the number of servers (or assets) can grow and shrink on demand based on load. Another example might be procedures governing infrastructure installation qualification, which are often document-centric and don't take advantage of infrastructure as code (IaC).

There are also certain processes that now span the shared responsibility model and so will need updating.

Change management

Change management is a bidirectional process when dealing with a cloud service provider. On the one hand, AWS is continually making changes to improve its services as mentioned earlier in this paper. On the other hand, you can make feature requests, which is highly encouraged because 90% of AWS service features are a direct result of customer feedback.

Customers typically use a risk-based approach appropriate for the type of change to determine the subsequent actions.

Changes to AWS services that add functionality are not usually a concern because no application will be using that new functionality yet. However, new functionality might trigger an internal assessment to determine if it affects the risk profile of the service and should be allowed for use. If mandated by your QMS, this might trigger a re-qualification of building blocks prior to allowing the new functionality.

Deprecations are considered more critical because they could break an application. A deprecation might include a third-party library, utility, or version of languages such as Python. The deprecation of a service or feature is rare. When you receive the notification of a deprecation, you should trigger an impact assessment. If an impact is found, the application teams should plan changes to remediate the impact. The notice period for a deprecation should allow for time for assessment and remediation. AWS will also help you understand the impact of the change.

There are other changes such as enhancements and bug fixes that don't change the functionality of the service and don't trigger notifications to customers. These types of changes are synonymous with *standard* changes in ITIL which are usually pre-authorized, low risk, relatively common and follow a specific procedure.

Incident management

The Amazon Security Operations team employs industry-standard diagnostic procedures to drive resolution during business-impacting events. Staff operators provide continuous coverage to detect incidents and to manage the impact and resolution. As part of the process, potential breaches of customer content are investigated and escalated to AWS Security and AWS Legal. Affected customers and regulators are notified of breaches and incidents where legally required. You can subscribe to the AWS Security Bulletins webpage, which provides information regarding identified security issues. You can subscribe to the Security Bulletin RSS Feed to keep abreast of security announcements on the Security Bulletins webpage.

You are responsible for reporting incidents involving your storage, virtual machines, and applications, unless the incident is caused by AWS.

For more information see AWS Vulnerability Reporting.

Customer support

AWS develops and maintains customer support procedures that include metrics to verify performance. When you contact AWS to report that AWS services do not meet your quality objectives, your issue is investigated and, where required, commercially reasonable actions are taken to resolve it. Where AWS is the first to become aware of a customer impacting issue, procedures exist for notifying impacted customers according to their contract requirements or through the AWS Service Health Dashboard.

You should ensure that your policies and procedures align to the customer support options provided by AWS. Additional details can be found in the Customer Complaints and Customer Training sections in this document.

Cloud platform and regulated landing zone qualification

A landing zone is a well-architected, multi-account AWS environment that's based on security and compliance best practices.

One of the main functions of the landing zone is to provide a solid foundation for development teams to build on, and address as many regulatory requirements as possible, thus removing the responsibility from the development teams.

The GAMP IT Infrastructure Control and Compliance guidance document follows a platform-based approach to the qualification of IT infrastructure which aligns perfectly with a customer's need to qualify their landing zone.

The landing zone includes capabilities for centralized logging, security, account vending, and core network connectivity. There are 2 main ways of establishing a landing zone:

1. Create one using AWS Control Tower
2. Use the AWS Landing Zone Accelerator

We recommend that you use the Landing Zone Accelerator (LZA) for Healthcare which is an industry specific deployment of the LZA solution architected to align with AWS best practices and in conformance with multiple, global compliance frameworks. It is available for free as open source software on GitHub. One reason why the LZA is a good choice for a regulated landing zone is that it is configured through a template that can easily be demonstrated to be kept under a state of control. In the same way infrastructure as code (IaC) can be managed and controlled by using source code repositories and deployment automation, so can the LZA configuration.

The LZA solution expands upon the basic AWS Control Tower framework, which includes Management, Audit, and Log Archive accounts, and adds extra resources and guardrails designed to strengthen your platform's security, compliance, and operational readiness. Keep in mind that deploying LZA for Healthcare is not a complete compliance solution on its own. Instead, think of it as a strong foundation that can be augmented with other solutions to meet your needs. Your organization remains responsible for thoroughly reviewing, testing, and validating that the solution aligns with your specific security requirements, including your unique tools, features, and system configurations. The objective of the landing zone, and the team owning it, should be to provide the guardrails and features that free the developers to use the *right tools for the job* and focus on delivering differentiated business value rather than on compliance.

The LZA for Healthcare Solution incorporates security and compliance controls aligned with multiple international healthcare standards, including HIPAA (United States), NCSC (United Kingdom), ENS High (Spain), C5 (Germany), and Fascicolo Sanitario Elettronico (Italy). Before implementing the LZA on AWS for Healthcare solution, we recommend discussing your specific compliance requirements with your AWS representative to ensure proper alignment.

This solution implements comprehensive security controls in AWS environments through both detective (AWS Config rules) and preventative (service control policies) guardrails. AWS Config monitors and records AWS resource configurations, evaluates compliance against rules, and tracks changes for auditing and security. Service control policies are AWS Organizations rules that limit permissions across accounts, controlling which services and actions users can access. The configuration framework addresses three critical areas: organizational controls, security settings, and global parameters. Organizational controls include policies for service management, resource tagging, and backup strategies, while security settings enable core AWS security services and implement detective guardrails through multiple compliance frameworks such as AWS Security Hub, CIS Benchmarks, NIST guidelines, and HIPAA compliance rules. Global settings manage configurations for active Regions and centralized logging through CloudTrail and CloudWatch services. All these components work together to create a robust security foundation, though organizations should review and customize these configurations to align with their specific compliance requirements and operational needs.

Tooling and automation

Many customers include common tooling and automation as part of a landing zone so it can be qualified and validated once and used by all development teams. This common tooling is often within the shared services account of the landing zone.

For example, standard tooling around requirements management, test management, CI/CD, and so on need to be qualified and validated.

Similarly, any automation of IT processes also needs to be validated. For example, it's possible to automate the infrastructure verification step of your computer systems validation process.

Using managed services

Instead of building and operating a landing zone yourself, you have the option of delegating this responsibility. This delegation could be to AWS by using AWS Managed Services or to a partner within the AWS Partner Network (APN). This means the service provider is responsible for

building a landing zone based on AWS best practices, operating it in accordance with industry best practices and providing sufficient evidence to you in meeting your expectations.

Maintaining the landing zone's qualified state

After the landing zone is live, it must be maintained in a qualified state. Unless the operations are delegated to a partner, you typically create a Cloud Platform Operations and Maintenance SOP based on Section 6 of GAMP IT Infrastructure Control and Compliance.

According to GAMP, there are several areas where control must be shown, such as change management, configuration management, security management, and others. GAMP guidance also suggests that 'automatic tools' should be used whenever possible. The following sections cover these control areas and how AWS services can help with automation.

Change management

Change management processes control how changes to configuration items are made. These processes should include an assessment of the potential impact on the GxP applications supported by the landing zone. As mentioned earlier, all of the landing zone components are deployed using an automated pipeline. Therefore, after a change has been approved and committed in the source code repository tool, the pipeline is triggered and the change deployed. There will likely be multiple pipelines for the various parts that make up the landing zone.

The landing zone is made up of infrastructure and automation components. Now, using infrastructure as code, such as through AWS CloudFormation, there is no real difference between how these different components are deployed.

We recommend a continuous deployment methodology because it ensures changes are automatically built, tested, and deployed, with the goal of eliminating as many manual steps as possible. Continuous deployment seeks to eliminate the manual nature of this process and automate each step, allowing development teams to standardize the process and increase the efficiency with which they deploy code. In continuous deployment, an entire release process is a *pipeline* containing stages. AWS CodePipeline can be used along with AWS CodeCommit, AWS CodeBuild, and AWS CodeDeploy. For customers needing additional approval steps, CodePipeline also supports the inclusion of manual steps.

AWS Config provides continuous monitoring and assessment of AWS resource configurations. It helps track changes to resources over time and evaluate them against desired configurations, with ability to set up rules for compliance checking.

All changes to AWS services, either manual or automated, are logged by AWS CloudTrail.

AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services. This event history simplifies security analysis, resource change tracking, and troubleshooting. In addition, you can use CloudTrail to detect unusual activity in your AWS accounts. These capabilities help simplify operational analysis and troubleshooting.

Of course, customers also want to be alerted about any unauthorized and unintended changes. You can use a combination of AWS CloudTrail and AWS CloudWatch to detect unauthorized changes made to the production environment and even automate immediate remediation. Amazon CloudWatch is a monitoring service for AWS Cloud resources and alarms can be used to trigger responses to AWS CloudTrail events.

Amazon Event Bridge is a serverless event bus service that can detect changes in AWS services and trigger automated responses. It is useful for creating automated workflows around change management processes.

AWS Systems Manager Change Manager is a framework for requesting, approving, implementing, and reporting on operational changes to application configuration and infrastructure. It includes pre-approved change templates and automated approval workflows.

Configuration management

Going hand in hand with change management is configuration management. Configuration items (CIs) are the components that make up a system, and CIs should only be modified through the change management process.

Infrastructure as code (IaC) brings automation to the provisioning process through tools like AWS CloudFormation. Rather than relying on manually performed steps, both administrators and developers can instantiate infrastructure using configuration files. Infrastructure as code treats these configuration files as software code. These files can be used to produce a set of artifacts, namely the compute, storage, network, and application services that comprise an operating environment. Infrastructure as Code eliminates configuration drift through automation, thereby increasing the speed and agility of infrastructure deployments.

AWS Resource Groups and tagging let you organize your AWS landscape by applying tags at different levels of granularity. Tags allow you to label, collect, and organize resources and components within services.

The Tag Editor lets you manage tags across services and AWS Regions. Using this approach, you can globally manage all the application, business, data, and technology components of your target landscape.

A Resource Group is a collection of resources that share one or more tags. It can be used to create an enterprise architecture view of your IT landscape, consolidating AWS resources into a per-project (that is, the on-going programs that realize your target landscape), per-entity (that is, capabilities, roles, processes), and per-domain (that is, Business, Application, Data, Technology) view.

AWS Config is a service that lets you assess, audit, and evaluate the configurations of AWS resources. AWS Config continuously monitors and records your AWS resource configurations and lets you automate the evaluation of recorded configurations against desired configurations. With AWS Config, you can review changes in configurations and determine their overall compliance against the configurations specified in your internal guidelines. This enables you to simplify compliance auditing, security analysis, change management, and operational troubleshooting. In addition, AWS provides conformance packs for AWS Config to provide a general-purpose compliance framework designed to enable you to create security, operational or cost-optimization governance checks using managed or custom AWS Config rules and AWS Config remediation actions, including a conformance pack for 21 CFR 11.

AWS Systems Manager serves as a unified interface for infrastructure control across AWS and on-premises environments. It includes several key components: Parameter Store for secure configuration storage, State Manager for maintaining consistent configurations, automation for scripting common maintenance and deployment tasks, and Fleet Manager for remote node management. This integrated approach makes it a powerful tool for comprehensive configuration management across diverse environments.

AWS Service Catalog facilitates the creation and management of approved IT services and configurations catalogs. Through version-controlled templates and robust access controls, it enforces standardization across the organization. The service supports delegated administration and self-service provisioning within approved boundaries, while integrating with AWS Organizations for comprehensive multi-account governance.

AWS Resource Access Manager (AWS RAM) enables secure and controlled sharing of AWS resources across accounts. It maintains consistent configurations across shared resources while

providing centralized control over resource sharing permissions. Its integration with AWS Organizations makes it particularly valuable for managing resource sharing in complex organizational structures.

AWS Proton focuses on managing and deploying container and serverless applications using standardized templates. It maintains consistency through versioned service templates and infrastructure configurations, enables automated updates across multiple environments, and provides centralized configuration management specifically designed for microservices architectures. This makes it particularly valuable for organizations adopting modern application architectures.

Security management

AWS has defined a set of best practices for customers who are designing the security infrastructure and configuration for applications running in AWS.

Best Practices for Security, Identity, & Compliance provides security best practices that will help you define your information security management system (ISMS) and build a set of security policies and processes for your organization so you can protect your data and assets in the AWS Cloud.

Security, Identity, and Compliance on AWS provides an overview of different security topics such as identifying, categorizing and protecting your assets on AWS, managing access to AWS resources using accounts, users, and groups and suggesting ways you can secure your data, operating systems, applications and overall infrastructure in the cloud.

AWS provides you with an extensive set of tools to secure workloads in the cloud.

If you implement full automation, it could negate the need for anyone to have direct access to any environment beyond development. However, if a situation occurs that requires someone to access a production environment, they must explicitly request access, have the access reviewed and approved by the appropriate owner, and upon approval, obtain temporary access with the least privilege needed and only for the duration required. You should then track their activities through logging while they have access. See Temporary security credentials in IAM for further information.

Problem and incident management

With AWS, you get access to many tools and features to help you meet your problem and incident management objectives. These capabilities help you establish a configuration and security baseline that meets your objectives for your applications running in the cloud.

When a deviation from your baseline does occur (such as by a misconfiguration), you might need to respond and investigate. To successfully do so, you must understand the basic concepts of security incident response within your AWS environment and the issues you need to consider to prepare, educate, and train your cloud teams before security issues occur. It is important to know which controls and capabilities you can use, to review topical examples for resolving potential concerns, and to identify remediation methods that can be use automation and improve response speed.

Because security incident response can be a complex topic, we encourage you to start small, develop runbooks, use basic capabilities, and create an initial library of incident response mechanisms to iterate from and improve upon. This initial work should include teams that are not involved with security and should include your legal departments, so that they are better able to understand the impact that incident response (IR), and the choices they have made, have on your corporate goals.

For a comprehensive guide, see the AWS Security Incident Response Technical Guide.

Backup and restore

The ability to back up and restore is required for all validated applications. It is therefore a common capability that can be centralized as part of the regulated landing zone. Backup and restore should not be confused with archiving and retrieval.

For a cloud-based backup and restore capability, consider AWS Backup.

AWS Backup is a fully managed backup service that makes it easy to centralize and automate the backup of data across AWS services. Using AWS Backup, you can centrally configure backup policies and monitor backup activity for AWS resources, such as Amazon EBS volumes, Amazon EC2 instances, Amazon RDS databases, Amazon DynamoDB tables, Amazon EFS file systems, Amazon FSx file systems, and AWS Storage Gateway volumes. AWS Backup automates and consolidates backup tasks previously performed service-by-service, removing the need to create custom scripts and manual processes. With a few clicks in the AWS Backup console, you can create backup policies that automate backup schedules and retention management. AWS Backup provides a fully managed, policy-based backup solution, simplifying your backup management, enabling you to meet your business and regulatory backup compliance requirements.

Disaster recovery

In traditional on-premises situations, disaster recovery (DR) involves a separate data center located a certain distance from the primary data center. This separate data center only exists in

case of a complete disaster impacting the primary data center. Often the infrastructure at the DR site sits idle, or at best hosts pre-production instances of applications thus running the risk of it being out-of-sync with production. With the advent of the cloud, DR is now much easier and cheaper.

The AWS global infrastructure is built around AWS Regions and Availability Zones (AZ). AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected by low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

With Availability Zones, it is easy to create a multi-AZ architecture capable of withstanding a complete failure of one or more zones. For even more resilience, multiple AWS Regions can be used. With the use of infrastructure as code, the infrastructure and applications in a DR Region do not need to run all of the time. In case of a disaster, the entire application stack can be deployed into another Region. The only components that must run all the time are those keeping the data repositories in sync.

AWS Elastic Disaster Recovery (AWS DRS), formerly known as CloudEndure Disaster Recovery, is a highly automated disaster recovery service that enables quick and reliable recovery of physical, virtual, and cloud-based servers into AWS. The service continuously replicates your servers (including operating system, system state configuration, databases, applications, and files) into a low-cost staging area in your target AWS account and preferred Region. During an actual disaster or disaster recovery drill, AWS DRS can automatically convert your servers to run natively on AWS, typically within minutes of initiating the recovery, with minimal data loss and near-zero downtime.

Helpdesk

The helpdesk serves as a critical component in maintaining the validated state of GxP systems by functioning as the primary point of contact for incident management, change requests, and user support while ensuring regulatory compliance. It must operate under documented procedures that clearly define incident classification, escalation protocols, and resolution workflows, with established severity levels and response time requirements specifically tailored for GxP-critical systems. The helpdesk staff must be properly trained to understand the impact of changes and maintain a comprehensive record of all actions taken in compliance with predefined procedures. Communication protocols play a vital role in helpdesk operations, requiring clear channels with users and stakeholders, regular status updates on critical

incidents, and well-defined escalation procedures for high-priority issues. The helpdesk must maintain strong integration with the change control board and quality assurance team, while implementing effective feedback mechanisms for system improvements. This ensures that all stakeholders are properly informed and involved in maintaining the validated state of GxP systems.

Regular performance evaluation through established metrics, quality reviews, and compliance assessments helps ensure the helpdesk maintains its qualified state and continues to support the validated environment effectively. This includes monitoring response times, resolution rates, documentation quality, and compliance with established procedures. The helpdesk must maintain detailed records of all activities, ensuring traceability and providing evidence of maintained system validation status for regulatory inspections and audits.

Amazon Connect is a cloud-based contact center service that can be used to create a comprehensive helpdesk solution. The service offers an omnichannel experience, supporting voice, chat, and task management through a unified interface, while enabling seamless integration with existing CRM systems and other business applications. It provides intelligent routing, real-time and historical analytics, and integration capabilities with other AWS services and external systems. Amazon Connect offers features like call recording, queue management, and interactive voice response (IVR), making it suitable for organizations requiring a full-featured helpdesk phone system with scalability.

Performance monitoring

Performance monitoring in GxP-regulated environments requires a systematic approach to continuously track, analyze, and maintain system performance while ensuring compliance with regulatory requirements. Organizations must establish comprehensive monitoring programs that include defined performance metrics, thresholds, and alert mechanisms for critical system parameters that could impact product quality or data integrity. The monitoring system should provide real-time visibility into system performance, with automated alerts for deviations from established parameters, and maintain detailed audit trails of all monitoring activities and responses to performance issues.

Amazon CloudWatch serves as the primary monitoring service, collecting and tracking metrics, logs, and events across AWS resources and applications. It provides real-time monitoring capabilities with customizable dashboards and alerting mechanisms, enabling comprehensive performance tracking and automated responses to performance issues. The detailed metrics and analytics capabilities of CloudWatch help organizations understand system behavior and identify potential performance bottlenecks before they impact operations.

Amazon CloudWatch Application Signals provides application performance monitoring for distributed applications, offering end-to-end visibility into request traces, service dependencies, and latency issues across microservices architectures. Supporting OpenTelemetry as the instrumentation standard, Application Signals enables service maps, automated SLO tracking, and detailed performance optimization of complex distributed systems.

Amazon Managed Service for Prometheus and AWS Managed Grafana provide powerful monitoring and visualization capabilities for containerized environments. These services enable collection of high-cardinality metrics and creation of rich dashboards for monitoring container and microservice performance, while offering scalable, managed solutions for metric storage and visualization.

Periodic review

While our objective is to establish continuous compliance, while manual processes still exist periodic reviews are required and represent a critical quality assurance process designed to ensure systems maintain their qualified state and continue to operate effectively within regulatory requirements. This systematic assessment involves comprehensive evaluation of system performance, security measures, compliance status, and documentation to identify any gaps or areas requiring improvement. The process helps organizations maintain regulatory compliance while ensuring systems remain fit for their intended use through regular assessment of operational effectiveness and control mechanisms. Periodic reviews must be conducted according to established schedules based on system criticality and risk assessment, with all activities documented to demonstrate ongoing compliance.

On top of AWS Config, Amazon CloudWatch, and Amazon Inspector mentioned previously, AWS Audit Manager simplifies the review process by continuously auditing AWS usage and simplifying risk and compliance assessments. It provides pre-built frameworks for common industry standards and maintains evidence collection and audit documentation. Amazon Macie enhances security reviews by automatically discovering and protecting sensitive data in AWS, providing ongoing monitoring and risk assessment of data security, and generating detailed reports for security and compliance reviews.

AWS Trusted Advisor offers real-time guidance for following AWS best practices, providing recommendations across cost optimization, security, fault tolerance, and performance domains. Systems Manager supplements this by offering tools for viewing and controlling infrastructure state, providing operational insights and patch compliance status, and enabling automated compliance reporting and system maintenance. AWS Organizations facilitates management of

multiple AWS accounts centrally for governance, enabling policy-based management and compliance across accounts.

Qualifying building blocks

Customers frequently want to know how AWS gives developers freedom to use any AWS service while still maintaining regulatory compliance and fast development. To address this problem, you can use technology, but this also involves changes in process design to move away from blocking steps and towards guardrails. The changes required to your processes and IT operating model are beyond the scope of this whitepaper. However, we cover the core steps of a supporting process to qualify building blocks which is one tactic for maintaining regulatory compliance more efficiently.

The infrastructure building block concept as defined by GAMP is an approach to qualify individual components or combinations of components which can then be put together to build out the IT infrastructure. In GxP environments, building blocks are pre-defined, reusable infrastructure components (such as an Amazon S3 configuration or an Amazon RDS deployment pattern) that have been documented, tested, and approved for use in regulated workloads. The approach is applicable to AWS services.

Qualifying a building block means demonstrating once—through documented evidence—that the component consistently meets its intended use and complies with applicable security and data integrity requirements. This allows organizations to reuse it across multiple workloads without repeating full infrastructure qualification each time. In the AWS shared responsibility model:

- AWS already qualifies the standard, documented features of its services through its own controlled processes, audits, and certifications (for example, SOC 2, ISO 27001, HITRUST).
- Customers do not need to re-qualify those base services.
- Customers must qualify their configuration of the service to meet their specific intended use, risk profile, and regulatory needs.

Service approval

Before qualifying a building block, organizations determine whether the underlying AWS service is suitable for regulated use. Customers often consider multiple regulations when approving a service for use by development teams. For example, you might allow all services to be used in

sandbox accounts but restrict the services in an account to only HIPAA-eligible services if the application is subject to HIPAA regulations.

Service approval is implemented through the use of AWS Organizations and Service Control Policies.

You could take this approach to allow services to be used as part of GxP relevant applications. For example, a combination of ISO, PCI, SOC, and HIPAA-eligibility might provide sufficient confidence. Sometimes, customers want to implement automated controls over the approved service as described in Approving AWS services for GxP workloads.

After a service meets compliance coverage, customers evaluate the risk profile for data integrity and security.

1. **Security layer**: Security controls prevent unauthorized access or modification:
    - Encryption at rest (AWS KMS)
    - Encryption in transit (TLS 1.2+)
    - IAM least-privilege access
    - VPC network segmentation
2. **Data integrity layer**: Integrity controls ensure the system complies with ALCOA+ principles:
    - Attributable, for example, CloudTrail logs and Amazon S3 object versioning to track who made changes
    - Legible store in readable, non-proprietary formats
    - Contemporaneous, for example, automatic timestamps on data creation (CloudTrail, Amazon RDS logs)
    - Original, for example, S3 Object Lock in compliance mode to preserve original data
    - Accurate: Validation rules in application and database layers
    - Complete: Ensure all transactions are logged and retained
    - Consistent: Enforce standard date/time formats and naming conventions
    - Enduring: For example, durable storage (S3 with replication and retention policies)
    - Available: For example, retrieval within defined SLAs; multi-AZ or cross-Region replication
    - Traceable (common extension): Link each record to its source via audit logs

Example: For Amazon S3, qualifying the building block would focus on enabling versioning, Object Lock, KMS encryption, logging, and SCP rules blocking public access—rather than revalidating the durability or encryption algorithms of S3.

You might prefer to follow a more rigorous qualification process like the following building block qualification.

Building block qualification

The qualification of AWS service *building blocks* follows a process based on the GAMP IT Infrastructure Control and Compliance guidance documents *Infrastructure Building Block Concept* (Section 9 / Appendix 2 of GAMP IT).

> According to EU GMP, the definition of qualification is: "Action of proving that any equipment works correctly and actually leads to the expected results." The equipment also needs to continue to lead to the expected results over its lifetime.

In other words, your process should show that the building block works as intended and is kept under control throughout its operational life. There will be written procedures in place and, when executed, records will show that the activities actually occurred. Also, the staff operating the services need to be appropriately trained. This process is often described in an SOP describing the overall qualification and commissioning strategy, the scope, roles, and responsibilities, in addition to a deliverables list and any good engineering practices that will be followed to satisfy qualification and commissioning requirements.

With the number of AWS services, it can be difficult for you to qualify all AWS services at once. An iterative and risk-based approach is recommended where services are qualified in priority order. Initial prioritization will take into account the needs of the first applications moving to cloud and then the prioritization can be reassessed as demand for cloud services increases.

Design stage

Requirements

The first activity is to consider the requirements for the building block. Define functional and non-functional needs, explicitly covering both security and ALCOA+ integrity requirements. One approach is to look at the service API definition. Each AWS service has a clearly documented API describing the entire functionality of that service. Many service APIs are extensive and support

some advanced functionality. However, not all of this advanced functionality might be required initially, so any existing business use cases can be considered to help refine the scope.

For example, when noting Amazon S3 requirements, you include the core functionality of creating and deleting buckets and the ability to put, get, and delete objects. However, you might not include the lifecycle policy functionality because this functionality is not yet needed. These requirements are captured in the building block requirements specification or requirements repository.

It's also important to consider non-functional requirements. To ensure suitability of a service you can look at the services SLA and limits.

Gap analysis

Where application requirements already exist, in the same way you can restrict the scope, you can also identify any gaps. Either the gap can be addressed by including more functionality for the building block, like bringing the Amazon S3 bucket Lifecycle functionality into scope, or the service is not suitable for satisfying the requirements and an alternate building block should be used.

If no other service seems to meet the requirements, you can develop a custom service or make a feature request to AWS for service enhancement.

Risk assessment

Infrastructure is qualified to ensure reliability, security, and business continuity for the validated applications running on it. These three dimensions are usually included as part of any risk assessment. The published AWS SLA provides confidence in AWS services reliability. Data regarding the current status of the service plus historical adherence to SLAs is available from the AWS service health dashboard. For confidence in security, the AWS certifications can be checked for the relevant service. For business continuity, AWS builds to guard against outages and incidents, and accounts for them in the design of AWS services, so when disruptions do occur, their impact on customers and the continuity of services is as minimal as possible.

This step is also not only for GxP qualification purposes. The risk assessment should include any additional checks for other regulations such as HIPAA.

When assessing the risks for a cloud service, it's important to consider the relationship to other building blocks. For example, an Amazon RDS database might have a relationship to the Amazon VPC building block because you decided a database is only allowed to exist within the private subnet of a VPC. Therefore, the VPC is taking care of many of the risks around access control. These dependencies will be captured in the risk assessment and then focus on additional risks

specific to the service, or residual risks which cannot be catered for by the surrounding production environment.

Each cloud service building block goes through a risk assessment that identifies a list of risks. For each identified risk, a mitigation plan is created. The mitigation plan can influence one or more of the following components:

- Service control policy
- Technical design or infrastructure as code template
- Monitoring and alerting of automated compliance controls

A risk can be mitigated by using service control policies (SCPs) where a service or specific operation is deemed too risky and its use explicitly denied through such a policy. For example, you can use an SCP to restrict the deletion of an Amazon S3 object through the AWS Management Console. Another option is to control service usage through the technical design of an approved infrastructure as code (IaC) template where certain configuration parameters are restricted or parameterized. For example, you can use an AWS CloudFormation template to always configure an Amazon S3 bucket as private. Finally, you can define rules that feed into monitoring and alerting. For example, if the policy states S3 buckets cannot be public, but this configuration is not enforced in the infrastructure template, then the infrastructure can be monitored for any public S3 buckets. When an S3 bucket is configured as public, an alert triggers remediation, such as immediately changing a bucket to private.

Technical design

In response to the specified requirements and risks, an architecture design specification will be created by a cloud infrastructure architect describing the logical service building block design and traceability from risk or requirement to the design. This design specification will, among other things, describe the capabilities of the building block to the end users and application development teams.

Design review

To verify that the proposed design is suitable for the intended purpose within the surrounding IT infrastructure design, a design review can be performed by a suitably trained person as a final check.

Construction stage

The logical design can be captured in a document, but the physical design is captured in an infrastructure as code (IaC) template, like an AWS CloudFormation template. This IaC template is always used to deploy an instance of the building block ensuring consistency. For one approach, see Automating GxP compliance in the cloud: Best practices and architecture guidelines.

The IaC template will use parameters to deal with workload variances. As part of the design effort it will be determined, often by IT Quality and Security, which parameters affect the risk profile of the service and so should be controlled and which parameters can be set by the user. For example, the name of a database can be set by the template user and generally does not affect the risk profile of a database service. However, any parameter controlling encryption does affect the risk profile and therefore is fixed in the template and not changeable by the template user.

The template is a text file that can be edited. However, the rules expressed in the template are also automated within the surrounding monitoring and alerting. For example, the rule stating that the encryption setting on a database must be set can be checked by automated rules. Therefore, a developer might override the encryption setting in the development environment, but that change isn't allowed to progress to a validated environment or beyond.

At this point, automated test scripts can be prepared for executing during the qualification step to generate test evidence. The author of the automated tests must be suitably trained and a separate and suitably trained person performs a code review and/or random testing of the automated tests to ensure the quality level.

The automated tests ensure the building block initially functions as expected. These tests can be run again to ensure the building block continues to function as expected, especially after any change. However, to ensure nothing has changed after the building block is in production, you should identify and create automated controls. Using the Amazon S3 example again, all buckets should be private. If a public bucket is detected, it can be switched back to private and an alert raised and notification sent. You can also determine the individual that created the S3 bucket and revoke their permissions.

The final part of construction is the authoring and approval of any additional guidance and operations manuals. For example, how to recover a database would be included in the operations manual of an Amazon RDS building block.

Qualification and commissioning stage

It's important to note that infrastructure is deployed in the same way for every building block, that is, through AWS CloudFormation using an infrastructure as code template. Therefore, there is usually no need for building block specific installation instructions. Also, you can be confident that every deployment is done according to specification and has the correct configuration.

Automated testing

If you want to generate test evidence, you can demonstrate that the functional requirements are fulfilled and that all identified risks have been mitigated—thus indicating the building block is fit for its intended use—through the execution of the automated tests created during construction. The output of these automated tests is stored into a secure repository and can be used as test evidence.

This automation deploys the building block template into a test environment, executes the automated tests, captures the evidence, and then destroys the stack again avoiding any ongoing costs.

Testing might only make sense in combination with other building blocks. For example, the testing of a NAT gateway can only be done within an existing VPC. One alternative is to test within the context of standard archetypes, that is, a complete stack for a typical application architecture.

Handover to operations stage

The handover stage ensures that the cloud operation team is familiar with the new building block and is trained in any service specific operations. Once the operations team approves the new building block, the service can be approved by changing a Service Control Policy (SCP). The infrastructure as code template can be made available for use by adding it into the AWS Service Catalog or other secure template repository.

If the response to a risk was a SCP or monitoring rule change, then the process to deploy those changes is triggered at this stage.

Computer Systems Assurance

Computer Software Assurance (CSA) in GxP environments are guidelines referring to a modern, risk-based approach to software validation. These guidelines are designed to ensure product quality, patient safety, and data integrity across all stages of the product lifecycle. Historically,

the industry has relied on Computer System Validation (CSV) as a compliance framework, particularly in accordance with the U.S. Food and Drug Administration (FDA) requirements outlined in 21 CFR Part 11 for electronic records and signatures and 21 CFR Part 820 for quality system regulations in medical device manufacturing. However, the FDA's guidance on Computer Software Assurance for Production and Quality System Software, published in September 2025, marks a significant shift toward a more flexible and risk-based approach to software validation.

The FDA's guidance proposes CSA as a modernized alternative to traditional CSV, focusing on risk-based assurance rather than exhaustive documentation. This shift reflects the FDA's intent to streamline validation processes without compromising regulatory compliance or patient safety. The guidance emphasizes that manufacturers should focus on ensuring that software functions as intended, particularly when used within production or quality systems as defined under 21 CFR 820.70(i). This regulation mandates that manufacturers validate software for its intended use when it is part of the production or quality system, highlighting the importance of functionality in maintaining compliance.

A core principle of CSA, as outlined in the FDA guidance, is to shift effort from documentation to *critical thinking*, reasoned decisions, and testing that assure quality, not only paper trails. Assurance activities should be proportional to the risk associated with the software's failure. This requires manufacturers to assess the software's intended use and the potential consequences if it fails. If a failure could directly affect patient safety, product quality, or data integrity, it demands more rigorous assurance activities. For example, software used to control critical production parameters—such as temperature, pressure, or humidity—poses a high risk, because failure could compromise product safety. Such systems, under 21 CFR Part 820.70, must undergo comprehensive validation. Conversely, administrative software, like scheduling tools or email systems, falls outside the scope of 21 CFR 820.70(i) and typically requires minimal validation, as their failure would not directly impact product quality or patient safety.

In this section we will apply some critical thinking to find alternate ways of achieving some of our control objectives but in a much more efficient way, using current cloud engineering good practices.

Installation verification

Traditionally, this meant creating an installation qualification protocol and having various system administrators execute their parts, sign, date, and so on. This worked well to demonstrate control over a manual installation process. However, good engineering practices now automate this whole procedure.

This verification and documentation demonstrate a system is installed according to a pre-approved specification. Verification is achieved through testing that shows that the installation and configuration of software and hardware is correct.

First, let's consider the *pre-approved specification*. It is an AWS best practice to define the required infrastructure through the use of an infrastructure as code (IaC) template or code. This template describes the required resources and their configuration. The configuration might change slightly between environments (QA, UAT, PRE-PROD, and PROD) and Regions, but the core IaC template does not. This flexibility is achieved through template parameters. This template is then used by AWS CloudFormation to provision the resources.

These templates are controlled in a similar way as source code. By storing them in a source code repository it enables us to version the template and keep a complete history of its evolution over time.

Another key part of that phrase is *pre-approved*. There are many ways that a customer can handle the approval. For example, a Jira workflow or a pull request approval in their source code repository. Whatever the method it will be vetted and approved by the customers Quality IT or Compliance team. The net result is a specific version of the template in the source code repository being recorded as approved.

So, you now have a pre-approved specification describing the resources you want deployed. An automated pipeline is then triggered to deploy the resources. You then need to look at the next requirement, which is to demonstrate the installation was correct. This can be done at various points in the pipeline.

Testing application code during development is a common practice. If using AWS CDK you can perform local testing during the development cycle of your infrastructure code.

CloudFormation Hooks is a feature that helps ensure that your CloudFormation resources, stacks, and change sets comply with your organization's security, operational, and cost optimization best practices. CloudFormation Hooks can also ensure this same level of compliance for your AWS Cloud Control API resources. With CloudFormation Hooks, you can provide code that proactively inspects the configuration of your AWS resources before provisioning. If non-compliant resources are found, CloudFormation either fails the operation and prevents the resources from being provisioned or emits a warning and allows the provisioning operation to continue.

CloudFormation is the service that provisions your resources. Its sole purpose is to help you model and provision AWS infrastructure as code (IaC). It allows you to manage, provision, and

update related AWS and third-party resources in an orderly and predictable fashion. If there is no error during provisioning, then confidence is high the resources were provisioned as expected. If there is an error, then manual intervention is required but we have improved process efficiency and positioned manual review *by exception* only.

Post-deployment, automated tests can be included in the pipeline to test provisioned resources. These tests can generate evidence that the resources are configured correctly and functioning as expected. In addition, it is a best practice to use AWS Config Rules which evaluate the configuration settings of your provisioned AWS resources. AWS Config rules continuously evaluate your AWS resource configurations for compliance with desired settings helping ensure resource configurations *remain* compliant. A rule can run when AWS Config detects a configuration change to an AWS resource or at a periodic frequency that you choose.

With all these checks at various points in the automated pipeline, and all the evidence they generate, it could be argued that traditional manual testing and verification reporting is no longer needed. Should a report still be needed to adhere to SOPs, it can be generated from the data generated by the automation.

Inspection readiness through IT records

Under the CSA framework the focus shifts from documentation toward collecting only the evidence necessary to demonstrate that the software is fit for its intended use. The FDA encourages using *electronic records, audit trails, and system logs* to capture objective evidence efficiently.

Development tools generate extensive compliance documentation during the engineering process. Version control systems maintain complete histories of every code change with developer attribution and timestamps. Continuous integration pipelines capture test results in machine-readable formats that provide better evidence than manual screenshots. Deployment automation documents every change from development to production with full traceability. These tools create richer compliance documentation by preserving the actual evolution of software rather than collecting signatures on point-in-time snapshots.

The basic premise is to have access to all the data generated by our IT tooling. Rather than generating separate documents as evidence, use the IT records from the tooling. This data can be gathered into a data lake or joined virtually. The main requirement being that nobody should be able to change the data—that it is immutable—because it forms the basis of your regulatory evidence. This data effectively enables you to create a digital twin of your software development lifecycle (SDLC). After this data is available it can be queried to answer quality

related questions or demonstrate control effectiveness. For example, data from requirements management and test management tooling can be combined to create a traceability matrix showing requirements have been tested and those tests are passing. The data can answer inspection questions like what features were included in a specific release or demonstrate that server patching is happening and that all servers are on the latest release.

Along with answering questions, *dashboards* can be created to show the current compliance status at all times, helping move from a static point-in-time view of compliance to a continuous compliance stance. It also helps show operational healthiness, for example, team velocity or end-to-end productivity.

The next logical progression is that this data foundation enables agentic AI quality management. For example, if a defect is detected or reported it can trigger an agentic workflow to determine the root cause and suggest remediation steps. It could even execute those remediation steps.

Good technical documentation mechanisms

Documentation remains a crucial aspect of GxP compliance. The documentation should include the intended use of the software, risk assessment results, assurance activities performed, and any issues discovered during testing, along with their resolutions.

During a typical development cycle there is a risk that documentation gets out of sync with reality. One mechanism that has proven effective to combat this, for technical documentation usually maintained by the technical team, is the use of markup stored in the code repository alongside the code (application code, infrastructure as code, LZA template, and so on.). As the code is updated, the markup is also updated, and this is confirmed during code review or approval. As the code is built and deployed, the documentation is also deployed to a wiki. This technique has the advantage that documentation is stored in the same place as the code, updated at the same time, and maintained using the same development tools as the code, reducing the risk of documentation getting out of sync with reality.

Development tools generate extensive compliance evidence during the engineering process. Version control systems maintain complete histories of every code change with developer attribution and timestamps. Continuous integration pipelines capture test results in machine-readable formats that provide better evidence than manual screenshots. Deployment automation documents every change from development to production with full traceability. These tools create richer compliance evidence by preserving the actual evolution of software rather than collecting signatures on point-in-time snapshots.

Validation during cloud migration

One important point that might be covered in a cloud strategy is the overarching approach to computer system validation (CSV) during migration. If you are embarking on a migration effort, part of the analysis of the application portfolio will be to identify archetypes, or groups of applications with similar architectures. A single runbook can be developed and then repeated for each of the applications in the group, speeding up migration.

At this point, if the applications are GxP relevant, the CSV or migration strategy can also be defined for the archetype and repeated for each application.

For further information about the migration of GxP workloads to the AWS cloud, see this series of blog posts.

Conclusion

If you are a Life Science customer with GxP obligations, you retain accountability and responsibility for your use of AWS products, including the applications and virtualized infrastructure you develop, validate and operate using AWS Products. Using the recommendations in this whitepaper, you can evaluate your use of AWS products within the context of your quality system and consider strategies for implementing the controls required for GxP compliance, as a component of your regulated products and systems.

Contributors

Contributors to this document include:

- Sylva Krizan PhD, Security Assurance, AWS Global Healthcare and Life Sciences
- Rye Robinson, Solutions Architect, AWS Global Healthcare and Life Sciences
- Ian Sutcliffe, Principal Solutions Architect, AWS Global Healthcare and Life Sciences
- Denny Daugherty, Principal Technical Account Manager, AWS Enterprise Support
- Koushik Das, Sr. Technical Account Manager, AWS Enterprise Support
- Subrat Bora, Sr. Solutions Architect, AWS Global Healthcare and Life Sciences

Further reading

For additional information, see:

- AWS Compliance
- Healthcare & Life Sciences on AWS

Document revisions

Date	Description
March 2026	Updated to reflect modern Computer Systems Assurance (CSA) approach, expanded layered architecture guidance with enhanced Landing Zone qualification methodology, updated AWS service portfolio and certifications, and strengthened Infrastructure as Code and continuous compliance monitoring practices
March 2021	Updated to include more elements of AWS Quality System Information and updated guidance on customer approach to GxP compliance on AWS
January 2016	First publication

Appendix: 21 CFR 11 Controls – Shared responsibility for use with AWS services

Applicability of 21 CFR 11 to regulated medical products and GxP systems are the responsibility of the customer, as determined by the intended use of the systems or products. AWS has mapped some of these requirements based on the AWS Shared Responsibility Model, however, customers are responsible for meeting their own regulatory obligations.

Below, we have identified each subpart of 21 CFR 11 and clarified areas where AWS services and operations and the customer share responsibility in order to meet 21 CFR 11 requirements.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
11.10 Controls for closed systems. Persons who use closed systems to create, modify, maintain, or transmit electronic records shall employ procedures and controls designed to ensure the authenticity, integrity, and, when appropriate, the confidentiality of electronic records, and to ensure that the signer cannot readily repudiate the signed record as not genuine. Such procedures and controls shall include the following:		

11.10(a) Validation of systems to ensure accuracy, reliability, consistent intended performance, and the ability to discern invalid or altered records.	AWS Services are built and tested to conform to IT industry standards, including SOC, ISO, PCI, and others https://aws.amazon.com/compliance/programs/. AWS compliance programs and reports provide objective evidence that AWS has implemented several key controls, including, but not limited to: Control over the installation and operation of AWS product components, including both software components and hardware components; Control over product changes and configuration management; Risk management program; Management review, planning, and operational monitoring; Security management of information availability, integrity, and confidentiality; and Data protection controls including mechanisms for data backup, restore and archiving. All purchased materials and services intended for use in production processes are documented, and documentation is reviewed and approved prior to use and verified to be in conformance with the specifications. Final inspection and testing is performed on AWS Services prior to their release to general availability. The final service release review procedure includes a verification that all acceptance data is present and that all product requirements were met. Once in production, AWS Services undergo continuous performance monitoring. In addition, AWS's significant customer base, authorization for use by government agencies, and	AWS products are basic building blocks that allow you to create private, virtualized infrastructure environments for your custom software applications and commercial-off-the-shelf applications. In this way, you remain responsible for enabling (i.e. installing), configuring, and operating AWS products to meet your data-, application-, and industry-specific needs like GxP software validation and GxP infrastructure qualification as well as validation to support 21 CFR Part 11 requirements. AWS products are, however, unlike traditional infrastructure software products in that they are highly automatable, allowing you to programmatically create qualified infrastructure via version controlled JSON[2] scripts instead of manually-executed paper protocols, where applicable. This automation capability not only reduces effort, it increases control and consistency of the infrastructure environment such that continuous qualification[3] is possible. Installation qualification of AWS Services into your environment, operational and performance qualification (IQ/OQ/PQ) are your responsibility, as are the validation activities to demonstrate that systems with GxP workloads managing electronic records are appropriate for the intended use and meet regulatory requirements.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
	recognition by industry analysts as a leading cloud services provider are further evidence of AWS products delivering their documented functionality https://aws.amazon.com/documentation/. Relevant SOC2 Common Criteria: CC1.2, CC1.4, CC3.2, CC7.1, CC7.2, CC7.3, CC7.4	
11.10(b) The ability to generate accurate and complete copies of records in both human readable and electronic form suitable for inspection, review, and copying by the agency. Persons should contact the agency if there are any questions regarding the ability of the agency to perform such review and copying of the electronic records.	Controls are implemented subject to industry best practices in order to ensure services provide complete and accurate outputs with expected performance committed to in SLAs.; Relevant SOC2 Common Criteria: A1.1	AWS has a series of Security Best Practices (https://aws.amazon.com/security/security-resources/) and additional resources you may reference to help protect data hosted within AWS. You ultimately will verify that electronic records are accurate and complete within your AWS environment, and determine the format by which data is human and/or machine readable and is suitable for inspection by regulators, per the regulatory requirements.

(c) Protection of records to enable their accurate and ready retrieval throughout the records retention period.	Controls are implemented subject to industry best practices in order to ensure services provide complete and accurate outputs with expected performance committed to in SLAs.; Relevant SOC2 Common Criteria: A1.1 AWS has identified critical system components required to maintain the availability of our system and recover service in the event of outage. Critical system components are backed up across multiple, isolated locations known as Availability Zones, and backups are maintained. Each Availability Zone is engineered to operate independently with high reliability. Backups of critical AWS system components are monitored for successful replication across multiple Availability Zones. Refer to the AWS SOC 2 Report CC A1.2. The AWS Resiliency Program encompasses the processes and procedures by which AWS identifies, responds to, and recovers from a major event or incident within our environment. This program builds upon the traditional approach of addressing contingency management, which incorporates elements of business continuity and disaster recovery plans and expands this to consider critical elements of proactive risk mitigation strategies such as engineering physically separate Availability Zones (AZs) and continuous infrastructure capacity planning. AWS service resiliency plans are periodically reviewed by members of the Senior Executive management team and the Audit Committee of the Board of Directors. The AWS Business Continuity Plan outlines measures to avoid and lessen environmental disruptions. It includes operational details about	AWS has a series of Security Best Practices (https://aws.amazon.com/security/security-resources/) and additional resources you may reference to help protect your data hosted within AWS. You are responsible for implementation of appropriate security configurations for your environment to protect data integrity as well as ensure data and resources are only retrieved by appropriate permission. You are also responsible for creating and testing record retention policies as well as backup and recovery processes. You are responsible for properly configuring and using the Service Offerings and taking your own steps to maintain appropriate security, protection, and backup of your Customer Content, which may include the use of encryption technology (to protect your content from unauthorized access) and routine archiving. Using Service Offerings such as Amazon S3, Amazon Glacier, and Amazon RDS, in combination with replication and high availability configurations, AWS's broad range of storage solutions for backup and recovery are designed for many customer workloads. https://aws.amazon.com/backup-recovery/ AWS Services provide you with capabilities to design for resiliency and maintain business continuity, including the utilization of frequent server instance back-ups, data redundancy replication, and the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. You need to architect your AWS usage to take advantage of multiple regions and availability zones. Distributing applications across multiple availability zones provides the ability to

21 CFR Subpart	AWS Responsibility	Customer Responsibility
	steps to take before, during, and after an event. The Business Continuity Plan is supported by testing that includes simulations of different scenarios. During and after testing, AWS documents people and process performance, corrective actions, and lessons learned with the aim of continuous improvement. AWS data centers are designed to anticipate and tolerate failure while maintaining service levels. In case of failure, automated processes move traffic away from the affected area. Core applications are deployed to an N+1 standard, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites. Refer to the AWS SOC 2 Report CC3.1, CC3.2, A1.2, A1.3.	remain resilient in the face of most failure modes, including natural disasters or system failures. The AWS cloud supports many popular disaster recovery (DR) architectures, from “pilot light” environments that are ready to scale up at a moment’s notice to “hot standby” environments that enable rapid failover. You are responsible for DR planning and testing.

(d) Limiting system access to authorized individuals.	AWS implements both physical and logical security controls. Physical access to all AWS data centers housing IT infrastructure components is restricted to authorized data center employees, vendors, and contractors who require access in order to execute their jobs. Employees requiring data center access must first apply for access and provide a valid business justification. These requests are granted based on the principle of least privilege, where requests must specify to which layer of the data center the individual needs access, and are time-bound. Requests are reviewed and approved by authorized personnel, and access is revoked after the requested time expires. Once granted admittance, individuals are restricted to areas specified in their permissions. Access to data centers is regularly reviewed. Access is automatically revoked when an employee's record is terminated in Amazon's HR system. In addition, when an employee or contractor's access expires in accordance with the approved request duration, his or her access is revoked, even if he or she continues to be an employee of Amazon. AWS restricts logical user access privileges to the internal Amazon network based on business need and job responsibilities. AWS employs the concept of least privilege, allowing only the necessary access for users to accomplish their job function. New user accounts are created to have minimal access. User access to AWS systems requires approval from the authorized personnel, and validation of the active user. Access privileges to AWS systems are reviewed on a regular basis.	AWS provides you with the ability to configure and use the AWS service offerings in order to maintain appropriate security, protection, and backup of content, which may include the use of encryption technology to protect your content from unauthorized access. You maintain full control and responsibility for establishing and verifying configuration of access to your data and AWS accounts, as well as periodic review of access to data and resources. Using AWS Identity and Access Management (IAM), a web service that allows you to securely control access to AWS resources, you must control who can access and use your data and AWS resources (authentication) and what data and resources they can use and in what ways (authorization). IAM is a feature of all AWS accounts offered at no additional charge. You will be charged only for use of other AWS Services by your users, https://aws.amazon.com/iam/. IAM Best Practices can be found here: http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html. Maintaining physical access to your facilities and assets is solely your responsibility.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
	When an employee no longer requires these privileges, his or her access is revoked. Refer to the AWS SOC 2 Report C1.2, C1.3, and CC6.1-6.6 to verify the AWS physical and logical security controls.	

(e) Use of secure, computer-generated, time-stamped audit trails to independently record the date and time of operator entries and actions that create, modify, or delete electronic records. Record changes shall not obscure previously recorded information. Such audit trail documentation shall be retained for a period at least as long as that required for the subject electronic records and shall be available for agency review and copying.	AWS maintains centralized repositories that provide core log archival functionality available for internal use by AWS service teams. Leveraging S3 for high scalability, durability, and availability, it allows service teams to collect, archive, and view service logs in a central log service. Production hosts at AWS are equipped with logging for security purposes. This service logs all human actions on hosts, including logons, failed logon attempts, and logoffs. These logs are stored and accessible by AWS security teams for root cause analysis in the event of a suspected security incident. Logs for a given host are also available to the team that owns that host. A frontend log analysis tool is available to service teams to search their logs for operational and security analysis. Processes are implemented to protect logs and audit tools from unauthorized access, modification, and deletion. Refer to the AWS SOC 2 Report CC5.1, CC7.1	Verification and implementation of audit trails, as well as back up and retention procedures of your electronic records are your responsibility. AWS provides you with the ability to properly configure and use the Service Offerings in order to maintain appropriate audit trail and logging of data access, use and modification (including prohibiting disablement of audit trail functionality). Logs within your control (described below) can be used for monitoring and detection of unauthorized changes to your data. Using Service Offerings such as AWS CloudTrail, AWS CloudWatch Logs, and VPC Flow Logs, you can monitor your AWS data operations in the cloud by getting a history of AWS API calls for your account, including API calls made via the AWS Management Console, the AWS SDKs, the command line tools, and higher-level AWS Services. You can also identify which users and accounts called AWS APIs for services that support AWS CloudTrail, the source IP address the calls were made from, and when the calls occurred. You can integrate AWS CloudTrail into applications using the API, automate trail creation for your organization, check the status of your trails, and control how administrators turn logging services on and off. AWS CloudTrail records two types of events: (1) Management Events: Represent standard API activity for AWS Services. For example, AWS CloudTrail delivers management events for API calls such as launching EC2 instances or creating S3 buckets. (2) Data Events: Represent S3 object-level API activity, such as Get, Put, Delete and List actions.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
		https://aws.amazon.com/cloudtrail/ https://aws.amazon.com/documentation/cloudtrail/ http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html
(f) Use of operational system checks to enforce permitted sequencing of steps and events, as appropriate.	Not applicable to AWS – this requirement only applies to the customer's system.	You are responsible for configuring, establishing and verifying enforcement of permitted sequencing of steps and events within the regulated environment.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
(g) Use of authority checks to ensure that only authorized individuals can use the system, electronically sign a record, access the operation or computer system input or output device, alter a record, or perform the operation at hand.	Not applicable to AWS – this requirement only applies to the customer's system.	AWS provides you with the ability to configure and use the AWS Service offerings in order to maintain appropriate security, protection, and backup of content, which may include the use of encryption technology to protect your content from unauthorized access. You maintain full control and responsibility for establishing and verifying configuration of access to your data and AWS accounts, as well as periodic review of access to data and resources. Using AWS Identity and Access Management (IAM), a web service that allows you to securely control access to AWS resources, you must control who can access and use your data and AWS resources (authentication) and what data and resources they can use and in what ways (authorization). IAM is a feature of all AWS accounts offered at no additional charge. You will be charged only for use of other AWS Services by your users, https://aws.amazon.com/iam/. IAM Best Practices can be found here: http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html.
(h) Use of device (e.g., terminal) checks to determine, as appropriate, the validity of the source of data input or operational instruction.	Not applicable to AWS – this requirement only applies to the customer's system.	You are responsible for establishing and verifying the source of the data input into your system is valid, whether manually, or, for example, by enforcing only certain input devices or sources are utilized.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
(i) Determination that persons who develop, maintain, or use electronic record/electronic signature systems have the education, training, and experience to perform their assigned tasks.	AWS has implemented formal, documented training policies and procedures that address purpose, scope, roles, responsibilities, and management commitment. AWS maintains and provides security awareness training to all information system users on an annual basis. The policy is disseminated through the internal Amazon communication portal to all employees. Relevant SOC2 Common Criteria: CC1.3, CC1.4, CC2.2, CC2.3	You are responsible for ensuring your AWS users—including IT staff, developers, validation specialists, and IT auditors—review the AWS product documentation and complete the product training programs you have determined are appropriate for your personnel. AWS products are extensively documented online, https://aws.amazon.com/documentation/, and a wide range of user training and certification resources are available including introductory labs, videos, self-paced online courses, instructor lead training and AWS Certification https://aws.amazon.com/training/. Adequacy of training programs for your personnel, as well as maintenance of documentation of personnel training and qualifications (such as training record, job description and resumes) are your responsibility.
(j) The establishment of, and adherence to, written policies that hold individuals accountable and responsible for actions initiated under their electronic signatures, in order to deter record and signature falsification.	Not applicable to AWS – this requirement only applies to the customer's system.	Establishment and enforcement of policies to hold personnel accountable and responsible for actions initiated under their electronic signatures is your responsibility, including training and associated documentation.
(k) Use of appropriate controls over systems documentation including:		

21 CFR Subpart	AWS Responsibility	Customer Responsibility
(1) Adequate controls over the distribution of, access to, and use of documentation for system operation and maintenance.	AWS maintains formal, documented policies and procedures that provide guidance for operations and information security within the organization and the supporting AWS environments. Policies are maintained in a centralized location that is only accessible by employees. Security policies are reviewed and approved on an annual basis by Security Leadership, and are assessed by third-party auditors as part of our audits. Refer to SOC2 Common Criteria CC2.2, CC2.3, CC5.3	You are responsible to establish and maintain your own controls over the distribution, access and use of documentation and documentation systems for system operation and maintenance.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
(2) Revision and change control procedures to maintain an audit trail that documents time-sequenced development and modification of systems documentation.	AWS policies and procedures go through processes for approval, version control, and distribution by the appropriate personnel and/or members of management. These documents are reviewed periodically and, when necessary, supporting data is evaluated to ensure the document fulfills its intended use. Revisions are reviewed and approved by the team that owns the document, unless otherwise specified. Invalid or obsolete documents are identified and removed from use. Internal policies are reviewed and approved by AWS leadership at least annually, or following a significant change to the AWS environment. Where applicable, AWS Security leverages the information system framework and policies established and maintained by Amazon Corporate Information Security. AWS service documentation is maintained in a publicly accessible online location so that the most current version is available by default. https://aws.amazon.com/documentation/ Refer to the AWS SOC 2 Report CC2.3, CC3.4, CC6.7, CC8.1	You are responsible for changes to your computerized systems running within your AWS accounts. System components must be authorized, designed, developed, configured, documented, tested, approved, and implemented according to your security and availability commitments and system requirements. Using Service Offerings such as AWS Config, you can manage and record your AWS resource inventory, configuration history, and configuration change notifications to enable security and governance. AWS Config Rules also enables you to create rules that automatically check the configuration of AWS resources recorded by AWS Config, https://aws.amazon.com/documentation/config/ Change records and associated logs within your environment may be retained according to your record retention schedule. You are responsible for storing, managing and tracking electronic documents in your AWS account and as part of your overall quality management system, including maintaining an audit trail that documents time-sequenced development and modification of systems documentation.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
§11.30 Controls for open systems. **Persons who use open systems to create, modify, maintain, or transmit electronic records shall employ procedures and controls designed to ensure the authenticity, integrity, and, as appropriate, the confidentiality of electronic records from the point of their creation to the point of their receipt. Such procedures and controls shall include those identified in §11.10, as appropriate, and additional measures such as document encryption and use of appropriate digital signature standards to ensure, as necessary under the circumstances, record authenticity, integrity, and confidentiality.**	Industry standard controls and procedures are in place to protect and maintain the authenticity, integrity and confidentiality of customer data. Refer to the AWS SOC 2 Report C1.1-C1.2	You are responsible for determining whether your use of AWS Services within your environment meets the definition of an open or closed system and whether these requirements apply. Refer to the responsibilities in §11.10 above for more information for recommended procedures and controls. Additional measures such as document encryption and use of appropriate digital signature standards are your responsibility to maintain data integrity, authenticity and confidentiality.
§11.50 Signature manifestations. **(a) Signed electronic records shall contain information associated with the signing that clearly indicates all of the following:** **(1) The printed name of the signer;** **(2) The date and time when the signature was executed; and** **(3) The meaning (such as review, approval, responsibility, or authorship) associated with the signature.** **(b) The items identified in paragraphs (a)(1), (a)(2), and (a)(3) of this section shall be subject to the same controls as for electronic records and shall be included as part of any human readable form of the electronic record (such as electronic display or printout).**	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications meet the signed electronic records requirements identified.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
§11.70 Signature/ record linking. **Electronic signatures and handwritten signatures executed to electronic records shall be linked to their respective electronic records to ensure that the signatures cannot be excised, copied, or otherwise transferred to falsify an electronic record by ordinary means.**	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the signature/record linking requirements identified, including any required policies and procedures.
Subpart C—Electronic Signatures **§11.100 General requirements.** **(a) Each electronic signature shall be unique to one individual and shall not be reused by, or reassigned to, anyone else.**	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the general electronic signature requirements identified, including any required policies and procedures to enforce electronic signature governance.
(b) Before an organization establishes, assigns, certifies, or otherwise sanctions an individual's electronic signature, or any element of such electronic signature, the organization shall verify the identity of the individual.	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the general electronic signature requirements identified, including any required policies and procedures to verify individual identity prior to use of an electronic signature.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
(c) Persons using electronic signatures shall, prior to or at the time of such use, certify to the agency that the electronic signatures in their system, used on or after August 20, 1997, are intended to be the legally binding equivalent of traditional handwritten signatures. **(1) The certification shall be submitted in paper form and signed with a traditional handwritten signature, to the Office of Regional Operations (HFC-100), 5600 Fishers Lane, Rockville, MD 20857.** **(2) Persons using electronic signatures shall, upon agency request, provide additional certification or testimony that a specific electronic signature is the legally binding equivalent of the signer's handwritten signature.**	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the general electronic signature requirements identified, including determining whether any required notification to the agency is required, and documenting accordingly.
§11.200 Electronic signature components and controls.		
(a) Electronic signatures that are not based upon biometrics shall:	Not applicable to AWS – this requirement only applies to the customer's applications.	

21 CFR Subpart	AWS Responsibility	Customer Responsibility
(1) Employ at least two distinct identification components such as an identification code and password. **(i) When an individual executes a series of signings during a single, continuous period of controlled system access, the first signing shall be executed using all electronic signature components; subsequent signings shall be executed using at least one electronic signature component that is only executable by, and designed to be used only by, the individual.** **(ii) When an individual executes one or more signings not performed during a single, continuous period of controlled system access, each signing shall be executed using all of the electronic signature components.** **(2) Be used only by their genuine owners; and** **(3) Be administered and executed to ensure that attempted use of an individual's electronic signature by anyone other than its genuine owner requires collaboration of two or more individuals.**		You are responsible for establishing and verifying that your applications/systems meet the electronic signature components and controls identified, including establishing the procedures for use of identifying components, and use by genuine owners.
(b) Electronic signatures based upon biometrics shall be designed to ensure that they cannot be used by anyone other than their genuine owners.	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the electronic signature components and controls identified, including establishing the procedures for use by genuine owners.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
§11.300 Controls for identification codes/passwords. **Persons who use electronic signatures based upon use of identification codes in combination with passwords shall employ controls to ensure their security and integrity. Such controls shall include:**		
(a) Maintaining the uniqueness of each combined identification code and password, such that no two individuals have the same combination of identification code and password.	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the electronic signature controls identified, including establishing the procedures and controls for uniqueness of password and ID code combinations.
(b) Ensuring that identification code and password issuances are periodically checked, recalled, or revised (e.g., to cover such events as password aging).	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the electronic signature controls identified, including establishing the procedures and controls for periodic review of password issuance.
(c) Following loss management procedures to electronically deauthorize lost, stolen, missing, or otherwise potentially compromised tokens, cards, and other devices that bear or generate identification code or password information, and to issue temporary or permanent replacements using suitable, rigorous controls.	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the electronic signature controls identified, including establishing the procedures and controls for loss management of compromised devices that generate ID code or passwords.
(d) Use of transaction safeguards to prevent unauthorized use of passwords and/or identification codes, and to detect and report in an immediate and urgent manner any attempts at their unauthorized use to the system security unit, and, as appropriate, to organizational management.	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the electronic signature controls identified, including establishing the procedures and controls to prevent, detect and report unauthorized use of ID codes and/or passwords.

21 CFR Subpart	AWS Responsibility	Customer Responsibility
(e) Initial and periodic testing of devices, such as tokens or cards, that bear or generate identification code or password information to ensure that they function properly and have not been altered in an unauthorized manner.	Not applicable to AWS – this requirement only applies to the customer's applications.	You are responsible for establishing and verifying that your applications/systems meet the electronic signature controls identified, including establishing the procedures and controls to periodically test devices that generate ID codes or passwords for proper functionality.

Notes

[1] Don't Blame Regulators: How Software Excellence Satisfies Compliance | AWS Executive in Residence Blog, https://aws.amazon.com/blogs/enterprise-strategy/stop-blaming-regulations-how-software-excellence-satisfies-compliance/

[2] In computing, JSON (JavaScript Object Notation) is the open-standard syntax used for AWS CloudFormation templates, https://aws.amazon.com/documentation/cloudformation/.

[3] https://www.continuousvalidation.com/what-is-continuous-validation/