

# Navigating HCLS Regulatory and Compliance Requirements on AWS

*April 13, 2022*



## Notices

Customers are responsible for making their 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.

© 2022 Amazon Web Services, Inc. or its affiliates. All rights reserved.

# Contents

- Introduction ..... 1
  - About AWS ..... 2
- Fundamentals ..... 3
  - Shared responsibility ..... 4
  - Confidentiality, integrity, and availability ..... 7
- HCLS scenarios ..... 9
  - Medical devices (MedTech) ..... 10
  - Genomics ..... 15
  - Biopharma ..... 21
  - Providers ..... 24
  - Payors ..... 27
- How can AWS support you? ..... 34
  - DIY route ..... 34
  - Partner route ..... 35
  - Professional service route ..... 35
  - Hybrid approaches ..... 36
  - AWS Professional Services – regulatory and compliance practice ..... 36
- Conclusion ..... 37
- Contributors ..... 38
- Further reading ..... 38
- Appendix A: HCLS regulatory and compliance principle concept reference ..... 40
  - Data provenance ..... 41
  - Data residency, transfers, and adequacy ..... 42
  - Data at rest and in motion ..... 45
  - Data privacy ..... 47
  - Change management ..... 49

Logging and auditing.....	50
Backup, recovery, and disaster recovery (DR).....	51
Consent management.....	54
Continuous compliance.....	57
Compliance at scale.....	59
Appendix B: Compliance, certification, and regulatory alignment.....	63
Document revisions.....	69

## Abstract

Amazon Web Services (AWS) customers and partners need information to understand how AWS supports worldwide healthcare and life sciences (HCLS) regulatory and compliance programs. There are common concepts that span all regulatory and compliance programs, such as shared responsibility, data at rest, data in motion, encryption, data residency, and change management. This whitepaper can help you understand common questions relative to AWS, and regulatory and compliance topics.

## Introduction

Amazon Web Services (AWS) provides cloud services designed to help customers run and manage their most sensitive healthcare and life sciences (HCLS) workloads in the cloud. The AWS Cloud enables integrations between electronic health record (EHR) systems, and provides tools that help customers meet and enforce regulatory requirements. AWS also offers industry-leading services and features that allow customers to control where their data is stored, who can access it, and help protect data, accounts, and workloads from unauthorized access.

This paper provides an overview of general regulatory and compliance concepts that cross all regulations that can be found in HCLS regardless of country (see [Appendix A: HCLS regulatory and compliance principle concept reference](#)). These concepts provide valuable insight into how data may be secured, transmitted, and stored, and how it can be used and disposed of, providing a prescribed set of common controls dictating the entire data lifecycle.

This paper then outlines how AWS shares the responsibility with its customers by protecting global infrastructure, building and offering services that have strong security features, and enabling HCLS customers to protect their own data. This paper examines the fundamental principles of the [AWS shared responsibility model](#) and how it applies to HCLS, and defines key concepts such as confidentiality, availability, and integrity.

Many customers already use industry guidance to influence their interpretation of these regulatory concepts. Therefore, the focus of this whitepaper is to help you understand common applications of regulatory and compliance concepts and how they align with key HCLS industry subsegments such as [Medical devices \(MedTech\)](#); [Genomics](#); [Biopharma](#); [Providers](#) and [Payors](#). This paper provides resources for a deeper understanding of each concept focus area outlined in Appendix A: HCLS regulatory and compliance principle concept reference, depending on the global regulatory framework of interest. This paper gives you an understanding of these fundamental concepts that are common across all frameworks and are critical to how AWS and the cloud, in general, operate. Next this paper provides reference information describing and helping you to understand the aforementioned regulatory and compliance concepts common to all regulatory and compliance frameworks.

Finally, this paper describes several of the key global frameworks, critical laws and regulations, and certifications that are HCLS industry-recognized as important to how protected data is handled (refer to [Appendix B: Compliance, certification, and regulatory alignment](#)).

## About AWS

In 2006, AWS began offering on-demand information technology (IT) infrastructure services to businesses in the form of web services with pay-as-you-go pricing. Today, AWS provides a highly reliable, scalable, low-cost infrastructure platform in the cloud that powers hundreds of thousands of businesses in countries worldwide. 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 provides customers the ability to take advantage of a broad set of global cloud-based products, including compute, storage, databases, networking, security, analytics, mobile, developer tools, management tools, Internet of Things (IoT), and enterprise applications.

### AWS healthcare and life sciences

AWS started its dedicated Genomics and Life Sciences Practice in 2014 in response to the growing demand for a cloud services provider with experience and reliability in the healthcare and life sciences (HCLS) industry. Today, the AWS Life Sciences Practice team consists of members who have performed industry health and life sciences roles including: Chief Medical Officer, Chief Digital Officer, physician, radiologist, and researcher, among others. The AWS Genomics, Healthcare, and Life Sciences industry vertical serve many customers, including pharmaceutical, biotechnology, medical device, genomics, start-ups, universities, government institutions, and healthcare. A complete list of AWS customer case studies can be found at [Healthcare & Life Sciences Case Studies](#).

In addition to the resources available within the HCLS Practices at AWS, customers can work with AWS Life Sciences Competency Partners to drive innovation and improve efficiency across the HCLS 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 sciences solutions on AWS. A full list of Competency Partners can be found at [AWS Life Sciences Competency Partners](#) and [AWS Healthcare Competency Partners](#).

## Fundamentals

The landscape of IT concerning HCLS has consistently focused on the security of the network over which data flows, the security at the location where the data is stored, and the ability to monitor who and what has access and is accessing the data in near real-time with visibility across all assets. Regulations and certifications exist to verify that the systems transmitting and storing this information remain secure, both from an infrastructure and a software perspective. Across these regulations, a common set of control categories have emerged that have been used colloquially to group and categorize the functions of various regulations and frameworks with respect to their impact on private health data and information systems. These common conceptual groupings include:

- *Data at rest and in motion*
- *Data privacy*
- *Data provenance*
- *Data residency, transfers, and adequacy*
- *Consent management*
- *Change management*
- *Compliance at scale*
- *Backup, recovery, and disaster recovery (DR)*
- *Logging and auditing*
- *Continuous compliance*

When you examine how privacy and security are implemented in the cloud, specifically the AWS Cloud, they can be described using the [AWS shared responsibility model](#). The AWS shared responsibility model explains the AWS set of responsibilities *of* the cloud, and the customer responsibilities *in* the cloud. This is expanded upon through an HCLS lens later in this section.

An overview of key concepts such as confidentiality, availability, and integrity wrap up this section. These concepts are critical because they guide policies for information security, defining the rules that limit access to information, provide assurance that information is trustworthy and accurate, and that the information can only be accessed

by users who are authorized to do so. For more information, see [Best Practices for Security, Identity, & Compliance](#) in the AWS Architecture Center.

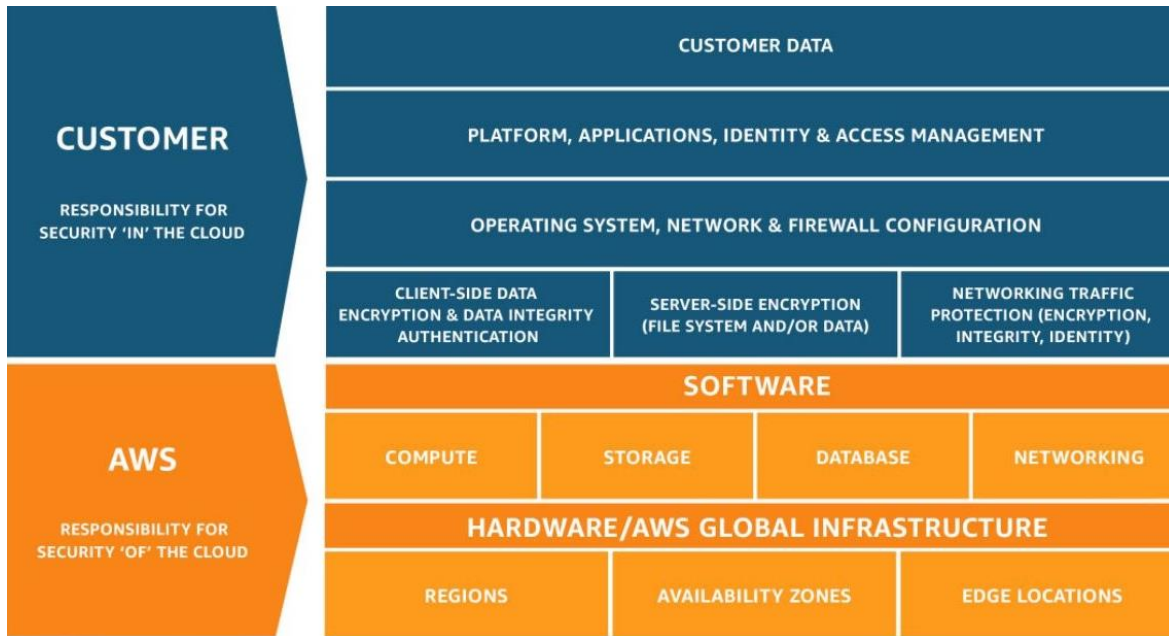
## Shared responsibility

At a high level, AWS is responsible for the security *of* the cloud itself, and the customer is responsible for their security *in* the cloud—the elements the customer has control over. For example, if the customer is using an [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) instance, the customer is responsible for the following aspects:

- The operating system, vulnerabilities, and implementing security patches.
- The applications, vulnerability, coding practices (such as not storing passwords or access keys within the code), and encryption of data in transit.
- The persistent layer— whether it consists of a database, or a file system, as well as the choice to encrypt the data at rest.
- Who and what can access the instance itself, for instance or application management.


AWS provides certifications, attestations, and regulatory control alignments with worldwide frameworks and laws as evidence that AWS appropriately manages and protects the infrastructure *of* the cloud. AWS is responsible for protecting the infrastructure that runs all of the services offered in the AWS Cloud. This infrastructure is composed of the hardware, software, networking, and facilities that run AWS services. AWS makes third-party attestation reports available through the [AWS Artifact](#) service within the [AWS Management Console](#). Builders of applications on AWS can use these third-party reports as evidence in their own compliance documentation.

The following figure is an excellent visual of that separation for AWS services.



*AWS shared responsibility model*

It is important to understand that AWS offers a multitude of services that can be categorized as *infrastructure* services, *container* services, or *managed* services. Each category of service has different aspects that are under the customer's control, from a configuration and management perspective. The following diagram from [Applying the AWS Shared Responsibility Model to your GxP Solution](#) depicts the three categories, and which aspects fall under the AWS responsibility of the cloud and the customer responsibility *in* the cloud under the AWS shared responsibility model.



*Spectrum of shared responsibility of different service categories*

For example, with the infrastructure service Amazon EC2, AWS is responsible for the compute, storage, and network infrastructure. Simultaneously, the customer is responsible for everything that goes into an Amazon EC2 instance, such as the operating system, networking configuration, access configuration, and platform application management.

When using the container service [Amazon Relational Database Service \(Amazon RDS\)](#), there is a separation of concerns between what database administrators need to understand about infrastructure and lower-level application configuration. AWS is responsible for infrastructure services and tasks, such as the database installation, management, and database-specific networking aspects associated with replication and redundancy. The customer is responsible for the database’s configuration, client-side encryption, network traffic protection, and importing their data into the database.

For managed services, such as [Amazon Simple Storage Service \(Amazon S3\)](#), AWS is responsible for security and compliance aspects of the cloud. However, the customer is responsible for tasks such as those associated with controls for the data they bring, the encryption they choose to use, and who and what services have access to the data.

In this manner, AWS provides a continuum across the responsibility spectrum between *of* the cloud and *in* the cloud scenarios, allowing solution developers to focus on their innovation while AWS performs the undifferentiated lifting relative to infrastructure and service management. This spectrum provides customers with the flexibility to choose the most appropriate services for both their solution and regulatory requirements.

A managed database like Amazon RDS may not always be the best fit. AWS has services to match the customer's business needs, and although a managed database shifts several responsibilities to AWS, it is not quite as flexible as to when the database runs in an Amazon EC2 instance. The customer has complete control to adjust and tune the database to their needs—by offering all three of these models, AWS offers the most flexible platform for each specific scenario.

By using good cloud design practices as described in the [AWS Well-Architected Framework](#), and understanding how the AWS shared responsibility model may support your regulatory certification requirements, you can allow AWS to perform more of the infrastructure management *of* the cloud while maintaining a strong compliance posture. Additionally, you can make use of various [AWS reference architectures](#) and [AWS Quick Starts](#) that exist for building compliant solutions, such as [GxP Compliance Automation](#) or [HIPAA Reference Architecture on the AWS Cloud QuickStart](#). As a result, you can focus on your solutions' added value instead of infrastructure management.

## Confidentiality, integrity, and availability

*Confidentiality, integrity, and availability*, also known as the [CIA triad](#), is a model designed to guide policies for information security within an organization. The model is also sometimes referred to as the *AIC triad* (availability, integrity, and confidentiality) to avoid confusion with the Central Intelligence Agency (U.S. government agency).

In this context, *confidentiality* is a set of rules that limit access to information, *integrity* is the assurance that the information is trustworthy and accurate, and *availability* grants reliable access to the information by authorized people.

## Confidentiality

Confidentiality measures are designed to prevent sensitive information from unauthorized access attempts. It is common for data to be categorized according to the amount and type of damage that can be done if it falls into the wrong hands. More or fewer stringent measures can then be implemented according to these categories.

Sometimes, safeguarding data confidentiality involves special training for users with access to sensitive documents. Training can help familiarize authorized people with risk factors and how to guard against them. Further aspects of training may include strong passwords and password-related best practices, and information about social engineering methods to prevent users from bending data-handling rules with good intentions, yet having potentially disastrous results.

Besides verifying that sensitive data is encrypted in transit and at rest, limiting who has access to sensitive information should also use two-factor authentication, and applicable technologies such as biometric verification, security tokens, key fobs, or soft tokens. In addition, customers can take precautions to minimize the number of places where information appears and the number of times it is transmitted to complete a required transaction. Extra measures might be taken in the case of extremely sensitive documents, such as storing them only on air-gapped computers or disconnected storage devices.

## Integrity

Integrity involves maintaining the consistency, accuracy, and trustworthiness of data over its entire lifecycle. Data must not be changed in transit, and steps must be taken to verify that unauthorized people cannot alter data (for example, if there is unintended broad access or disclosure).

These measures include file permissions and user access controls. Version control may be used to prevent erroneous changes or accidental deletion by authorized users from becoming a problem.

Data might include checksums, even cryptographic checksums, for verification of integrity. Backups or redundancies must be available to restore the affected data to its correct state. Furthermore, digital signatures can be used to provide effective nonrepudiation measures, meaning to capture evidence of logins, messages sent, electronic document viewing, and sending so that a chain of evidence exists and cannot be denied.

## Availability

Availability means information should be consistently and readily accessible for authorized parties and involves properly maintaining hardware and technical infrastructure and systems that hold and display the information. To enable availability, rigorously maintain all hardware, perform hardware repairs immediately when needed, and maintain a properly functioning operating system environment, to validate that it is free of software conflicts. It's also important to keep current with all necessary system upgrades. Providing adequate communication bandwidth and preventing the occurrence of bottlenecks are equally important tactics. Other ways to mitigate for unforeseen events is for the architecture to be highly resilient, to include redundancy and failover strategies, and to consider leveraging multiple geographical regions where appropriate.

Fast and adaptive disaster recovery (DR) is essential for worst-case scenarios; that capacity relies on a comprehensive DR plan. Safeguards against data loss or interruptions in connections must include unpredictable events such as natural disasters and fire. For example, a cloud-enabled DR strategy can include active-active multi-Region deployment for mission critical workloads. Layered security capabilities can provide defense in depth such as firewalls and proxy servers, which can guard against downtime and unreachable data blocked by malicious denial-of-service attacks and network intrusions.

For more information about which AWS services can be used, see the [AWS Security Reference Architecture](#).

## HCLS scenarios

To better describe how the categories covered in [Appendix A: HCLS regulatory and compliance principle concept reference](#) are applied to healthcare and life sciences in the real world, the following sections describe real-world scenarios across multiple sub-segments from genomics, biopharma, and medical devices to those sub-segments that specifically focus on providers and payors. The scenarios described will help you to understand how the categories described in the principal concept reference section apply to that sub-segment. Additionally, where applicable, guidance on how AWS can help reduce the regulatory burden or references to further reading will be given.

These scenarios should be used as guideposts on your journey to understanding the applicability of these concepts to your scenarios and illustrate how AWS certifications and shared responsibility may apply to you.

## Medical devices (MedTech)

The United States Food and Drug Administration (FDA) [defines a medical device](#), per Section 201(h) of the Food, Drug, and Cosmetic Act, as:

An instrument, apparatus, implement, machine, contrivance, implant, in vitro reagent, or other similar or related article, including a component part or accessory which is:

1. Recognized in the official National Formulary, or the United States Pharmacopoeia, or any supplement to them,
2. Intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease, in man or other animals, or
3. Intended to affect the structure or any function of the body of man or other animals, and which does not achieve its primary intended purposes through chemical action within or on the body of man or other animals and which does not achieve its primary intended purposes through chemical action within or on the body of man or other animals and which is not dependent upon being metabolized for the achievement of its primary intended purposes. The term “device” does not include software functions excluded pursuant to section 520(o).

From this definition, a medical device can be many things. For example, it can be a skillfully created solution to diagnose, cure, treat, or prevent a disease. Medical devices can be hardware only (such as joint replacements), Software as a Medical Device (SaMD), or a combination of both. They can be used for therapeutic purposes (such as smart injectors) or diagnostic purposes (such as blood pressure or glucose readings); heavily regulated or lightly regulated; controlled by the patients directly, or have the adjustable settings controlled by a remote care team. The [universal scope of medical devices](#) is vast and can be part of all aspects of healthcare.

The term SaMD is [defined by the International Medical Device Regulators Forum \(IMDRF\)](#) as:

- Software intended to be used for one or more medical purposes that perform these purposes without being part of a hardware medical device.

SaMD is software with a medical purpose, whether it's included or not with a physical device or other SaMD solution. It runs on a general-purpose computing device (the hardware it's running on does not have a specific medical purpose), even if this general computing device is part of the medical device itself (for example, a general-purpose personal computer included with, and attached to, a medical image device that runs the user interface for the technician).

Embedded software that is required by a medical device to perform its function is not considered SaMD. Also not considered as part of SaMD, is software that provides input to a SaMD solution or processes the output of a SaMD solution where such software does not have a medical purpose. Examples of this include database queries, search results, and data encryption software that stores the output of an SaMD solution.

There are several regulatory programs and laws that impact medical devices across the globe. It is up to the users of AWS, based on the [AWS shared responsibility model](#), to understand which regulatory programs are relevant for their product and delivery areas. Examples of these programs may include FDA Guidelines, U.S. 510(k), EU Medical Devices Directive (93/42/EEC), Medical Device Regulation (MDR) 2017/745, HIPAA, and GDPR.

The intersection of the medical device universe and AWS exists relative to the software associated with a hardware device, SaMD, and in all cases, the collection, storage, presentation, and analysis of data associated with all types of devices. Some illustrative examples include:

- A database allowing the lookup of an implant's serial numbers, which are associated with patient's information, and when they were used. Although the device has no software, the support software would be subject to regulation by privacy laws and related authorities.
- A smart inhaler that collects information about each time it's used (when, where, and how much medicine was dispensed) and the device then connects to the cloud to store, process, and present the information about the device usage to the user and the care team.
- A blood glucose monitor connected to the cloud to store, process, and present the information back to the user and their care team. In addition, the software will analyze past readings and make an insulin unit recommendation to allow the patient and care team to keep the patients glucose reading within a target range.
- An application running on a smartphone or other device for viewing magnetic resonance imaging (MRI) for diagnostic purposes.

The following section discusses the principal categories from this paper and how they relate to medical devices relative to AWS.

## Applicability to principle concepts

The following sections provide examples of how each of the principle concepts outlined in the [Appendix A: HCLS regulatory and compliance principle concept reference](#) section apply to medical devices.

### Data provenance

Regulations regarding electronic records, such as [FDA Title 21 CFR Part 11](#), focus heavily on data integrity, making it a priority for medical device companies. Regulations often stipulate that the original record, in this case the data elements generated or originating on or from a medical device, should always be available and therefore, any changes should not obscure previously recorded information. In addition, it is common for the data generated from a medical device to go through conversions or transformation as the data moves through a system, and HCLS regulations often require the preservation of such changes for audit purposes. For example, if an update is made to a device that converts units to a different format, the original values may need to be preserved, perhaps recorded in an audit trail. In this way, every operation performed on the data would be tracked to enable a regulatory inspector to trace back through the data lineage to the original values.

As part of this data integrity and provenance, ensuring that the data is not manipulated in-flight is a consideration for manufactures and developers. In addition, reliable and accurate tracking of the source of the data as part of its providence is another consideration.

### Data residency, transfers, and adequacy considerations

Citizenship and the physical geographic location of a person who uses a device must be taken into consideration, because this may impact which AWS Region is used to store the device data. Suppose a country requires all clinical data for a citizen to remain in its geographic boundaries. In this case, the design of the medical device needs to verify that it stores data within that boundary, even if the person using the device has traveled outside its borders, in some cases.

AWS provides software architects and device designers the ability to fully control where data is stored to meet various regulatory requirements. Based on the [AWS shared responsibility model](#), it's the device designer's responsibility to use these mechanisms to verify that their user's data is stored appropriately to meet regulatory data residency

needs, and it's the responsibility of AWS to certify that these mechanisms reliably function as designed.

### **Data at rest and in motion**

As with many of the other scenarios discussed, various regulations address the privacy and integrity of protected health information (PHI) in solutions and define the requirements for data associated with medical devices. For information about methods to manage these requirements in AWS, see the [Data at rest and in motion](#) section in *Appendix A*.

### **Data privacy**

Data privacy is a central tenet of many regulatory programs. These programs dictate that patients control their personal health records and require healthcare entities not to share protected health information (PHI) and personally identifiable information (PII) without proper consent procedures. For more information, see the [Data privacy](#) section in *Appendix A*.

### **Change management**

The international standard IEC 62304 defines the lifecycle process requirements for medical devices. IEC 62304 Part 8 addresses the software configuration management process, which includes change management. For information about methods to manage these requirements in AWS, see the [Change management](#) section in *Appendix A*.

### **Logging and auditing**

As with other scenarios where GxP regulations have requirements, the need to store and maintain immutable audit logs for creating an audit trail of any actions which might impact data integrity is of primary importance. It is critical to maintain logs for different types of activities, such as system events, data access, and user activity. These activities include, but are not limited to, tracking all operations on user accounts and data; logging relevant metadata such as data transformations or when applicable, device settings changes relevant to the data processing; and monitoring and logging access and attempted access to data, and applications. In addition, as discussed in the [Data provenance](#) section, the logging and ability to follow data back through the system to its original source is crucial for audit support.

## Backup, recovery, and disaster recovery (DR)

Regulations for medical devices focus on patient safety, which is driven by many factors, including product quality and data integrity. As a result, a validated backup and recovery procedure may be needed to meet the regulatory requirements to enable data integrity and quality, allowing the solution to meet its intended use during a failure or recovery scenario as described in the *Appendix A* section on [Backup, recovery, and disaster recovery \(DR\)](#), subsection [GxP considerations](#).

## Consent management

In several areas of HCLS, such as clinical trials or patient support programs, it is essential to have appropriate management of participants and their data. Current regulations mandate that collecting and using personal data from participants must comply with rigorous standards. To help ensure privacy, the ability for an individual to grant permission for other groups or individuals to access participant data may be regulated by HIPAA, GDPR, or other localized HCLS legal authorities. Refer to the discussion in the *Appendix A: HCLS regulatory and compliance* principle concept reference section on [Consent management](#) for information on methods to manage these requirements in AWS.

As in the *Biopharma* scenario, with clinical trials or patient support programs, it is essential to have appropriate management of participants and their personal data. Refer to the *Biopharma* section on [Consent management](#) for a discussion on clinical trials and consent management.

## Continuous compliance

As described in [Compliance at scale](#), historically, point-in-time reviews are used to verify that the configuration of a medical device solution meets its compliance standards per the requirements outlined in an organization's quality management system. In addition, AWS provides a set of services discussed further in the *Appendix A: HCLS Regulatory and compliance principle concept reference* section on [Continuous compliance](#) to implement a model of continuous compliance in the cloud.

## Compliance at scale

Being able to demonstrate that your environment is in a state of control is at the core of many regulations that apply to the medical device industry. Because your environment grows, this becomes a more significant challenge. Refer to the *Appendix A: HCLS regulatory and compliance* principle concept reference section on [Compliance at scale](#) for considerations relative to both enterprise-wide and dynamic resource scaling relative to GxP and other regulatory requirements.

## Genomics

The reduced cost of DNA sequencing technology has caused a massive splurge in next-generation sequencing (NGS) data. Several initiatives, including large-scale population genomics projects, are collecting, storing, sharing, and analyzing NGS data to predict risk for different diseases, determine personalized treatment regimes, and improve patient outcomes and care. Because these initiatives pave the way for advanced scientific discovery, enforcing and maintaining the privacy and security of genomic data is of utmost importance. In particular, genomic data poses specific challenges that must be considered to comply with regulatory guidelines.

Even after de-identification and anonymization of genomic data collected from human subjects, certain DNA sequences can uniquely identify an individual. This increases the risk of potentially linking and correlating genomic data with other “-omic” (such as proteomic, metabolomic, and transcriptomic) datasets or data modalities (such as medical insurance, electronic health records, and medical imaging), which can inadvertently expose patient-level details.

Exposing genomic data not only reveals sensitive information about a person but also kinship. Preserving the privacy of genomic data collected from identifiable populations, such as certain ethnic groups or patients with a rare disease, can lead to discrimination or stigmatization and is even more challenging.

### Applicability to principle concepts

The following sections provide examples of how each of the principles outlined in the [Appendix A: HCLS regulatory and compliance principle concept reference](#) section apply to genomics.

#### Data provenance

Advancements in sequencing technology and algorithms have led to a significant increase in the number and scale of genomics projects. To harness the potential of genomic data, it is crucial to maintain its provenance and integrity. Controlling the quality of data throughout its lifecycle further enhances the reproducibility of genomic data analysis, which, in turn, can save the underlying cost and time incurred by large-scale projects.

Because researchers often struggle to trace the data associated with published research, setting up data repositories that are able to meet regulatory guidelines can help share data and results more widely. Genomics medicine research relies on

associating genotypic and phenotypic data to determine disease susceptibility and treatment regime. To facilitate this, interoperability and integration of genomic data with other data formats are essential. This further indicates the need for maintaining the provenance and quality of genomic data.

The inability to maintain proper documentation and metadata can lead to inconsistencies in data analysis and interpretation. In this context, Amazon S3 supports version control on a bucket, which offers a level of traceability at the object storage level, however this would only be one link in the data provenance chain. Amazon S3 and [AWS CloudTrail](#) can be used to maintain the provenance of data and the actions performed during its lifecycle.

### **Data residency, transfers, and adequacy considerations**

The emerging field of population genomics has been attributed to large-scale national genomic initiatives. Many nations have invested in launching population-scale genomic initiatives focusing on rare diseases and cancer in the past decade. One of the preliminary challenges nations often face is developing infrastructure, particularly common standards and policies for data sharing. Scientific research has proven the value of cross-border data sharing. This includes international consortia and multinational research projects like the International Cancer Genome Consortium (ICGC), which involves cross-country sharing of genomic, epigenomic, and transcriptomic data among the world's leading genomic and cancer researchers. Similarly, the Human Cell Atlas was launched to create a reference map of all human cells.

Although a coordinated effort can help disseminate resources and knowledge, it must be noted that different countries enforce different regulatory guidelines regarding data residency and sharing. For example, GDPR aims at securing a high level of protection of personal data in all EU member states. When sharing data with another country, they also have to follow the regulations imposed by the respective country. According to the AWS shared responsibility model, customers own their data and choose the AWS Region for where it will be stored. The data does not move unless the customer decides to move it. They can further decide if and how to encrypt the data at rest or in transit.

As noted in the whitepaper [Data Residency: AWS Policy Perspectives](#), abiding by the data residency requirements does not necessarily prevent unauthorized data access. To help protect against such access, customers can opt for encryption, tokenization, data decomposition, and cyber deception, which are designed to render content unintelligible to parties seeking data access.

## Data at rest and in motion

Data owners must strive to protect sensitive genomic data, including PII, to comply with the requirements of data protection legislation and standards and preserve their competitive advantage. Furthermore, the [National Institutes of Health \(NIH\) Genomic Data Sharing Policy](#) and [Security Best Practices for Controlled-Access Data Subject to the NIH Genomic Data Sharing Policy](#) specify instructions to protect genomic data being shared with the NIH. Such data must be protected against inadvertent data exposure. Encryption, tokenization, and access control are the most common approaches to protect genomic data at rest.

Encryption involves transforming data into another form using a secret key, whereas tokenization creates a token to represent an otherwise sensitive piece of information. [AWS Key Management Service \(AWS KMS\)](#) integrates with most services to let customers manage the encryption keys, including both server-side and client-side encryptions. Customers can control when genomics data and metadata are encrypted, by whom, and under which conditions.

To protect data in motion, using a multi-level approach is advisable. At the physical layer, all traffic between AWS data centers is encrypted. At the network layer, traffic inside [Amazon VPC](#) and across [peered VPCs](#) can also be encrypted. At the application layer, customers can choose to implement encryption by using TLS protocol and create a secure HTTPS connection. In addition to AWS KMS, AWS CloudHSM (cloud-based hardware security module) and ACM can be used to encrypt data at rest and in motion.

## Data privacy

Increased availability of patient-level genomic data poses a significant concern for personal privacy. Even after de-identification and anonymization of genomic data, certain genomic sequences can identify an individual. This increases the risk of re-identifying patients, associating them with phenotypic data, and revealing kinship. The gaining popularity of direct-to-consumer DNA testing can also lead to sharing genomic data in a less regulated environment.<sup>1</sup>

The level of data privacy to be enforced often depends on the intent of using genomic data, for instance, healthcare, research, and direct-to-consumer services. The U.S. Genetic Information Nondiscrimination Act (GINA) in 2007, policies enacted by the Council of Europe in 2008, and the U.S. Presidential Report on genome privacy, 2012, among other authorities, govern protection of genomic data privacy. However, their requirements are often difficult to enforce as the use of genomic data may not be detected *a priori*. In addition to legal and policy mechanisms, implementing access

control, anonymization, and cryptographic methods can lead to a higher level of data privacy with a potential trade-off of reducing the utility of the data. Customers can use encryption, access, and logging features offered by services such as [AWS Identity and Access Management \(IAM\)](#), AWS KMS, [AWS Organizations](#), and AWS CloudTrail.

### **Change management**

Manual configuration and deployment of infrastructure for genomic data analysis are often prone to errors and misconfigurations. According to industry best practices, infrastructure as code (IaC) can mitigate this challenge by automating infrastructure deployment and reducing the associated risks. It also reduces the risk of knowledge gaps caused by hierarchical changes or employee turnover in an organization.

IaC facilitates deploying servers in a reproducible manner, reducing time and cost incurred by manual intervention, providing traceability of changes and configurations, and allowing for the codification of security standards that then can be applied in every deployment. The CI/CD aspect of it allows making changes in a nearly continuous flow rather than in large batches. Because in many cases frequently reviewing and testing code are required before deploying to a production environment, IaC provides efficient and safer change management across various infrastructures. AWS offers [AWS CloudFormation](#) as a built-in choice for IaC. It allows provisioning AWS and third-party resources, and managing them throughout their lifecycle.

### **Logging and auditing**

Regulatory guidelines defined by HIPAA and GDPR require storing and maintaining immutable audit logs to create an audit trail. In case of a genomic data exposure, a detailed log of data and activities would likely need to be analyzed to understand the root cause of the event. Genomics projects involve multiple parties and stakeholders, such as researchers, hospitals, insurance companies, pharmaceutical companies, and CROs. Therefore, regular audit logging, analysis, and system-level inspection are necessary to monitor and restrict suspicious behavior, and security events.

It is vital to maintain logs for different types of activities, such as genomic data and metadata access, system-level event, and user activity. These include, but are not limited to, tracking all operations on user account and data, logging relevant metadata, storing informed consent, and monitoring and logging access, and attempted access to restricted genomic data and applications. AWS CloudTrail can be used to log, monitor, and retain account activity across AWS infrastructure.

CloudWatch can be used to group audit logs from different systems, applications, and AWS services. CloudWatch also offers [CloudWatch Logs Insights](#) to search and analyze log data interactively. In addition, customers can use Audit Manager to continually audit their AWS usage and assess risk and compliance with regulations and industry standards.

### **Backup, recovery, and disaster recovery (DR)**

The advancement of biomedical research and large-scale collaborations have led to the significant growth of genomic data. [It is estimated that genomic research will generate between two and 40 exabytes of data within the next decade.](#) The cost of storing raw data will soon surpass the cost of sequencing and analyzing. Backup and recovery of such large-scale genomic data are critical for organizations to combat malicious attacks or natural disasters and meet the RTO and RPO for their workload.

Traditional backup methods, such as tape libraries and secondary standby sites, are often inadequate and cost-prohibitive. AWS offers scalable and advanced backup and recovery infrastructure for different data types, including block storage with [Amazon Elastic Block Store \(Amazon EBS\)](#), object storage with Amazon S3 and Amazon S3 Glacier, and file storage with [Amazon Elastic File System \(Amazon EFS\)](#). Common file formats, such as Binary Base Call (BCL) and Binary Alignment Map (BAM), which are accessed infrequently, can be stored in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) to reduce storage cost.

For long-term storage and backup of genomic data, [Amazon S3 Glacier](#) and Amazon S3 Glacier Deep Archive are suitable options. Copies of all data uploaded to Amazon S3 and Amazon S3 Glacier are stored across multiple [Availability Zones within an AWS Region](#) and can also be replicated across multiple [AWS Regions](#). AWS managed database services, such as [Amazon Relational Database Service \(Amazon RDS\)](#), [Amazon DynamoDB](#), [Amazon Aurora](#), and [Amazon Neptune](#), also have a built-in backup feature. AWS offers four approaches for DR, ranging from low cost and low complexity to complex strategies using multiple active Regions. These include:

- [backup and restore](#)
- [pilot light](#)
- [warm standby](#)
- multi-site [active-active](#)

## Consent management

Studies that involve large-scale genomic data collection, such as whole-genome sequencing (WGS), whole-exome sequencing (WES), and epigenetic and microbiotic profiling, generate massive amount of personal information about participants.

Genomics data is considered one of the most sensitive personal data so consent is required, and special care should be taken to protect it, as specified in regulatory guidelines such as HIPAA and GDPR.

It is important to protect the rights and interests of participants or human subjects who contribute genomic samples and related information to research projects. The National Institute of Health Genome Data Sharing (NIH GDS) Policy requires informed consent documents for prospective data collection to state the data type (such as genomic, phenotype, or health information) to be shared, the purpose of sharing (such as disease-specific research use or general research use), and mode of sharing (such as unrestricted or controlled-access databases).

As of January 2021, the National Human Genome Research Institute (NHGRI) requires all human data generated by NHGRI-supported research to document explicit consent for future research use and broad data sharing. For customized use cases, managed relational database services such as Amazon RDS and Amazon Aurora, or NoSQL database services such as DynamoDB, can be used to store and manage informed consent. However, [AWS Marketplace](#) can also provide solution offerings for more common use cases.

## Continuous compliance

Genomics organizations must adhere to industry regulations and compliance requirements, such as those set forth by HIPAA, GDPR, Global Alliance for Genomics and Health (GA4GH), and Clinical Laboratory Improvement Amendments (CLIA). This requires defining and implementing controls to maintain continuous compliance for organizations. [HIPAA requires](#) organizations to conduct an initial risk assessment and maintain a nearly continuous risk analysis mechanism.

To maintain continuous compliance, an organization must regularly review its records, set up timely or triggered scans, look for security events, and enable real-time monitoring to address new risks as soon as possible. A [continuous compliance workflow based on CI/CD and IaC](#) can be beneficial. [AWS Config](#) can be used to continuously record, monitor, compare, and react to changes in an environment. Integrating custom AWS Config rules with [AWS Lambda](#) functions can enforce compliance.

## Biopharma

Biopharma companies operate within a highly regulated industry and are subject to regulatory compliance from local and federal agencies to verify patient safety and product quality. In addition, biopharma companies use the cloud to innovate across their entire value chain to bring drugs to market faster and cheaper to help patients.

Regulations require a company to maintain certain records and submit specific information to the regulatory agency as part of compliance. Although these rules originally applied to paper records with handwritten signatures, [FDA Title 21 CFR Part 11](#) and similar regulations are also applicable to electronic records and signatures used for compliance purposes.

The cloud provides excellent efficiency and benefits, but you also have to verify that the workloads and data in the cloud continue to meet all life sciences regulatory requirements.

### Applicability to principle concepts

The following sections provide examples of how each of the principle concepts outlined in the [Appendix A: HCLS regulatory and compliance principle concept reference](#) section apply to biopharma.

#### Data provenance

Regulations surrounding electronic records, such as FDA Title 21 CFR Part 11, focus heavily on data integrity, making it a priority for biopharma companies. Regulations tend to stipulate that the original record should always be available, and any changes should not obscure previously recorded information. For example, if an update is made to some data, the original values must still be available, perhaps recorded in an audit trail. In addition, every operation on the original data should be tracked to enable a regulatory inspector to trace back through the data lineage to the original values.

#### Data residency, transfers, and adequacy considerations

Biopharma companies are often international in nature, be it collaborative research, clinical trials conducted on a global level to collect enough data to demonstrate safety and efficacy, or global distribution of those drugs. The data captured from such operations linked to people is subject to personal data protection laws such as the GDPR. Such data is generally considered sensitive personal data under United States data protection laws, or a special category of personal data within the EU GDPR.

Particular attention should be taken to determine if the laws in all regions of operation include data residency requirements.

### **Data at rest and in motion**

The [FDA Title 21 CFR Part 11.10](#) states that systems used to create, modify, maintain, or transmit electronic records shall employ procedures and controls designed to validate the authenticity, integrity, and confidentiality of electronic records at all times. Encryption of the data at rest and in motion is a critical part of satisfying those requirements.

### **Data privacy**

Personal data may be used at various stages in the biopharma value chain, from research to patient support. Under many regulations, the processing of data purely related to research activities must be distinguished from processing related to protecting health. Moreover, data may be subject to several regulations at the same time, such as Clinical Trials Regulation (EU) 536/20141 and the GDPR (EU) 2016/679, which share some common goals with respect to the protection of PHI.

### **Change management**

After a computer system moves into production, operational change management should continue until the system is retired. If the data needs to be retained after system retirement, that data will continue to be subject to change control.

All changes should be reviewed, impact and risk assessed, authorized, documented, tested, and approved before implementation. The amount of testing performed should be commensurate to the risks to patient safety, product quality, and data integrity introduced by the change. Testing should not only prove that the new or changed system functions as specified but also that the change has not introduced any defects.

If the change requires it, user training should also be updated and delivered along with any changes to standard operating procedures.

Configuration management is closely associated with change management. The configuration of all hardware and software assets should be managed throughout the life of the system. The configuration information should be sufficient to allow the system to be effectively and efficiently rebuilt in the event of complete system loss.

### **Logging and auditing**

GxP regulations often require storing and maintaining immutable audit logs to create an audit trail of any actions which may impact data integrity. It is important to maintain logs

for different types of activities, such as system events, data access, and user activity. These include, but are not limited to, tracking all operations on user accounts and data, logging relevant metadata, and monitoring and logging access and attempted access to data and applications.

### **Backup, recovery, and disaster recovery (DR)**

Regulations generally have a focus on patient safety, which means considering many factors, including product quality and data integrity. A validated backup and recovery procedure may be helpful for fulfilling regulatory requirements related to data integrity and quality, and for allowing the solution to meet patient needs as discussed in the [GxP considerations](#) subsection of the [Backup, recovery, and disaster recovery \(DR\)](#) section.

### **Consent management**

In several areas of the Biopharma business, it is essential to have appropriate management over personal data. As mentioned earlier, for clinical trials, data might be subject to several regulations at the same time, such as the Clinical Trials Regulation (EU) 536/20141 and the GDPR (EU) 2016/679. However, both are requirements requiring clinical trial participants to be asked for consent for data processing.

Current regulations mandate that collecting and using personal data from participants must comply with rigorous standards. Among such standards are those which govern consent, which may require clinical trial managers to obtain freely given, specific, informed, and unambiguous consent before collecting participant data. Consent requirements may be set by local data privacy regulations like GDPR or be based in [Clinical Trials Regulation](#), among other authorities.

### **Continuous compliance**

As described in the [Compliance at scale](#) principal reference topic in this paper, point-in-time reviews are historically used to verify that the configuration of a biopharma solution meets its compliance standards per the organization's quality management system. AWS provides a set of services discussed in the [Appendix A: HCLS Regulatory and compliance principle concept reference](#) section related to [Continuous compliance](#) to implement a model of continuous compliance in the cloud.

### **Compliance at scale**

Consistently demonstrating that your environment is in a state of control is at the core of many regulations that apply to the biopharma industry. Because your environment grows, this becomes a more significant challenge. See the discussion in the [Appendix A: HCLS Regulatory and compliance principle concept reference](#) section related to

*Compliance at scale* for considerations relative to both enterprise-wide and dynamic resource scaling relative to the GxP and other regulatory requirements

## Providers

Providers have been experiencing a growing need to scale quickly, securely, and handle an increasing number of patients while improving the quality of care delivery. To meet this demand, providers often look to AWS to use the ever-growing number of [HIPAA eligible services](#) to innovate on behalf of their patients and staff. Although there are various deeper use cases that providers have expressed interest in, this section focuses on several common high-level examples of providers interested in using AWS to meet their needs.

### Applicability to principle concepts

The following sections provide examples of how each of the principle concepts outlined in the [Appendix A: HCLS regulatory and compliance principle concept reference](#) section apply to providers.

#### Data provenance

There has been an explosion of data in healthcare over the past decade. Patient rooms are becoming increasingly connected, with medical devices, electronic health records, medical images, and more, capturing tremendous amounts of data and making it difficult to ascertain what it all means about a patient. Surfacing this data to clinicians, and using AWS to derive additional insights, can build a better clinical picture of patients by empowering clinicians.

The HL7 organization's [Fast Healthcare Interoperability Resources \(FHIR\)](#) standard is an increasingly popular standard that customers are adopting to provide RESTful API access to clinical data and help with data sharing. [Amazon HealthLake](#), a fully managed FHIR data store with built-in medical comprehension, is a service that can be used to provide interoperability options and built-in medical comprehension so that both structured and unstructured data about patients can be brought together and provide a more holistic view of a patient population.

#### Data residency, transfers, and adequacy considerations

One primary challenge that providers face regarding data residency is that countries around the world have different requirements for healthcare data of their citizens. Some countries allow their citizen's healthcare data to leave their country, and for providers in those countries, it opens the door to the compute, storage, databases, and other AWS

services available because they can choose what Region they wish to use for their clinical data. Other countries have stricter data residency requirements, such that healthcare data for their citizens may not go beyond physical borders. Providers in those areas can use a variety of the AWS Edge services available to use the latest compute, storage, and database offerings.

Another challenge providers face today is collaborating with researchers and other clinicians around the world. Sharing data, in particular de-identified patient data, is critical for finding cures to diseases and providing more effective care. [AWS Data Exchange](#) is one service that can help with this by making it easier to find, subscribe to, and use third-party data in the cloud.

### **Data at rest and in motion**

One of the core components for adhering to nearly all healthcare compliance requirements and standards (such as those set forth by HIPAA or [HITRUST](#)) is encrypting data at rest and in motion. Beyond compliance, encrypting clinical data is also a good practice because, for healthcare providers, there is patient data everywhere. For example, encrypting data secures data from outside access but can also help prevent unintended disclosure of PHI on disparate systems and clinicians' personal devices that may access the data. For databases, especially where employees might have access, encrypting databases and enabling row and column level security controls also provide additional levels of control and security on critical data.

### **Data privacy**

Securing PHI is critical for maintaining patient privacy, but using social media and other devices by patients and visitors has made re-identification possible. Even staff can inadvertently disclose patient data through the use of social media (that is, taking photos of themselves at work with a patient in the background). Providers need to take extra precautions to verify data privacy and access. Even location-based games, where users collect rewards for visiting virtual landmarks in the real world, which are played by patients and visitors to healthcare organizations, can prove challenging to maintain patient privacy.

Data privacy is also challenging for ensuring that only those individuals with required access rights are the ones that can view patient data. Various types of patients and data access scenarios necessitate due diligence by healthcare providers to protect their clinical data, through measures such as robust logging, encryption, and access controls: celebrity patients, employee patients, minors, and non-care team workers or family members accessing patient data are a few examples.

## Change management

Broken IT systems can lead to adverse events if not working properly. The wrong data provided to a clinician can lead to a misinformed decision, or system outages can lead to system-wide issues that impact hospital operations. Having proper change management is vital to ensuring that healthcare organizations continue to function and that systems are working more effectively. One key to this is deploying IaC (for example, using AWS CloudFormation) so that manual processes and mistakes can be decreased. Using multiple AWS accounts can also provide better change control for providers so that the scope of impact of a system issue or security event is minimized.

Another key aspect of change management in provider organizations is that it helps align and enforce the other security principles mentioned in this document. For example, insider access to patient data is one of the most significant risks that healthcare organizations must face, and manually deploying code to servers provides admin access to employees who shouldn't have access to patient data stored there. Through proper change management and using AWS controls and services, providers can reduce the number of people who need to interact with production systems.

## Logging and auditing

Healthcare providers need robust logging and auditing capabilities because they must know who has viewed what clinical and business data, what systems were accessed (both inpatient care and back-office operations), how to debug systems, drive real-time alerts, and perform reporting to verify that all systems are performing efficiently. Larger provider systems often contain mission control centers where staff monitors all systems in real-time so that staff can be called in at any time for fixing system outages. By using [Amazon CloudWatch](#) and [AWS CloudTrail](#), customers can create real-time dashboards and alerts for meeting all of their logging and auditing needs.

Beyond the business operations, logging and auditing can also be used for advanced analytics, creating new artificial intelligence and machine learning (AI/ML) models, and ultimately be used for improving care delivery, patient outcomes, financials, and improving system-wide operations.

## Backup, recovery, and disaster recovery (DR)

Clinical and business data for healthcare providers is one of their most important assets. In particular, disaster recovery (DR) for providers could potentially mean life or death for patients if systems cannot be brought back online quickly. For data retention, clinical records for patient care must be backed up and available for typically six years per HIPAA guidelines. However, this can change depending on the country and the

provider's requirements. Cost concerns are also of importance given the volume of data that providers are required to store. [Amazon S3](#) and [Amazon S3 Glacier](#) are ideal for this, given that they provide 99.999999999% of durability, are highly cost-effective, and can have retention policies enforced by using Amazon S3 Lifecycle policies.

### Consent management

Privacy management and complying with patient consent are critical for ensuring that only proper health record is shared with the respective parties that have been given access. Violating consent can lead to hefty fines for healthcare organizations, so it is vital that healthcare customers validate that patients have granted proper consent for viewing their data, and procedures and various options when it comes to how care is delivered. For this, AWS has solutions on [AWS Marketplace](#) that are focused on patient consent, all of which are built using HIPAA eligible services and were technically validated by both the solution provider and AWS.

### Continuous compliance

Ensuring that all systems used to treat patients meet desired requirements and perform as expected is essential for patient safety. Healthcare providers have rigorous standards and checklists to validate systems and routinely verify that they are working as expected. Often this involves including clinicians and clinical leaders early in software and device procurement processes.

Provider organizations, who also tend to purchase systems, will also initiate many requests for proposals throughout the year and expect vendors, who respond to those, will have validated and verified their solutions while also meeting regulatory and compliance requirements. AWS Artifact, AWS Config, and Audit Manager are three common tools software vendors and healthcare customers use to meet and enforce their compliance needs. For example, AWS Config offers a [conformance pack focused on HIPAA](#) so that customers can quickly deploy AWS Config rules that enforce requirements for HIPAA.

## Payors

Payors are one of the most highly regulated entities of the U.S. healthcare system and are subject to many compliance regulations from state and federal agencies to enable data privacy and information security. Payors are using the cloud to innovate more quickly and deliver better and faster health outcomes for their members. Payors collect, store and manage health, financial and demographic data to support their members and

have access to the arguably most holistic patient view (as compared with other covered entities), collecting and maintaining the following types of data:

- Medical and pharmacy claims
- Membership eligibility and demographics
- Benefit plans and contracts
- Provider networks and provider profiles
- Member health and wellness program data

The cloud provides great efficiency and benefits to manage the above data types, but payors have to validate those workloads are [compliant with regulatory requirements](#). Payors are subject to various frameworks, guidelines, and regulations \ which they have to comply with to do business, such as:

- [FedRAMP \(NIST guidelines\)](#)
- HIPAA and HITRUST
- Institutional review boards and ethics committees
- GDPR

In accordance with the AWS shared responsibility model, it's the responsibility of the payors to secure their workloads and data within their AWS Cloud environment while adhering to the compliance requirements set forth by the regulators. This section explains the different compliance principle concepts. Payors need to take this into account while operating their workloads in the AWS Cloud. The section also explores services and best practices, which can help payors secure and run sensitive workloads in alignment with their compliance requirements.

## Applicability to principles

The following sections provide examples of how each of the principle concepts outlined in the [Appendix A: HCLS regulatory and compliance principle concept reference](#) section apply to payors.

### Data provenance

Data provenance regulations stem from GDPR in the EU, and HIPAA in the U.S. Generally, the practice of data provenance requires the recording of data origin and its transformation history. The data provenance regulations can facilitate an audit trail of data housed by payors to ensure integrity, data privacy, and information security. To

have a good data provenance in practice, payors should capture origin for the data originated internally and externally by entities such as providers, labs, and pharmacies. Payors commonly store data containing PHI and/or PII in the following categories, and these categories are subject to data provenance regulations:

- Member demographics and relationships
- Medical claims and financial information
- Electronic health records
- Medical devices data and images

Meeting data provenance requirements fall under the responsibility of the payor customers, who have to make sure there is proper logging of data coming in and any translation, transformation is captured so that validation of the records can be performed. Consider an example when a payor receives a medical claim from a provider (following the EDI 837 standard). The payor can use Amazon S3 object storage service to store this claim file in the original format, with versioning enabled on the Amazon S3 bucket. A transformation or translation performed on the original claim will result in a new version of the Amazon S3 object, which can support data provenance compliance requirements.

### **Data residency, transfers, and adequacy considerations**

Data residency compliance requires storing data in a particular geographic location, based on the country or region of origin. Payors have vast amounts of sensitive data to support their members and adhere to strict data residency regulations to [safeguard their members' privacy](#). To further protect their member's privacy, payors often have requirements to route network traffic within a country or geographic region while processing workloads. To support their adherence to these requirements, payors can use [AWS Regions](#), [AWS Local Zones](#), [AWS Outposts](#), and [AWS Direct Connect](#) to deploy their workloads in targeted geographic locations.

Additionally, data residency can present a challenge for analytics workloads; for example, payors create risk pools to segment their members, which helps calculate accurate premiums. Risk pools require the payor to access a large set of claims data, but if claims data can't be shared across multiple geographies, the models are less accurate, resulting in potential high premiums. Payors can use [Amazon Comprehend Medical](#) as part of their data pipeline to remove PHI and de-identify data to help mitigate these challenges while still adhering to data residency requirements.

## Data at rest and in motion

In the U.S., according to [reports published by the Department of Health and Human Services \(HHS\) of disclosures of unsecured PHI](#), patients have been impacted by inadvertent or unintended disclosures. Encryption at rest and in motion is typically the minimum requirement to adhere with applicable regulations when it comes to storing health data. Payors have sensitive PHI and PII from their members, and it is the payor's responsibility to safeguard this data. AWS offers a number of general-purpose data storage services, which payors can use to store numerous types of data. Payors can employ AWS services such as [server-side encryption with Amazon S3-managed keys \(SSE-S3\)](#), [client-side encryption \(CSE\) with Amazon S3](#), and [Transparent Data Encryption \(TDE\) with Amazon RDS](#) when storing data which might have PHI or PII.

Consider an example of data at motion. As part of the 21st Century Cures Act, [Centers for Medicare and Medicaid Services \(CMS\)](#) rules require health data interoperability between healthcare entities, and require payors to share complete patient health and claims history. When sharing data with other health entities, payors can employ AWS services (such as AWS KMS) to encrypt data at rest while ensuring all traffic uses Secure Sockets Layer (SSL) with TLS in the Amazon API Gateway to make sure data in motion is encrypted as well. Payors can also use predefined rules or create their own rules in AWS Config, automatically alerting them of any noncompliant resources in their AWS account.

## Data privacy

With social media and fitness trackers, PHI data is collected, transmitted, analyzed and shared for the benefit of the user but also for advertising purposes. Data privacy is a central tenet of several regulations (e.g. HIPAA and GDPR) which provide for and enforce patient control over their personal health records and require healthcare entities such as payors not to share PHI and PII without proper consent.

Regulations require payors to take extra precautions in protecting their member's privacy. Because they have a vast variety of health, claims, and financial records in their possession, securing sensitive data from unauthorized access is critical. Payors have to implement robust access controls to ensure only payors with required access can view patient data. Payors can use AWS services such as [Amazon CloudWatch Logs](#) for logging of access logs, AWS KMS to encrypt PHI data, and AWS IAM service to implement robust access controls. Additionally, payors can leverage a number of AWS Marketplace solutions which can provide a turn-key solution for monitoring against regulatory compliance (such as HITRUST) and validate that only the right resources have access to sensitive data.

## Change management

The [Sarbanes-Oxley \(SOX\) Act of 2002](#) established guidelines for IT system governance to prevent fraudulent accounting activities, including guidelines in the healthcare setting, particularly to payors. Today, payors manage and run vast amounts of financial systems used for accounting, claims processing, and payment fulfillment, and these systems are subject to SOX compliance. Payors are constantly innovating on behalf of their members and making it easier to receive care effectively.

To adhere to compliance requirements, change management guidelines often have to be incorporated in the software development life cycle (SDLC) process. Payors may find multiple AWS services helpful for change management requirements for workloads running in the AWS Cloud. It's a best practice to use [AWS Organizations](#) to segregate production and non-production AWS accounts. For example, in a production environment, payors can implement change management procedures in order to control the delivery of changes to the solution in the AWS account.

Services such as AWS Config and [AWS Systems Manager Change Manager](#) also can help payors to verify compliance after a change has been implemented. One of the greatest advantages of the cloud is the ability to use IaC, where changes can be tracked and managed in version control systems such as [Git](#). AWS offers several services to enable this capability, such as [AWS CloudFormation](#) and [AWS Cloud Development Kit \(AWS CDK\)](#). To minimize human error during changes, payors can use a comprehensive suite of AWS DevOps tools (such as [AWS CodeDeploy](#), [AWS CodeCommit](#), and [AWS CodeStar](#)) to automate deployments into higher environments.

## Logging and auditing

Several regulations require logging, and these requirements can include logging of healthcare systems, audit trails and logs of system and configuration changes, applications, and user network activity. For example, in the U.S., payors must comply with [HIPAA logging and auditing requirements](#) due to the nature of data they have in their possession and the type of systems they run.

For example, when a payor receives a medical claim from a provider, the provider may need to log that network transaction and payload, along with the system processing of the claim. In addition, if there are any errors or warnings in the processing of the claim, it is considered best practice for those issues to be logged and stored for troubleshooting and root-cause analysis purposes. For logging of processes, applications and user interactions, payors can use [Amazon CloudWatch Logs](#) (general-purpose logging service) to capture telemetry data from the system. Logs can

eventually be stored in Amazon S3, or Amazon S3 Glacier for long-term archival, allowing historical auditing using tools such as [Amazon Athena](#) and [AWS Glue](#).

For network traffic logging, payors can use [Amazon Virtual Private Cloud \(Amazon VPC\)](#) with [VPC Flow Logs](#) (network logging service) to capture incoming and outgoing network requests along with the security protocol, HTTP headers, and the payload, giving auditors visibility of history network traffic in case of a security event.

Payors can use [AWS CloudTrail](#) to make sure all cloud configuration changes are logged and available for auditing needs. For a more automated auditing solution, payors can use AWS Audit Manager to continuously audit their usage. Audit Manager simplifies how to access risk and compliance by utilizing a framework of prebuilt compliance templates which makes it easier to validate sensitive workloads.

### **Backup, recovery, and disaster recovery (DR)**

Regulations such as the [HIPAA Security Rule](#) require covered entities such as payors to have a backup, recovery, and disaster recovery (DR) plan in place. Additionally, to mitigate site-wide outages, especially when it comes to PHI, which might impact patient health outcomes, regulations require an off-site backup along with an off-site DR plan. [RTO and RPO](#) must be taken into consideration to minimize patient impact in case of a system failure.

All data backed up must be stored in an encrypted manner and properly safeguarded to prevent unauthorized access. With the growing number of incidents (such as unintended disclosures) in recent years, payors must have comprehensive backup, recovery, and DR practices. Payors manage sensitive member financial data, along with PHI. It's essential to have comprehensive backup policies and procedures in place such that in the event of a security event, systems can be brought back online with minimum disruption to patient care.

Payors can use AWS Backup to centrally manage and automate backups across AWS services on a given frequency and with a data retention policy. To store sensitive backup data in an encrypted manner, payors can use Amazon S3 server-side encryption (S3-SSE) to store data, and use retention policies to move data to an archival state in Amazon S3 Glacier.

For mission-critical workloads, such as claims processing, payors can use multiple [AWS Regions](#) to distribute their workloads using active-passive or active-active configurations. Data can be replicated across AWS Regions on a given frequency using [AWS Batch](#) or Amazon S3 cross-Region replication, or use Amazon RDS cross-Region replicas. Additionally, payors can use IaC services such as CloudFormation and AWS

CDK to automatically recreate degraded runtimes in a new target Region to minimize RTO.

## Consent management

Several regulations, including the HIPAA Privacy Rule, afford patients the right to approve or deny sharing of their PHI with other healthcare entities. Healthcare entities often have to manage patient consent at a global level and make sure the consent is updated across all their systems when there are changes. Payors have to make sure patient consent is managed centrally when sharing PHI and PII information with other healthcare entities.

Payors can store consent using multiple mechanisms, and depending on the complexity of the solution the payor needs, the payors can leverage relational/object databases when needing a to manage uncommon consent requirements. AWS Marketplace can also provide solution offerings for more common use cases.

As an example, if a patient updates their consent for a given service, it may need to be propagated across the payor's IT systems; [Amazon Simple Notification Service \(AWS SNS\)](#) can be used in a [publish-subscribe](#) (pub-sub) model, which can help systems propagate consent. AWS Marketplace offerings can also be used to accomplish this as well.

## Continuous compliance

Laws require healthcare entities to maintain compliance with rules outlined in HIPAA and GDPR throughout the lifecycle of a workload, and as an operational process. Being out of compliance can result in hefty fines along with the risk of data or security events. Payors have to maintain compliance on all systems which handle PHI, PII, financial, and claims data. Payors need to serve their members with new and innovative experiences, but while making those changes, constant compliance is a must.

After the payor workloads are deployed to the cloud, AWS offers tools to validate and verify them. Payors can use IaC tools such as CloudFormation and AWS CDK to define and provision their infrastructure. CloudFormation includes the ability to run drift checks against the current environment, which helps validate the deployed environment. Payors can also use Audit Manager to continuously audit their AWS usage to simplify how they assess risk and compliance.

For more holistic configuration management, payors can use AWS Config, a service that enables them to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records customers' AWS resource

configurations and allows customers to automate the evaluation of recorded configurations against desired configurations. Payors also have access to the compliance reports of the AWS Cloud, using AWS Artifact, providing on-demand access to AWS’s security and compliance reports and select online agreements.

## How can AWS support you?







AWS has four intrinsic ways to help a customer align with a compliance and security framework in AWS:

- Do it yourself (DIY) route
- Partner route
- Professional service route
- Hybrid of two or more routes

### DIY route

For the DIY route, AWS offers many services to help customers protect their data and align with security and governance frameworks.

#### AWS security, identity, and compliance solutions

 Identity & access management	 Detection	 Infrastructure protection	 Data protection	 Incident response	 Compliance
AWS Identity & Access Management (IAM) AWS Single Sign-On AWS Organizations AWS Directory Service Amazon Cognito AWS Resource Access Manager	AWS Security Hub Amazon GuardDuty Amazon Inspector Amazon CloudWatch AWS Config AWS CloudTrail VPC Flow Logs AWS IoT Device Defender	AWS Firewall Manager AWS Network Firewall AWS Shield AWS WAF – Web application firewall Amazon Virtual Private Cloud (VPC) AWS PrivateLink AWS Systems Manager	Amazon Macie AWS Key Management Service (KMS) AWS CloudHSM AWS Certificate Manager AWS Secrets Manager AWS VPN Server-Side Encryption	Amazon Detective CloudEndure DR AWS Config Rules AWS Lambda	AWS Artifact AWS Audit Manager

#### DIY route

The extensive [AWS documentation](#) is a great way to get started with the DIY or *build* approach. Keep in mind that AWS has published many whitepapers and solutions to facilitate getting started:



- [AWS Whitepapers & Guides](#)
- [AWS Solutions Library](#)

## Partner route

AWS has a comprehensive cloud partner network, which includes many partners who are verified experts in a particular subject area where customers need help.

The [Amazon Partner Network \(APN\)](#) is the global community of Partners who use AWS to build solutions and services for customers. AWS helps Partners build, market, and sell their AWS offerings by providing valuable business, technical, and marketing support.

There are tens of thousands of AWS Partners across the globe. More than 90% of Fortune 100 companies and the majority of Fortune 500 companies use AWS Partner solutions and services. As a result, AWS Partners are uniquely positioned to help businesses take full advantage of all that AWS offers and accelerate their journey to the cloud.

APN Partners are categorized based on the [path](#) upon which their business is based:

- **Software path partners** are organizations that produce and sell software that runs on or is integrated with AWS.
- **Hardware path partners** are organizations that develop hardware devices that work on AWS. Hardware Path Partners include original equipment manufacturers, semiconductor manufacturers, and network carriers.
- **Service path partners** are professional, consulting, or managed services firms that help customers of all types and sizes accelerate their journey to the cloud.
- **Training path partners** are organizations that sell, deliver, and incorporate AWS training.
- **Distribution path partners** are organizations that recruit, onboard, and support their Partners to resell and develop AWS solutions.

## Professional service route

Adopting the AWS Cloud can provide you with sustainable business advantages, and supplementing your team with specialized skills and experience can help you achieve those results. The [AWS Professional Services](#) organization is a global team of experts

that can help you realize your desired business outcomes when using the AWS Cloud. AWS Professional Services works with the customer's team and chosen members of the APN to implement your enterprise cloud computing initiatives.

## Hybrid approaches

Not every customer has the resources and knowledge to do everything on the cloud on their own. Often it is best to dedicate resources to the specific value add or intellectual property the customer brings and let someone else handle the undifferentiated heavy lifting of services not tied to the company's code value add or intellectual property. But even in that case, it does not mean that the customer is ready to give complete control to a third party, so a hybrid approach is often a common architectural decision where the customer retains substantial control of the AWS environments while contracting a third party to (for example) handle the security operations center with *actual people* monitoring the environments 24/7, handling the smaller alerts and security events, and involving the customer's core team for the more complex events and tasks.

## AWS Professional Services – regulatory and compliance practice

AWS HCLS Professional Services GxP compliance practice supports HCLS customers for seamless migration or development of regulated workloads into AWS. The AWS HCLS Professional Services GxP compliance practice provides advisory service and guidance on building robust, compliant, secured, and highly available regulated and controlled systems on AWS.

[AWS Professional Services](#), AWS HCLS Professional Services, GxP compliance team, is comprised of customer-facing GxP compliance and AWS experts who partner with the AWS field teams (service teams, local practice managers, overlay sales teams, and Partners) to help customers in their compliance journey by providing guidelines to the team of consultants, providing guidance on security and compliance, audit management, cloud governance, and operation related best practices for regulated workloads. AWS HCLS Professional Services GxP compliance practice team acts as a trusted advisor partner for our customer to migrate or build GxP compliant workloads on AWS using automation strategies.

AWS HCLS Professional Services GxP compliance offerings:

- Lead data analytics and strategies towards data privacy and security with customers.

- Drive automation of audit management expectation and governance.
- Facilitate and drive [GxP assessment of existing workloads on AWS](#).
- Identify the scope of GxP compliance opportunities within customer accounts.
- Help customer's challenges in GxP compliance, cloud operation, and governance for managing regulated workloads.
- Help migrate their GxP workload (GMP, GCP, GLP) to AWS.
- Help customers in [automating Installation Qualification \(IQ\)](#) and Operational Qualification (OQ) reports and achieve [continuous GxP compliance](#).

The AWS HCLS Professional Services GxP compliance practice team, in collaboration with other AWS Professional Services SMEs, created and released various offerings to migrate GxP workloads seamlessly into AWS. Teams can use this as a starting point for GxP compliance engagements with best practices. For example:

- Align offering – Building GxP Compliance on AWS
- Launch offering – Rapid GxP compliant apps on AWS
- GxP allow listing of AWS Professional Services
- Regulated landing zone
- Rehost qualification for migration

## How to get engaged with the GxP compliance practice team?

AWS Field Teams such as Service Teams, Customer Practice Teams, and Account Teams, based on the customer's needs, engage with AWS HCLS Professional Services GxP Compliance Practice Team to help customers in their compliance journey.

## Conclusion

This whitepaper is intended to provide a baseline, and as such, the expectation is that the reader of this document will be inspired to continue their learning journey by continuing to seek out additional training, contribute to future improvements in this whitepaper, and to share this knowledge with others.

The content within this paper is not meant to be exhaustive, but hopefully comprehensive enough to give you an understanding of many of the concerns and considerations associated with implementing solutions using AWS services in regulated

projects, regardless of which Region your project will run within. Although this paper lists many of the more common global regulations and frameworks, this paper does not dive deeply into the specific regulations required for a specific country. Therefore, it is left up to the reader to do their due diligence within a specific local region to understand any regulations that may affect the concepts discussed.

Finally, the world is constantly changing, as are the regulations and certifications needed to operate within it. Therefore, this document is intended to be updated over time to encompass new scenarios and enhance how AWS approaches each principle concept documented here.

## Contributors

Contributors to this document include:

- Olivia Choudhury, PhD, Partner Solutions Architect, HCLS, AWS
- Senthil Gurumoorthi, Global HCLS Security Assurance Lead, Security & Compliance, AWS
- Umair Khalid, Partner Solutions Architect, HCLS, AWS
- Ajay Mohindra, PhD, Sr. Partner Solutions Architect, HCLS, AWS
- Will Quiles, Sr. Security Specialist Solutions Architect, AWS
- Tom Richards, Sr. Partner Solution Architect, HCLS, AWS
- Harvey Ruback, Sr. Partner Solution Architect, HCLS, AWS
- Ian Sutcliffe, World Wide Tech Lead, HCLS Compliance, AWS
- Brian Warwick, Sr. Partner Solution Architect, HCLS, AWS

## Further reading

For additional information, refer to:

- [Applying the AWS Shared Responsibility Model to your GxP Solution](#) (blog post)
- [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) (whitepaper)

- [Automating the Installation Qualification \(IQ\) Step to Expedite GxP Compliance](#) (blog post)
- [AWS Best Practices for Security, Identity, & Compliance](#)
- [AWS Compliance](#)
- [AWS Data Privacy Center](#)
- [AWS Data Privacy FAQ](#)
- [AWS Data Residency: AWS Policy Perspectives](#) (whitepaper)
- [AWS GDPR Center](#)
- [AWS GDPR Data Processing Addendum](#)
- [AWS GDPR Data Processing Addendum – Now Part of Service Terms](#) (blog post)
- [AWS Privacy Notice](#)
- [AWS Reference Architectures](#)
- [AWS Shared Responsibility Model](#)
- [AWS Sub-processors \(GDPR\)](#)
- [AWS Well-Architected Framework](#)
- [GxP Systems on AWS](#) (whitepaper)
- [HIPAA and HITRUST on AWS](#) (blog post)
- [HIPAA Eligible Services Reference](#)
- [HIPAA on AWS QuickStart](#)
- [HITRUST Reference Architecture](#)
- [Logical Separation on AWS](#)
- [Navigating GDPR Compliance on AWS](#) (whitepaper)

## Appendix A: HCLS regulatory and compliance principle concept reference

A typical method for understanding regulatory and compliance considerations with respect to data is to look at the regulations, certifications, and frameworks that act as the guardrails that define what is important. These governing laws and best practices are based on feedback from the public, but also from lessons learned, best practices, and proven patterns. Over time, this body of knowledge has been distilled down into a common set of concept categories that are principally self-evident regardless of the framework or set of laws defining them.

The following section does not attempt to be exhaustive with defining all of these concepts, however for the HCLS industry, several categories stand out as central to the concerns of the industry, relevant stakeholders, and—in particular—the patients.

The 10 principal category areas covered in this section are:

- *Data provenance*
- *Data residency, transfers, and adequacy*
- *Data at rest and in motion*
- *Data privacy*
- *Change management*
- *Logging and auditing*
- *Backup, recovery, and disaster recovery (DR)*
- *Consent management*
- *Continuous compliance*
- *Compliance at scale*

You can probably think of other common concepts that might be important additions to this list; however, these principal areas stand out as critical and paramount for ensuring that interactions between systems on behalf of a patient, practitioners, and other stakeholders is done in a way that is safe and secure. These concepts are vital focal areas that must be held to the highest standards when implemented using AWS tools and technologies. Having a solid understanding of this will help you understand how AWS has embraced these concepts across the entire infrastructure. In addition, it will

provide you with the background you need to understand how important the AWS shared responsibility model principles are across these concepts.

## Data provenance

Data provenance, also known as data lineage, tracks data from its source to the point where it is used for decision making. Metadata tracks the data from creation through movement, modification, and—ultimately—deletion.

This is often done to prove the integrity of the data (that is, that the data is complete, consistent, and accurate). In addition, the data should be shown to be [attributable, legible, contemporaneously recorded, original or a true copy, and accurate \(ALCOA\)](#).

## Tools and techniques

Tracking data lineage so that the location and data source is tracked and known during further processing can be achieved using [AWS Glue](#). Customers can visually map the lineage of their data to understand the various data sources and transformation steps that the data has been through. Customers can also use metadata provided by [AWS Glue Data Catalog](#) to establish data lineage. Similarly, the [Amazon SageMaker Data Wrangler](#) data flow user interface provides a visual map of the complete data lineage.

Any operations performed on the data should be captured in an audit trail that provides a tamper-resistant, computer-generated, time-stamped electronic record that reconstructs the course of events relating to the creation, movement, modification, or deletion of data. The audit trail should capture who did what, when, and—often—why they did it.

[AWS CloudTrail](#) logs all service application programming interface (API) calls and can be used to track what operations were performed and can be configured to log [data events](#) for Amazon S3 and [Amazon DynamoDB](#). For an Amazon S3 bucket, this includes object-level API operations such as `GetObject`, `DeleteObject`, and `PutObject`. For DynamoDB, this includes object-level API activity on tables such as the `PutItem`, `DeleteItem`, and `UpdateItem` API operations.

Auditing changes to data in other data stores depends on the AWS service being used. [Amazon Redshift logs information about connections and user activities](#) in your database. [Amazon RDS for MySQL](#) supports the [MariaDB Audit Plugin](#), and [Amazon Aurora MySQL](#) supports [advanced auditing](#). The audit trail should be immutable, as described in [Logging and auditing](#).

Many regulations require a true copy of the original data to be maintained securely throughout the retention period. This is usually accomplished through the use of backups. The backup should also contain associated metadata and audit trail. This is covered in more detail in the [Backup, recovery, and disaster recovery \(DR\)](#) section.

## Data residency, transfers, and adequacy

The term *data residency* is used interchangeably with various others and can mean slightly different things to different organizations. However, data residency essentially refers to the physical location and processing of an organization's data. Different regions or countries then impose legal or regulatory requirements on the data. Examples include the GDPR or HIPAA.

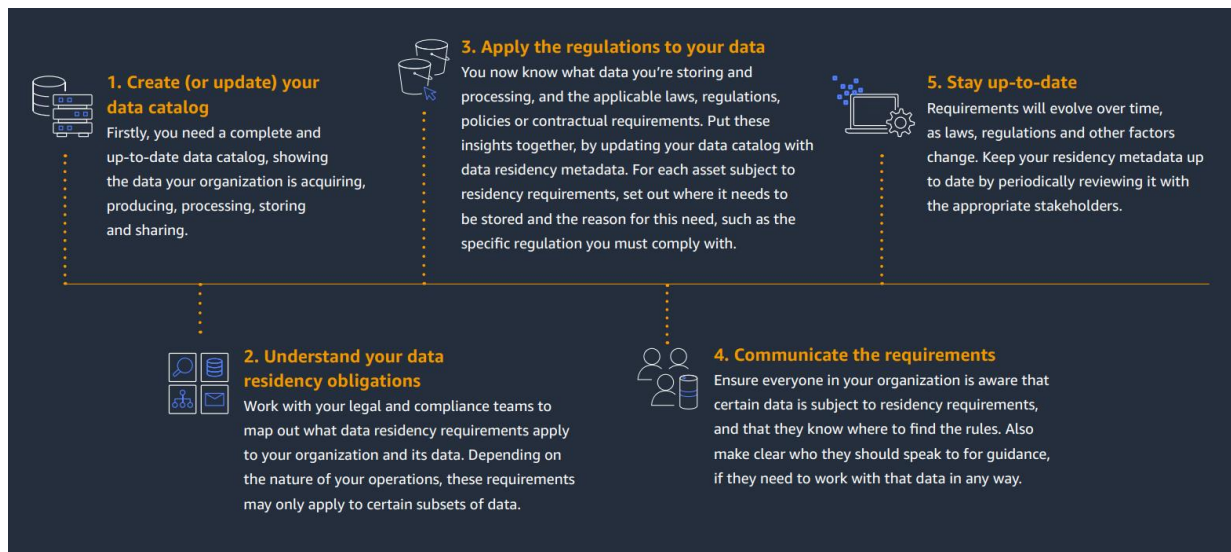
There are three terms which are commonly used to describe how data privacy is managed:

- **Data residency** – These requirements impose restrictions on where the data is physically stored.
- **Data sovereignty** – These requirements build on data residency by not only imposing restrictions on the location of the data but also the jurisdictional control over that data which affects the rights of access to the data.
- **Data localization** – Requires that specific data (or a copy) must remain within the borders of a country.

These rules not only impact where data is physically stored, but—often—also where data can flow, such as across country borders or across regions.

Some governments have determined that mandating data residency provides an extra layer of security; however, data residency often doesn't actually provide better security for a few reasons which are described in the [AWS Data Residency](#) whitepaper. Despite growing recognition of the limitations of data residency, there have been regulatory changes encouraging it in some specific areas. For example, DiGAV in Germany, which enables certain healthcare applications to be recognized as refundable under the German statutory health insurance system, builds on the German Social Code which states the data must be processed exclusively within the European Economic Area (EEA) or a country with an adequacy decision provided by the European Commission based on [GDPR Article 45](#). Because every organization will hold different data, and be subject to different laws, regulations, contracts, and policies, there's no one-size-fits-all checklist that will give a comprehensive understanding of your specific data residency

requirements. To establish what these requirements can be, AWS recommends going through the following [five-step process](#):



### *Five steps to consider*

Organizations worldwide store enormous amounts of data, so it's helpful to know where to start looking when carrying out your own data residency assessment. There are two main categories of information that typically fall under data residency requirements:

- **Personally Identifiable Information (PII)** – Virtually every organization stores or processes at least some PII, so this should be the first area you consider. For example, this might be financial information, such as an individual's transaction history or other personally identifiable details. Healthcare data is another example where organizations store information about patients and their medical histories. Public sector organizations, and their commercial partners, also hold large amounts of data about individuals created through their use of public services. This information can include addresses, dates of birth, financial data, details of someone's children or dependents, and their history of interactions with that organization.
- **National data** – This data is often considered sensitive and consequently subject to data residency requirements in certain jurisdictions. Examples include:
  - Geospatial information, such as maps and seismic data, which might be collected by either public sector bodies or private companies.
  - Data associated with the military, including intelligence, operations, and technology.

- Data about critical national infrastructure and resources, that might need to be stored within the country. This can include information about the design and operation of power-generation facilities, utility and communications networks, and transportation infrastructure.

After you establish and meet data residency requirements, you have four broad technology options for where to run your workloads and place your data:

- Run your own infrastructure in your chosen location
- Use the AWS Cloud
- Use AWS Local Zones
- [Bring cloud benefits to your own facility with AWS Outposts](#)

If you have specific data residency requirements, or have situations where you are not able to use an [AWS Region](#) that is closest to you, AWS can support your decision-making process when choosing a Region that best suits your regional requirements. Additionally, you retain complete control and ownership over the data in the Region where your data is physically located, making it easier to meet regional compliance and data residency requirements. A small number of AWS services involve the transfer of data, for example, to develop and improve those services, however in these cases you can opt-out of the transfer, or choose not to use the service because transfer is an essential part of the service (such as a content delivery service). You can refer to the [AWS Privacy Features](#) page for more information.

AWS complies with [many global data protection regulations](#), and AWS has services and tools to enable you to [address data residency](#) and build compliant infrastructure on top of AWS. Organizations from startups to enterprises and the public sector have access to infrastructure in their country to use advanced technologies including analytics, artificial intelligence, database, IoT, machine learning (ML), mobile services, and serverless, to drive innovation.

[AWS Local Zones](#) deploys AWS compute, storage, database, and other select services closer to a large population, industry, and IT centers where no AWS Region exists today. Although its primary purpose is to run latency-sensitive portions of applications close to end users, Local Zones can also be used to meet your data residency requirements.

[AWS Outposts](#) extends AWS infrastructure, services, APIs, and tools into your own data center, colocation space, or on-premises facility as a managed service. It's essentially an extension of the AWS Cloud, running on dedicated AWS infrastructure in a location

you specify. Outposts allow you to meet your data residency requirements while benefiting from AWS services, even where there's no AWS Region.

For an in-depth understanding of the security considerations related to data residency and how they apply to AWS, refer to [Data Residency: AWS Policy Perspectives](#).

## Data at rest and in motion

AWS recommends encryption as an additional access control to complement identity, resource, and network-oriented access controls. AWS provides additional services and features that enable customers to easily encrypt data and manage the encryption keys. All AWS services that store customer data offer the ability to encrypt that data. [AWS Key Management Service \(AWS KMS\)](#) integrates with many services to let customers control the lifecycle and permissions on the keys used to encrypt data on the customer's behalf. Customers can enforce and manage encryption across services integrated with AWS KMS through the use of policy and configuration tools.

Customers can control when data is decrypted, by whom, and under which conditions as it is passed to and from their applications and AWS services. Because access to encrypt or decrypt data within the service is independently controlled by AWS KMS policies that are set by customers, they can isolate control over access to the data, separate from access to the keys. This isolation model is a powerful additional logical separation control applied across a customer's AWS environment.

In addition to controlling how server-side encryption happens within AWS services, customers can choose to encrypt data within their own application environment using AWS KMS with client-side encryption, thereby taking AWS services out of their trust boundary. Application-level, client-side encryption can be used by the customer to enable a consistent security posture as data traverses within a customer's own service architecture, whether in AWS, on-premises, or in a hybrid model. The use of AWS KMS to manage the lifecycle of and permissions on keys provides a consistent access control mechanism for all encryption keys, regardless of where they are used.

To help prevent unauthorized use of encryption keys outside the boundary of AWS KMS, the service utilizes hardware security modules (HSMs) to protect customer key material while in use. These HSMs are validated under [Federal Information Processing Standard \(FIPS\) 140-2](#) with physical tamper response controls. The HSMs are designed so that plaintext keys cannot be used outside the HSM by anyone, including AWS employees. The only way keys can be used is when an authenticated and authorized customer request is received by the service. In response to the request, AWS KMS

enables the customer's key to be used within the HSM for an encryption or decryption operation.

The HSMs in AWS KMS are designed as multi-tenant in the sense that any customer's key can be used in any HSM within the Region. Like other AWS services that utilize multi-tenancy, AWS KMS is designed to isolate usage of keys only to the customer that owns the keys. There is no mechanism for an unauthorized user to cause a customer's key to be used. AWS KMS transparently manages the durability and availability of customer keys and can scale to support any number of keys at the rate customers' applications need to use them.

Customers simply manage the lifecycle and permissions on keys using the same authentication and authorization controls available to every other AWS service. Every request made of AWS KMS is logged to [AWS CloudTrail](#) to provide an audit of when keys were used and under what circumstances. AWS KMS is in scope for all accreditation programs supported by AWS that relate to data protection.

For customers with requirements to directly manage the HSM device that generates, stores, and uses their encryption keys, [AWS CloudHSM](#) is available as an option. CloudHSM offers a dedicated FIPS 140-2 Level 3 validated HSM and affords the flexibility of integrating with customer applications using industry-standard APIs such as PKCS#11, Java Cryptography Extensions (JCE), and Microsoft CryptoNG (CNG) libraries. It enables organizations to export keys to other commercially available HSMs for use in hybrid architectures. AWS automates the time-consuming administrative tasks around these HSMs, such as hardware provisioning, software patching, network routing, and creating encrypted backups of key stores.

Customers are responsible for scaling their CloudHSM environment and managing the crypto user accounts and credentials within the HSM. Like AWS KMS, CloudHSM is designed so that plaintext keys cannot be used outside the HSM by anyone, including AWS employees. Customers can combine the ease-of-use and integration with AWS services offered by AWS KMS with CloudHSM by using the AWS KMS custom key store option. Customers logically attach a CloudHSM cluster to an AWS KMS key identifier so that requests made to the key are authorized by AWS KMS but implemented on the customer's dedicated CloudHSM.

To protect data in transit, AWS encourages customers to use a multi-level approach. Network traffic between AWS data centers is transparently encrypted at the physical layer. All traffic within an [Amazon Virtual Private Cloud \(Amazon VPC\)](#) and between [peered VPCs](#) across Regions is transparently encrypted at the network layer when using supported Amazon EC2 instance types. At the application layer, customers have

a choice about whether and how to use encryption with a protocol like Transport Layer Security (TLS). All AWS service endpoints support TLS to create a secure HTTPS connection to make API requests.

**Note:** AWS has updated all AWS FIPS endpoints to a minimum TLS version of 1.2 across all AWS Regions, as of March 31, 2021. These updates revoke the ability to use TLS 1.0 and TLS 1.1 on all FIPS endpoints. No other AWS endpoints will be affected by this change.

For customer-managed infrastructure within AWS that needs to end TLS, AWS offers several options, including load balancing services (for example, [Elastic Load Balancing](#)), [Amazon CloudFront](#) (a content delivery network), and [Amazon API Gateway](#). To implement a TLS connection, each of these endpoint services allows customers to upload their own digital certificates to bind a cryptographic identity to the endpoint. Digital certificates are very difficult to manage at scale because they expire and need to be rotated. AWS simplifies the process of generating, distributing, and rotating digital certificates with [AWS Certificate Manager \(ACM\)](#). ACM offers publicly trusted certificates at no cost that can be used in AWS services that require them to terminate TLS connections from the internet. ACM also offers the ability to create a private certificate authority to automatically generate, distribute and rotate certificates to secure internal communication among customer-managed infrastructure.

Using services like AWS KMS, CloudHSM, and ACM, customers can implement a comprehensive [data at rest and data in transit encryption](#) strategy across their AWS environment to facilitate that all data of a given classification shares the same security posture. For further information, please refer to the whitepaper: [Logical Separation on AWS](#).

## Data privacy

At AWS, trust is built by not only demonstrating that the AWS Cloud is dependable, providing products and services with enhanced security features, but also by making certifications and attestations available for any customer to observe. AWS has a comprehensive [Security Assurance Program](#) that uses best practices for global privacy and data protection. These best practices help you to operate securely within AWS and to make the best use of the security control environment. These security protection and control processes are independently validated by multiple third-party independent assessments, and their attestations are available with [AWS Artifact](#).

AWS can support customers in fulfilling their data privacy requirements. For example, AWS affords customers full control of any content that they upload and the ability to configure access to resources and services, while AWS provides encryption, access, and logging features (for example, [AWS Identity and Access Management \(IAM\)](#), [AWS Organizations](#), and AWS CloudTrail). AWS will not access or use your content for any purpose without your agreement, nor will AWS ever use your content for marketing or advertising purposes. AWS will also not move your data outside of the Regions in which it is stored without your consent.

Security is a key to data privacy. One aspect of security relative to data privacy is ensuring your data at rest and in motion can be secured (refer to the [Data at rest and in motion](#) section). In both cases, your data can be encrypted with keys that you can create independently or AWS can create for you (through AWS KMS). There are data encryption capabilities available in more than 100 AWS services to support this scenario.

An important aspect of data privacy and the services provided by AWS, or any cloud provider, is that although a given service may be eligible for use across a set of regulations, it is possible for a service to be configured incorrectly for use with PHI data, bringing the usage of that service out of regulatory compliance. To help with this, AWS provides the whitepaper [Architecting for HIPAA Security and Compliance on Amazon Web Services](#), which outlines some of the best practices that would need to be implemented when using a service for healthcare, and possibly even life sciences, workloads.

Services such as [AWS Config](#), along with specific AWS Config [conformance packs](#), can support the creation of guardrails within an account to detect and alert on services that have been misconfigured. In addition to the tools, like AWS Config, which can be used to prevent or notify of changes to an account, there are [AWS Partners](#) that build solutions in addition to AWS specifically to help with regulatory compliance monitoring and governance.

Data privacy is central to why there are regulations to prevent PHI or PII from being accessed and used by individuals or organizations that should not have the right to do so. Regulations such as HIPAA, [HITECH](#), and [GDPR](#) (refer to the [Fundamentals](#) section) have controls that specifically address what an organization must do, at the very minimum, to prevent unintended broad access to information, and also what must happen if an unintended access does occur.

Customers should familiarize themselves with the [many local and international regulations](#) which pertain to the protection of data, and specifically PHI or PII. Most of

these regulations set limits and conditions on data use and disclosure and also give patients specific rights, such as getting access to their health records, requesting corrections, or even requesting the deletion of their data. Importantly, these regulations define what to do if there is a failure to comply, and this varies from country to country. For this reason, it is crucial to have attestations of compliance with regulations of the host countries where your data will be stored, and you can easily gain access to such documentation provided to you through AWS Artifact.

## Change management

A leading best practice for the deployment of environments, applications, and services is to remove as many manual steps as possible or for the deployment to be fully automated. The goal is to have no manual access to production environments except for break-glass situations.

Just as continuous integration and continuous delivery (CI/CD) refers to the idea of managing the lifecycle of software applications in an integrated, fully automated manner, the same ideas and techniques can be used to manage infrastructure. The ability to represent your infrastructure in a human readable format that can be interpreted by a computer system is called infrastructure as code (IaC). Because you have a code repository for software, containers, and so forth, you should also use the same tools to manage and store your IaC templates, whether you are using [AWS CloudFormation](#), [AWS Cloud Development Kit](#), or Terraform.

A powerful benefit of IaC is that by having templates that can be modified and reused, you can make sure that your security and compliance controls are in place as part of the template definitions. For example, you can enforce that encryption at rest is enabled, enforce TLS, and so forth.

As the requirements for the products, services, and applications change, both in and out of production, you need to define the process for evaluating, testing, and ultimately deploying the updated IaC template into production. This is what is referred to as change management, and at a minimum there are at least three areas to consider:

- **Low-risk change management** – Changes that happen often, follow standard operating procedures, and don't need approvals.
- **Standard change management** – This should be the process to manage planned releases. These can be scheduled, like every two weeks for small changes, and monthly or quarterly for more significant releases.

- **Break-glass change management** – Although uncommon, you should always be prepared for emergencies, security events, unintended data exposure, and ransomware. You should have at least a runbook of with standard operating procedures for dealing with these emergency situations where you have to react quickly.

As to specific recommendations from AWS, the [AWS Well-Architected Framework](#) Reliability Pillar has an excellent [Change Management](#) section. Lastly, note that [AWS Systems Manager](#) can be a powerful service to [automate and manage many change management tasks](#).

## Logging and auditing

The AWS recommendation for logging and auditing is to create a separate Log Archive account. The logs in this AWS account should be configured as read-only. This account is used as a consolidation point for log data that is gathered from [all the accounts](#) in the organization, and is primarily used by [your security, operations, audit, and compliance teams](#).

For example, in your primary account, AWS recommends that you consolidate AWS API access logs recorded in [AWS CloudTrail](#), changes to AWS resources recorded in AWS Config, and other logs that have security, compliance, and governance implications. If you use [VPC peering](#) between accounts, then you might also benefit from consolidating [VPC Flow Logs](#). It's a common practice to integrate this consolidated log data with a security information and event management (SIEM) solution.

If you use [AWS Control Tower](#) to manage your overall AWS environment, CloudTrail is automatically enabled in each account, and the CloudTrail logs are consolidated in an S3 bucket in a Log Archive account.

### Operational log data

Operational log data used by your infrastructure, operations, and workload owning teams often overlaps with the log data used by security, audit, and compliance teams. AWS recommends that you consolidate your operational log data into the Log Archive account. Based on your specific security and governance requirements, you may need to filter operational log data saved to this account. You may also need to specify who and what has access to the operational log data in the Log Archive account.

## Immutable log data

Not only should this data be encrypted at rest, but log data housed in the Log Archive account should be configured to be immutable (does not change). If using Amazon S3 for logs, with the proper IAM and bucket retention policies, Amazon S3 can provide for an effective log data store that is immutable. Another benefit of using Amazon S3, for long-term retention of logs, is to consider [Amazon S3 Glacier](#) with [Amazon S3 Glacier Vault Lock](#), which allows you to easily deploy and enforce compliance controls for individual Amazon S3 Glacier vaults with a vault lock policy. You can specify controls such as write once, read many (WORM) in a vault lock policy and lock the policy from future edits. After locked, the policy can no longer be changed.

## Managing access to this account

AWS recommends that you house log data only in the Log Archive account and refrain from including workloads in this account that act on the log data. By doing so, you can greatly limit access to this account. Workloads and tools that need to consume the consolidated log data are typically housed in your other accounts and are granted cross-account access through IAM roles to access the log data in a read-only, least privileged manner.

## Using this account

After creating a central logging archive, it can be used for several corporate needs, including audit, security, and business intelligence. You can use services such as [AWS Audit Manager](#) to collect evidence and create compliance reports. In addition, tools like [Amazon OpenSearch Service](#) (or like third-party Splunk) can then be used to monitor your data proactively. With Amazon OpenSearch Service, besides analysis, you can also set up alerting and anomaly detection. You can set up alerts to receive notifications when your data exceeds certain thresholds and pair anomaly detection with alerting to ensure you're notified as soon as an anomaly is detected.

## Backup, recovery, and disaster recovery (DR)

The concepts of [backup, recovery, and disaster recovery \(DR\)](#) are intertwined, and many AWS customers mix and match these terms often, even though the meaning often varies from customer to customer. Some customers might say they want DR, but when you dive deeper and learn what they are really after, you might find that a more straightforward backup solution is all the DR they need. Sometimes the terms come from a framework the customer is trying to comply with (for example, HIPAA or [PCI DSS](#)), which might help clarify the requirement.

## Backup

[Backup](#) loosely refers to making a copy of something with the expectation that the copy may be needed again in the future. A backup can be, for example, an hourly copy of a working database, or it can be a copy of an Amazon EC2 golden image, updated every time a new golden image is created or changed. A backup can also mean a copy of your [CloudFormation](#) template, perhaps part of a CI/CD pipeline. A backup can also mean taking the logs captured by [AWS CloudTrail](#), [Amazon CloudWatch](#), or other services, and copying those logs to another account or Region.

Another important aspect is maintaining a log of backups, including if they were successful (if not, evidence that the problem was resolved should be included). Also, it is important that the frequency and retention periods—both often independent—are clearly defined. Frequency refers to how often you take those backups.

- How often are backups taken? Are backups being taken fully or incrementally (only what changed from prior backup)?
- How long are these backups retained? What are the business reasons, governance, or compliance rules driving these retention periods? What mechanisms may be used to enforce these retention periods?

Understanding the answers to these questions may help drive what architecture choices need to be made to address customer needs.

[AWS Backup](#) is a service that can enable a centralized backup capability. [AWS Backup Audit Manager](#) provides built-in compliance controls and allows you to customize those controls to define your data protection policies. It is designed to automatically detect violations of those defined data protection policies and prompt you to take corrective actions. With AWS Backup Audit Manager, you can continuously evaluate backup activity and generate audit reports that can help demonstrate compliance with regulatory requirements.

## Recovery

Making backups is generally much easier than the matching action: recovery (or restoration). Recovery means taking a backup and making it the main object, file, software, or image. If you make backups and never test them, you will not know if those backups were correctly taken until you actually need to restore them. The worst time to test a backup and recovery process is when you actually need to use it to recover with, so ensuring that you have a properly tested recovery process in place is key.

## Disaster recovery

An environment or application faces technological and nature-related risks. As part of every solution, architects should list the applicable risks, and then identify how to deal with those risks. Although there are many ways of dealing with specific risks, these are four general categories as described by the (ISC)<sup>2</sup> blog post [Treating Risks](#):

- **Mitigate the risk** – Typically, by reducing the associated threat, vulnerability, and impact.
- **Avoid the risk** – By not doing something risky.
- **Transfer the risk** – For example, entering into a binding contract that places costs and liabilities on them for security events.
- **Accept the risk** – Accept that the costs of mitigating, avoiding, or transferring the risk outweigh the advantages (as you perceive them).

If a customer says that they need DR, this is a good opportunity to dig deeper and learn more what is behind that question. Likely this is just a small portion of a larger issue. You must first understand what risks they are trying to deal with and what expectations they have for each risk. You then need to understand how the customer plans to address each of those risks (mitigate, avoid, transfer, or accept).

In addition, when discussing DR strategies, it is imperative to understand the business objectives, which are often expressed in terms of Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs). RTO and RPO targets are based on business governance needs and are often based on service-level agreements (SLAs) with customers or suppliers.

- **RTO** – The maximum acceptable delay between the interruption of service and restoration of service. This determines what is considered a sufficient time window when service is unavailable.
- **RPO** – The maximum acceptable amount of time since the last data recovery point. This determines what is considered an acceptable loss of data between the last recovery point and the interruption of service.

Although availability (or high availability) and DR rely on similar practices for monitoring, deployment, and failover, availability focuses on components of the workload (like using managed services and several Availability Zones), whereas DR focuses on discrete copies of the entire workload. In other words, DR's objective is to focus on recovery after a disaster.

## GxP considerations

Regulations commonly require that patient safety, product quality, and data integrity should not be compromised by system failure or breakdown. Any regulated company should perform business continuity planning to actively protect its ability to continue to provide services to the public and comply with regulatory requirements. Part of the business continuity planning can be to use backups to restore data in case of failure. Therefore, regular backup of records, data, and software should be made to a safe storage location that is adequately separated from the primary storage location.

Data integrity is a crucial aspect of many regulations, for example, as described in the [FDA Title 21 CFR Part 11](#) or [EU Annex 11](#). An important aspect of data integrity is preventing the loss of any data and ensuring that data is readily available for review for the lifetime of the data. There are many ways to reduce the risk of data loss including scheduled backups and synchronous data replication, depending on the solution's individual requirements.

Regardless of the mechanism used, the processes and procedures should be verified when they are established and tested on a regular basis, including restore capabilities. Evidence of testing should be made available to inspectors, along with evidence showing backups or data synchronization functions are happening and are being retained according to data protection policies.

## Consent management

Consent management is a process used by covered entities, can be defined as anyone who provides treatment, payment, and operations in healthcare. Consent management allows patients to control and determine what PHI they are willing to share with stakeholders providing care (for example, providers, Payors, and clinical research organizations). Government regulations such as the GDPR and the [HIPAA Privacy Rule](#) provide legislated rights that an individual can express wishes concerning their personal health information, such as the *right to be forgotten*, the *right to withdraw*, and *consent to use the data*, among others. Consent management is a complex topic and often changes depending on governance structure and governmental jurisdiction.

In many global regions, consent management is a defining concept that must be considered. Examples include:

- **Individual Choice Principle (HIPAA Privacy Rule)** – Provides patients with the ability to determine how to manage and share their PHI data. These guidelines include the right to access certain health information maintained about the individual; the right to have certain health information amended; the right to receive an accounting of certain disclosures; the right to receive a covered entity’s notice of privacy practices; the right to agree or object to or authorize, certain disclosures; the right to request restrictions of certain uses and disclosures; and provisions allowing a covered entity to obtain consent for certain uses and disclosures.

The HIPAA Privacy Rule permits, but does not require, a covered entity to voluntarily obtain patient consent for uses and disclosures of PHI for treatment, payment, and healthcare operations. Covered entities have complete discretion to design a process that best suits their needs.<sup>2</sup>

- **[Article 7 of GDPR](#)** – Provides governance that where processing is based on consent, the controller (e.g. the natural or legal person, public authority, agency or other body) shall be able to demonstrate that the data subject (e.g. an identified or identifiable natural person) has consented to processing of his or her personal data. Further, the data subject shall have the right to withdraw his or her consent at any time.

The [AWS Partner Network \(APN\)](#) includes many companies providing healthcare, life science, and compliance-related solutions and services. For the latest listed partners that meet your specific use case needs, see [AWS for Health](#).

## Architectural considerations

The [AWS shared responsibility model](#) is key to understanding the roles that AWS and customers play in consent management. AWS is responsible for features and capabilities that are *of the cloud*, while the customer is responsible for everything *in the cloud*. This is a critical distinction because AWS does not provide specific consent management services; however, AWS does provide many purpose-built data stores and services that can be utilized to meet your individual solution use case requirements relative to consent management.

Many data storage services (such as Amazon S3) or any purpose-built databases (such as Amazon RDS and DynamoDB) can be configured to support your specific policy and role-based access needs. Services such as the [AWS Glue Data Catalog](#) can be used to track the lineage of PHI stored in a data lake or as part of a [modern data architecture \(formerly Lake House\)](#), which a custom consent management solution can use.

Consent management solutions may also be prebuilt, packaged and sold on [AWS Marketplace](#), which provides a centralized catalog of solutions built on AWS by Partners and customers.

Designing a consent management service or platform requires an understanding of specific decisions about how data can be organized, annotated, and shared. The design implications of consent management need to be considered early in the design process. Adding controls for consent management becomes more expensive the later in the product development cycle they are considered, if possible to add or modify at all. According to the [Privacy Patterns collective](#) many of these concerns are [codified into patterns](#)<sup>3</sup> to help communicate them better.

- Purpose-based consent is where a data controller secures consent from an individual for data for a specific purpose only.
- Implicit consent is where a data controller secures consent implicitly due to an individual doing business with the data controller.
- Explicit consent is where data controllers can build mechanisms to gain explicit consent before data can be processed to a third party following the patterns [Outsourcing \[with consent\]](#), [Obtaining Explicit Consent](#), [Lawful Consent](#), and [Informed Consent for Web-based Transactions](#).
- Understanding and identifying how data is organized within the domains it originates from ([User data confinement pattern](#) and [Personal Data Store](#)).
  - This can mean that when data is consolidated, the organizational and domain structures are maintained to aid in the process of identifying the originating source of the data, or this can mean annotating the data to facilitate the preservation of data lineage across centralized or disparate databases.
  - If a source of data is no longer available, then having methods in place to enable and verify that consent can still be managed and that owners of that data are notified that copies of data exist.
- Having the ability to remove records associated with an individual, regardless of which system it resides in ([Obligation Management](#), [Right to erasure: Article 17 GDPR](#), [Right to be Forgotten: Recital 66 GDPR](#)).

- Data is partitioned relative to the source of where the data comes from, and this can be in separate databases or locked through encryption at the highest organizational level. This makes it reasonably easy to manage data keys; however, it is challenging to delete data cryptographically. Alternatively, creating a metadata store that tracks each record and stores keys relative to the individual record can be implemented, but this can also introduce performance costs.

Consent management is complex and, in many cases, the local regulations may dictate what granularity must be pursued. There are good architectural best practices that should be understood and followed regardless of whether you are endeavoring to build your own solution or buy one that can be deployed into the cloud. This section has provided some of the basic building blocks of what makes up a consent management solution to help you understand the complexities of the problem and to what extent AWS can unburden you from managing your infrastructure, including what tools and services AWS can provide for you to implement such a solution.

## Continuous compliance

When dealing with security and governance, there are three areas to consider:

1. As part of the overarching cloud governance, incorporate guidance from the frameworks in scope (there can be more than one framework):
  - Corporate guidelines—not necessarily tied to a specific vertical or line of business, but required across the whole organization structure—often tied to internal audits.
  - Requirements coming from laws and regulations (for example, GxP or HIPAA in the U.S. for PHI data).
  - Requirements to do business (for example, for credit card data, [PCI DSS](#) is required by the credit card industry).
  - Customers or partners (for example, even though not required by law, some large healthcare payors require adherence to [HITRUST](#)).
2. When the environment is being deployed and configured, have the right controls and configuration in place to align to the chosen frameworks. For example, encrypt in transit, encrypt at rest, enable logging, enable backups, and apply roles with least access necessary for the task.

Services such as AWS CloudTrail, AWS IAM, AWS Config, AWS Control Tower, and [AWS License Manager](#) help customers to implement operational controls over their cloud resources. Controls can be implemented as AWS Config rules. AWS Config comes with numerous [conformance packs](#) which provide common controls that align to many compliance frameworks out of the box.

For example, you can configure AWS Config to stream configuration changes and notifications to an [Amazon Simple Notification Service \(Amazon SNS\)](#) topic. These notifications can then be used to alert staff for manual investigation and remediation, or automatically trigger remediation of non-conformance. For example, if an Amazon S3 bucket is made public, then an [AWS Lambda](#) function can switch it back to private. This automatic remediation minimizes the amount of time your environment is out of compliance.

Amazon CloudWatch and [AWS Security Hub](#) can then provide customers with a comprehensive overview of operational health and security posture across their AWS accounts at all times.

3. In case of an audit or inspection, enable the collection of auditable artifacts to:
  - Verify and validate that said framework is being followed.
  - Demonstrate alerts, notifications, and automation to remediate or rectify any deviations.

Audit Manager helps customers assess and collect evidence about their controls. Following is a specific example to make these controls easier to relate to. Suppose a customer is providing an application for a hospital and that the required regulatory framework is HIPAA (step 1). By studying the [AWS HIPAA whitepaper](#), the customer can learn the recommendations on how to use [HIPAA-eligible AWS services](#) and how those services should be configured in alignment with compliance requirements (step 2). You can then enable AWS Config from the [AWS Management Console](#) and deploy the [HIPAA conformance pack](#), which will establish a baseline of the assets in the environment.

This baseline can be used to compare against the HIPAA recommendations in order to identify any deviations (step 2), and potentially allow for remediation as appropriate when combined with [Amazon EventBridge](#) for automation.

Finally, Audit Manager can be used to collect and compile data to create a report that shows not only compliance status against HIPAA at any given time, but also a history of compliance over time (step 3). Continuous compliance should be the goal instead of simply performing an assessment or audit one time a year.

## GxP considerations

Traditional approaches towards [GxP](#) compliance involved point-in-time compliance checks, for example, as part of configuration management where conformance checks were made prior to changes being deployed into production or during regulatory scheduled internal audits. Periodic internal compliance audits typically are in place to verify the state of the system is still in compliance based on an organization's specific standard operating procedures as part of quality management systems in place.

Change is constant, and changes that happen between point-in-time verifications may expose your environment to compliance risks. By continually monitoring your environment, you can enable a continuous state of compliance. 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 FDA Title 21 CFR Part 11](#).

As application changes are being deployed, traditional steps like installation qualification (IQ) can now be automated through the use of IaC. Your infrastructure is now defined through a template like a CloudFormation template. This is the input you need to compare against the output, your provisioned infrastructure. When a deployment is detected, it's now possible to check the input specification against the provisioned infrastructure and make a comparison. If things match, the IQ step passes, and you can move on to operational qualification (OQ), which again can be automated using third-party automated testing tools.

## Compliance at scale

The topic of scale takes on several meanings depending on the audience. Compliance at scale can cover the topics of how to manage compliance across an entire organization and scale across several projects, including managing multiple regulatory and compliance frameworks. Another perspective on scale is how to manage compliance for a project that utilizes dynamic compute and resource scaling to address varying system load driven by usage. This is important to allow enterprises to efficiently manage the costs of their deployed solutions, accommodating a growing user base over time.

## Enterprise-wide compliance

For GxP or other regulatory programs, compliance considerations across an organization typically require implementing controls managed by a central team that is

continually monitored and maintaining them to verify that current regulatory program needs are satisfied. This can take the form of enforcing a baseline of resources and guardrails to enable compliance across all teams. Examples include golden machine images with the required controls in place by default; centrally managed security, action, and access policies; and default networking rules.

AWS provides a variety of services and features, allowing customers to implement their enterprise-wide compliance strategies. [Amazon Machine Images \(AMI\)](#) allows teams to create and utilize their own custom systems images for use by their development teams, helping to ensure the minimum compliance requirements are set for all solutions. Control over who can do what with the AMIs (use, update, delete) can be controlled with roles and policies defined companywide using the IAM service to assign which users are allowed to assume which roles and what policies each of those roles have relative to the services and resources of the system.

The current state of a system can be monitored for changes that can impact its regulatory and compliance posture, and if the designers choose, automatic remediation and various notifications can be generated and acted upon. AWS Config provides default compliance packages (conformance packs) that can be customized to a solution's specific needs. Audit Manager helps with the task of continuously gathering audit evidence to access the systems controls required for reoccurring audit requests. AWS Organizations is an account management service that consolidates the management of multiple AWS accounts.

Utilizing the services mentioned above allows for the creation of a hierarchical account structure for managing compliance, security, and budgetary needs. Another service that can enable an organization-wide compliance strategy is [AWS Control Tower](#). AWS Control Tower extends the capabilities of AWS Organizations by orchestrating the creation of landing zones with defined best practices, setting up resources on the user's behalf with the defined policies and procedures. Together, these services provide customers with the ability to define and manage their regulatory and compliance controls across their AWS accounts and manage them efficiently.

In addition, identifying common approved tooling and resources is an additional step many organizations may implement to avoid duplication of effort across their internal teams and promote a solid regulatory stance across the organization. This may also include a common repository of approved suppliers along with the results of the supplier assessments conducted to verify the named suppliers, such as AWS for cloud services, meet the individual organizations' regulatory and compliance requirements. AWS provides third-party attestations to support supplier assessments. For more information

about AWS Artifact and how to access the AWS security and compliance documents, see the [Shared responsibility](#) section.

With a solid plan in place, organizations can realize the efficiencies and economies of scale an enterprise-wide compliance solution brings as they scale out their regulatory compliance offerings.

## Dynamic resource scaling

Scale, from the lens of resource elasticity in the cloud, can pose challenges for some compliance teams. Historically, a complete system often had to be validated with all resources pre-allocated, enabled, and verified as part of the IQ, OQ, and performance qualification phases of the testing process to provide the required evidence for regulatory compliance. This approach can be replicated in the cloud by pre-allocating the maximum number of resources a solution is specified to require during the installation phase.

Typically, this involves reserving instances or leaving dynamic instances running for the life of the solution. Although these configurations are possible, they don't use one of the key value propositions of cloud computing, elasticity, or the ability to create and remove resources based on need dynamically. This helps to manage the cost of the solution over time as resources are not allocated until the system usage metrics indicate more resources are required.

Supporting elasticity, or resource scaling both up-and-down and in-and-out, in a regulated solution, starts with a set of system requirements defining the scaling behavior of the environment, as opposed to only the amount of required resources. Defining the minimum resources allowed, the rules for scaling based on system performance and usage metrics, and the maximum the system will be allowed to grow are key elements to create a testing plan with evidence to illustrate the stability of the systems. Having requirements that specify the minimum resource that will be allocated allows the base system to be validated and to mapped directly back to the requirements for a minimum system configuration.

The requirements for the maximum resources allowed provide evidence that the system performance is maintained when the maximum specified system usage is reached. A system with no upper bounds might experience performance issues over time for several reasons (for example, database retrievals, or [task thrashing](#)), providing requirements and appropriate validation for the upper bounds of the system promotes reliable operation and performance.

Lastly, having requirements that specify when new resources are allocated and brought into the active resource pool, and defining when resources are released, allows the validation tests to provide evidence that the system dynamically reacts to system usage as defined in the requirements. For example, compute resources might be added when the current usage reaches 80% capacity and reduces when the capacity is below 40%. These bounds depend on the specific solution profile, for example, the speed at which utilization increases, balanced against the time required to provision the additional resources required.

Many AWS services perform the undifferentiated heavy lifting tasks associated with elasticity as part of the service's normal operations. For example, when using Amazon S3 or [Amazon Simple Queue Service \(Amazon SQS\)](#), these services scale up and down to meet the demand. Other services such as Amazon EC2, [Amazon Elastic Container Service \(Amazon ECS\)](#), or [AWS Fargate](#) integrate with the [AWS Auto Scaling](#) service to help manage the scaling-up of resources to manage the on-demand load or the releasing of resources when the load decreases. For a set of best practices and further discussion of these resources, see the [AWS Auto Scaling Documentation](#).

Utilizing these approaches, organizations can use the cost efficiency associated with elasticity in the cloud while maintaining a verifiable system that aligns with their regulatory and compliance requirements.

## GxP considerations

The ability to dynamically scale resources can increase the challenges of demonstrating that your solution's environment is in a state of control. Relatedly, the ability to demonstrate control of your solution's environment is at the core of many regulations.

One of the significant advantages of the cloud is the freedom to use the right tool for the job and your ability to provision the resources within minutes. However, from a regulatory perspective, this could be viewed as a risk. With such freedom, how is control maintained? One technique is allow-listing of services. Each AWS service goes through a review process before being approved for use as part of regulated workloads. Subject matter experts (SMEs) check each service, make sure it is included in the required AWS compliance programs, perform a risk analysis, and produce a template configuration approved for use. This template might be for a single service, a combination of services, or even a complete stack like an n-tier web application.

This is synonymous with the concept of building blocks from [Good Automated Manufacturing Practice \(GAMP\)](#) good practice guides. The control over which services are allowed to be used can be achieved through service control policies (SCPs), which

can be set at the organizational unit or account level in AWS Organizations. SCPs control what is allowed to be used in the account, even by the root user. As for the building block templates, these can be made available using [AWS Service Catalog](#). So developers retain the self-service benefits of cloud to support business agility but operate within a controlled environment, which can be demonstrated to regulators.

Historically, this need to demonstrate control was often interpreted as having a fixed set of IT assets, such as servers, which was captured statically in a configuration management database. However, one of the other benefits of the cloud is its scalability. As discussed in the [Dynamic resource scaling](#) section, this ability is testable and can be included in service or building block qualification.

## Appendix B: Compliance, certification, and regulatory alignment

Central to earning the public trust is ensuring that you comply with the set of regulations that is required within each country and region where you operate or store data. These government and industry-accepted standards and regulations are created to validate that public information is properly handled and protected. Governments or professional standards organizations often provide ways in which companies or private institutions can prove compliance through some form of certification to show that they are meeting or exceeding those requirements. The following sections highlight key laws, standards, and regulatory frameworks of various countries and regions as well as compliance certifications that may be applicable to organizations operating in the healthcare and life sciences industries.

### Compliance

Compliance is the ongoing process of meeting or exceeding the legal, ethical, and professional standards applicable to a particular organization. HCLS compliance requires HCLS organizations and providers to develop effective processes, policies, and procedures to define appropriate conduct, train the organization's staff, and monitor adherence to the processes, policies, and procedures. HCLS compliance covers numerous areas including, but not limited to, patient care, billing, reimbursement, managed care contracting, the [Occupational Safety and Health Administration \(OSHA\)](#), Joint Commission on Accreditation of Healthcare Organizations, [General Data Protection Regulation \(GDPR\)](#), and [Health Information Portability and Accountability Act \(HIPAA\)](#) privacy and security, and [Good x Practice \(GxP\)](#) such as Good Pharmaceutical

Practices or Good Documentation Practices (quality control best practices for the manufacture of pharmaceuticals and medical devices), to name a few.

Each of the governmental agencies that regulate healthcare approaches its regulatory framework based upon its own area of control. Examples of well-known laws and regulations required by various international governments include:

**Family Educational Rights and Privacy Act (FERPA)** – A U.S. federal law enacted to protect the privacy of student education records.

**General Data Protection Regulation (GDPR)** – A legal framework that sets guidelines for collecting and processing personal information from individuals who live in the European Union.

**Health Information Portability and Accountability Act (HIPAA)** – A U.S. federal law that requires the creation of national standards to protect sensitive patient health information from being disclosed without the patient's consent or knowledge.

**Health Information Technology for Economic and Clinical Health (HITECH)** – A U.S. federal law enacted as part of the American Recovery and Reinvestment Act of 2009, was signed into law on February 17, 2009, to promote the adoption and meaningful use of health information technology.

**Protection of Personal Information Act (POPIA)** – Regulates the collection, use, and processing of personal information in South Africa.

**My Number Act** – The Japanese Social Security and Tax Number System is a social infrastructure that improves administrative efficiency, enhances public convenience, and helps realize a fairer and more just society.

**Data Protection Act 1998 (UK DPA - 1988)** – A law designed to protect personal data stored on computers or in an organized paper filing system. It enacted the EU Data Protect Directive 1995's provisions on the protection, processing, and movement of data.

**Personal Information Protection and Electronic Documents Act (PIPEDA)** – A Canadian law that governs the collection, use, and disclosure of personal information in a manner that recognizes the right of privacy of individuals with respect to their personal information and the need for organizations to collect, use, or disclose personal information for purposes that a reasonable person would consider appropriate in the circumstances.

## Certification

Certifications provide a means by which organizations can attest that they meet certain requirements needed to participate in that domain. For example, [Health Information Trust Alliance \(HITRUST\)](#) certification enables vendors to provide a third-party attestation that the HITRUST controls have been met which can be used to support an organization's security and privacy goals.

Examples of well-known certifications and attestations include:

**[Cloud Computing Compliance Controls Catalog \(C5\)](#)** – A German government-backed attestation scheme introduced in Germany by the Federal Office for Information Security (BSI) to help organizations demonstrate operational security against common cyberattacks when using cloud services within the context of the German government's [Security Recommendations for Cloud Computing Providers](#).

**[Federal Risk and Authorization Management Program \(FedRAMP\)](#)** – A U.S. government-wide program that delivers a standard approach to the security assessment, authorization, and continuous monitoring for cloud products and services.

**[Federal Information Processing Standard \(FIPS\) Publication 140-2](#)** – A U.S. and Canadian government standard that specifies the security requirements for cryptographic modules that protect sensitive information.

**[Health Information Trust Alliance Common Security Framework \(HITRUST CSF\)](#)** – A framework that uses nationally and internationally accepted standards and regulations such as GDPR, ISO, National Institute of Standards and Technology (NIST), Payment Card Industry Data Security Standard (PCI DSS), and HIPAA, to create a comprehensive set of baseline security and privacy controls.

**[ISO 9001:2015](#)** – International Organization for Standardization (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.

**[ISO/IEC 27001:2013](#)** – International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) Joint Technical Committee security management standard that specifies security management best practices and comprehensive security controls following the ISO/IEC 27002 best practice guidance.

**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.

**ISO/IEC 27018:2019** – A code of practice that focuses on protecting personal data in the cloud. It is based and provides implementation guidance on ISO/IEC 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/IEC 27002 control set.

**Payment Card Industry Data Security Standard (PCI DSS Level 1)** – A proprietary information security standard administered by the PCI Security Standards Council, which was founded by American Express, Discover Financial Services, JCB International, MasterCard Worldwide, and Visa Inc.

**System and Organization Controls Reports (SOC 1, SOC 2, SOC 3)** – Independent third-party examination reports demonstrating how an organization achieves essential compliance controls and objectives. The purpose of these reports is to help you and your auditors understand controls established to support operations and compliance.

**Hébergeur de Données de Santé (HDS)** – Introduced by the French governmental agency for health, Agence du Numérique en Santé (ANS), the HDS certification aims to strengthen the security and protection of personal health data. Website is in French.

**NHS Data Security and Protection Toolkit** – An online self-assessment tool that allows organizations to measure their performance against the National Data Guardian's 10 data security standards.

## Regulatory alignment

Regulation plays a significant role in the healthcare industry and health insurance coverage. The various regulatory bodies protect the public from many health risks by establishing rules and requirements, and oversee compliance and provide numerous public health and welfare programs. Together, these regulatory agencies protect and regulate public health at every level. Healthcare regulations are developed and implemented by all levels of government (federal, state, and local) and by private organizations.

Healthcare regulations and standards are necessary to verify compliance and provide safe medical care to every individual who accesses the system. In turn, the healthcare regulatory agencies monitor practitioners and facilities, provide information about industry changes, promote safety, and provide legal compliance and quality services.

In life sciences, the manufacture of pharmaceuticals and medical devices must be quality controlled by regulations such as [FDA Title 21 CFR Part 11](#), or the suite of regulations related to GxP.

Regulatory agencies often establish rules and regulations for the HCLS industries, and their oversight is mandatory. Some other agencies, such as those agencies for accreditation, require voluntary participation but are still important because they provide rankings or certification of quality and serve as additional oversight, ensuring that healthcare organizations promote and provide quality care.

Examples of well-known regulatory frameworks include:

**[Federal Information Security Management Act of 2002 \(FISMA\)](#)** – The act requires each federal agency to develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency, including security and systems provided or managed by another agency, contractor, or other sources.

**[G-Cloud](#)** – The G-Cloud Framework enables public bodies to procure commodity-based, pay-as-you-go cloud services on government-approved, short-term contracts through an online catalog called the Digital Marketplace. This streamlined procurement process supports the UK government's *Cloud First* policy, and is a crucial component in the government's ambition to operate a *cloud-native* digital architecture.

**[GxP](#)** –GxP is an acronym for *Good x Practice* and refers to the regulations and guidelines applicable to life sciences organizations that make food and medical products such as drugs, medical devices, and medical software applications. The overall intent of GxP requirements is to verify that food and medical products are safe for consumers and to validate the integrity of data used to make product-related safety decisions.

**Medicaid Information Technology Architecture Framework (MITA 3.0)** – The latest major release of MITA, a Center for Medicare and Medicaid Services initiative that fosters an integrated business and IT transformation across the Medicaid enterprise to improve the administration and operation of the Medicaid program.

**Public Health Regulations (PHR)** – Verified regulations with statistically higher individual noncompliance rates in establishments three months before a microbiological positive or a public health-related enforcement action than in establishments with no positives or enforcement actions. This statistical association does not inherently imply that a particular regulation constitutes a more severe food safety concern, but gives a statistical association to align scheduling criteria and agency resources better.

**Cloud Security Guidance published by NCSC** – This whitepaper provides a list of 14 essential cloud security principles to consider when evaluating cloud services and why these principles may be important to the public sector organization. Cloud service users should decide which of the principles are important, and how much (if any) assurance users require in implementing these principles.

**Israeli Ministry of Health Cloud Computing Circular** – Israeli Ministry of Health issued this circular to establish criteria for the proper operation of computing applications using cloud computing by healthcare organizations to encourage the introduction of advanced technologies for use by healthcare organizations.

## Document revisions

Date	Description
April 13, 2022	First publication

---

## Notes

---

- <sup>1</sup> Naveed, M., Ayday, E., Clayton, E. W., Fellay, J., Gunter, C. A., Hubaux, J.-P., . . . Wang, X. (2015). Privacy in the genomic era. *ACM Computing surveys (CSUR)* 48.1, 1-44.
- <sup>2</sup> U.S. Department of Health & Human Services, Office for Civil Rights. (2013, July 26). *What is the difference between "consent" and "authorization" under HIPAA Privacy Rule?* Retrieved from HHS.gov: <https://www.hhs.gov/hipaa/for-professionals/faq/264/what-is-the-difference-between-consent-and-authorization/index.html>
- <sup>3</sup> Doty, N., Gupta, M., Caiza, J. C., del Álamo, J. M., Martin, Y. S., Drozd, O., . . . Colesky, M. (2022). *Privacy Patterns*. Retrieved from Privacy Patterns, UC Berkley School of Information: [privacypatterns.org](http://privacypatterns.org)