

AWS Well-Architected Framework

Healthcare Industry Lens



Healthcare Industry Lens: AWS Well-Architected Framework

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Abstract and introduction	i
Introduction	1
Custom lens availability	1
Definitions	3
Learning path for healthcare	4
General design principles	6
Pillars of the Well-Architected Framework	7
Operational excellence	7
Design principles	7
Best practices	8
Key AWS services	16
Resources	17
Security	18
Design principles	18
Best practices	19
Key AWS services	30
Resources	32
Reliability	32
Design principles	32
Best practices	33
Key AWS services	35
Resources	36
Performance efficiency	36
Design principles	36
Best practices	37
Key AWS services	40
Resources	40
Cost optimization	41
Best practices	41
Key AWS services	43
Resources	44
Sustainability	44
Best practices	45
Key AWS services	48

Resources	48
Scenarios	49
Electronic healthcare record and revenue cycle systems	49
Questions	50
Healthcare interoperability	51
Reference architecture	52
Questions	55
Medical imaging	57
Reference architecture	58
Questions	60
Healthcare analytics	62
Reference architecture	63
Questions	65
Machine learning for healthcare	67
Machine learning resources	69
Reference architecture	69
Questions	71
Conclusion	73
Contributors	74
Document revisions	75
Notices	76
AWS Glossary	77

Healthcare Industry Lens

Publication date: **November 17, 2022** ([Document revisions](#))

This paper describes the Healthcare Lens for the AWS Well-Architected Framework, which enables customers to review and improve their cloud-based architectures and better understand the business impact of design decisions. We present general design principles and specific best practices aligned to the six pillars of the Well-Architected Framework.

Introduction

The [AWS Well-Architected Framework](#) helps you understand the pros and cons of decisions you make while building and deploying systems on AWS. By using the Framework you will learn architectural best practices for designing and operating reliable, secure, efficient, sustainable, and cost-effective systems in the cloud. It provides a way to consistently measure your architectures against best practices and identify areas for improvement. We believe that having Well-Architected systems greatly increases the likelihood of business success and improves the patient experience.

In this Lens we focus on how to design, deploy, and manage your healthcare workloads in the AWS Cloud. This Lens augments the Well-Architected Framework and highlights additional key principles that should be considered for healthcare workloads. See [Learning path for healthcare](#) for a series of foundational resources that can prepare you to use this document. We recommend that you read the preliminary resources listed in the learning path, like the AWS Well-Architected Framework whitepaper, before reading this Lens. Consider all the principles and questions in the AWS Well-Architected Framework whitepaper and this Lens when designing your architecture.

This document is intended for those in technology roles, such as chief technology officers (CTOs), chief information security officers (CISOs), architects, developers, compliance officers, and operations team members. After reading this document, you will understand AWS best practices and strategies to use when designing architectures for the development and deployment of healthcare applications.

Custom lens availability

Custom lenses extend the best practice guidance provided by AWS Well-Architected Tool. AWS WA Tool allows you to create your own [custom lenses](#), or to use lenses created by others that have been shared with you.

To determine if a custom lens is available for the lens described in this whitepaper, reach out to your Technical Account Manager (TAM), Solutions Architect (SA), or AWS Support.

Definitions

- **Health data/sensitive data:** Broadly defined as information relating to an identified or identifiable person. Example data elements include identification numbers, location data, genetic information, cultural and social attributes, and identifiable health records. Specific regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) or the General Data Protection Regulation (GDPR), may impose definitions and requirements. For example, HIPAA includes a definition of Protected Health Information (PHI), which defines identifiable data that is covered under the regulatory standard.
- **Healthcare payor:** Organizations who develop healthcare policy, manage risk, or provide healthcare networks, payment, and adjudication services for citizens, employers, or private individuals (for example, health plans, intermediaries, claims processing entities, and policy or regulatory entities at the state and federal level).
- **Healthcare provider:** Organizations (such as clinics, hospitals, and care networks) providing acute, ambulatory, ancillary, and retail healthcare services.
- **Healthcare ISV:** Technology providers who develop, maintain, and market technology solutions addressing the needs of healthcare organizations (for example, payors and providers).
- **Consumer health and wellness:** Technology provider companies who develop, maintain, and market health and wellness solutions targeting consumers.
- **Standards setting organizations:** Organizations responsible for establishing industry standards that are common across healthcare. Examples include Health Level 7 (HL7) for healthcare interoperability and the Health Information Trust (HITRUST) Alliance for data protection.
- **Regulatory bodies:** Organizations, often geography specific (FDA, ONC, EMA), that define regulations for controls that healthcare organizations must adopt in order to operate within that geography.

Learning path for healthcare

This Lens represents one of a body of resources that you can use to design and operate Well-Architected cloud-based healthcare workloads. This Lens augments, but does not repeat, all of the guidance offered in the foundational materials. The following learning path for healthcare lists an ordered collection of guides, documentation, blog posts, and other content to help customers be successful in bringing a Well-Architected solution to the cloud. We recommend considering each of resources below, and depending on your prior cloud experience, reviewing the prerequisites before reading this Lens.

Getting started

- [Introduction to AWS](#)
- [Overview of Amazon Web Services](#)

AWS fundamentals

- [Shared Responsibility Model](#)
- [Amazon Web Services: Risk and Compliance](#)
- [Management and Governance on AWS](#)
- [Management and Governance Cloud Environment Guide](#)
- [Introduction to AWS Security](#)
- [AWS Compliance Programs](#)

Architecting for healthcare

- [AWS Well-Architected](#)
- [AWS Well-Architected Lenses](#)
- [AWS Compliance Programs](#)
- [AWS Quick Starts](#)

Additional useful links

- [Healthcare Compliance in the Cloud](#)

- [How to protect sensitive data for its entire lifecycle in AWS](#)

General design principles

The Well-Architected Framework identifies a set of general design principles to facilitate good design in the cloud. In addition, the following design principles should also be considered for designing and operating healthcare workloads:

- **Align with applicable regulatory and quality frameworks:** Identify relevant regulations early, and architect solutions from the start to meet regulatory requirements.
- **Automation reduces operational risk:** Modern software practices, such as continuous integration and continuous delivery, allow for automated checks, like aligning to specific controls frameworks.
- **Encrypt all sensitive data:** Protecting data with encryption, at rest and in transit, is a best practice of the Well-Architected Framework. Further, many regulatory frameworks applicable to healthcare workloads specifically call out encryption of health data, and it is required by the AWS Business Associate Addendum. Implement encryption of all health data in your environments.
- **Log everything:** Logging allows you to monitor system and data access, and to verify that only authorized individuals are accessing the appropriate data. Implement immutability for logs for long term retention.
- **Implement least privilege for all data, not just health data:** Granting access to only the systems and data required for someone, or something, to do a job is a best practice in the Well-Architected Framework. Similar to the encryption design principle above, healthcare regulatory frameworks may require enforcing restrictions on access to health data. Restrict access to production systems and health data to only those who need it. Implement reviews to maintain least privilege over time.
- **Adopt modern software communication protocols:** Healthcare has many standards, some of which do not embrace modern software practices, such as APIs. Where possible, use data standards and communication protocols in-line with best practices to align with both current standards and potential future standards.
- **Promote interoperability:** Unlock new product development opportunities and improve patient outcomes with architectures that facilitate secure, governed access to health data across silos.
- **Plan to recover from failures automatically:** Healthcare workloads enable the delivery of care to patients. Consequently, failures may negatively impact patients. Identify critical workloads and the key performance indicators (KPIs) that describe workload health. Design architectures with monitoring and automated recovery processes to ensure that systems meet availability requirements.

Pillars of the Well-Architected Framework

This section maps the six pillars of the Well-Architected Framework to healthcare workloads. Each pillar discusses design principles, definitions, best practices, evaluation questions, considerations, key AWS services, and useful links.

Pillars

- [Operational excellence pillar](#)
- [Security pillar](#)
- [Reliability pillar](#)
- [Performance efficiency pillar](#)
- [Cost optimization pillar](#)
- [Sustainability pillar](#)

Operational excellence pillar

The operational excellence pillar provides guidance on running and monitoring systems to deliver business value and continually improving supporting processes and procedures.

The operational excellence pillar provides an overview of design principles, best practices, and questions. You can find prescriptive guidance on implementation in the [Operational Excellence Pillar](#) whitepaper.

Design principles

There are a number of principles that drive operational excellence in the cloud. Within healthcare, compliance is often a key consideration:

- **Develop a hub-and-spoke model for compliance controls:** Many healthcare customers are subject to regulatory requirements. Mapping to a central control framework, such as the National Institute of Standards and Technology (NIST) or HITRUST, simplifies mapping to multiple regulations in a hub-and-spoke model.
- **Align software and infrastructure development with applicable quality frameworks:** Align development processes and tools with the quality frameworks that apply to your workload, such as ISO 13485 and ISO 14971.

- **Take advantage of AWS fully managed services and approved third-party solutions:** Leverage managed cloud services and third-party solutions approved by your business criteria to simplify meeting regulatory requirements and maintaining a strong security posture. Managed cloud services also simplify the operations of managing your workloads.

Best practices

Definition

Operations teams must understand their business and customer needs so they can support business outcomes. Operations teams create and use procedures to respond to operational events, and they validate the procedures' effectiveness in supporting business needs. Operations teams collect metrics that are used to measure the achievement of desired business outcomes. In healthcare, operations also often include compliance. You may be subject to various healthcare regulations depending on what you do and where you operate. Plan for continual change in your business context, business priorities, and customer needs. It's important to design operations to support evolution over time in response to change, to incorporate lessons learned through their performance, and continuously demonstrate you are operating within the relevant regulations and frameworks.

The following are best practice areas for operational excellence in the cloud:

- [Organization](#)
- [Prepare](#)
- [Operate](#)
- [Evolve](#)

Organization

There are no operational excellence best practices for Organization specific to the Healthcare Industry Lens.

Prepare

HCL_OPS1. Have you defined a formal risk management program?

Create and document a risk management program

Many regulatory frameworks are intended to reduce risk in one way or another. Organizations usually understand that they must reduce their risk, but may struggle to determine what the appropriate risk appetite is and how to manage it. This is accomplished using a documented risk management program.

In healthcare, the risk management program is designed to safeguard patient data, as well as the overall organization's assets and reputation. For example, a healthcare provider's risk management program also covers clinical quality, which is critical to reducing potential patient risk. Healthcare organizations should create a comprehensive risk management program that includes all operational, clinical, strategic, financial, legal, environmental, and any other potential risk domains.

When designing your risk management program, ask questions similar to the following:

- Have you defined risk and compliance roles for the cloud?
- Have you created a risk management program for the cloud?
- Have you assessed your workload against regulatory needs?
- Have you performed a security risk assessment?
- Have you created a cloud governance program?
- Have you created a responsibility model?

Create a risk authority team

Creating an effective risk management program for the cloud should be defined by the appropriate risk authority team. The risk authority within the organization (for example, board of directors, chief risk officers, or business risk officers) must evaluate the criticality of a business process (and the underlying workloads that support that process) and specify the level of availability they require for the process. Consider the potential impact a disruption may have on the process, organization, and customers. Weigh the impact against the cost of operating the workload in a high availability mode, consequences for business agility, and pace of innovation. Working backwards from established risk appetites allows you to define operational priorities and corresponding cloud architectures that can meet your business objectives.

AWS publishes the [Amazon Web Services: Risk and Compliance whitepaper](#) that outlines the mechanisms used to manage risk on the AWS side of the shared responsibility model. This

whitepaper also provides tools that customers can use to ensure these mechanisms are being implemented effectively.

HCL_OPS2. What policies and procedures has your organization adopted for cloud governance?

Create policies and procedures to govern cloud workloads

Cloud governance is a set of policies and procedures that outline, or govern, how an organization manages their cloud workloads. A mature governance program requires understanding the compliance objectives and requirements and establishing a control environment that meets those objectives and requirements. Organizations that host and process healthcare data can be required to meet specific standards and regulations, such as HIPAA or General Data Protection Regulation (GDPR). A mature governance program can help verify that the necessary controls are implemented.

As outlined in the [Amazon Web Services: Risk and Compliance whitepaper](#), AWS customers are responsible for maintaining adequate governance over their entire IT control environment, regardless of how or where IT is deployed. Recommended practices include:

- Understanding the required compliance objectives and requirements (from relevant sources)
- Establishing a control environment that meets those objectives and requirements
- Understanding the validation required based on the organization's risk tolerance and applicable regulatory requirements
- Verifying the operating effectiveness of their control environment

Cloud deployments give organizations different options to apply various types of controls and various verification methods.

Strong customer compliance and governance on AWS should include the following:

- Reviewing the [AWS shared responsibility model](#), [AWS security documentation](#), [AWS compliance reports](#), and other information available from AWS, together with other customer-specific documentation. Try to understand as much of the entire IT environment as possible, and document all compliance requirements into a comprehensive cloud control framework.

- Designing and implementing control objectives to meet the enterprise compliance requirements as laid out in the [AWS shared responsibility model](#).
- Identifying and documenting controls owned by outside parties.
- Verifying that all control objectives are met and all key controls are designed and working.

Approaching compliance governance this way helps you better understanding you control environment. It can also delineate the verification activities that must be performed.

HCL_OPS3. How do you map security controls to compliance requirements?

Determine regulatory frameworks and security controls that are applicable to your business and your cloud workload

Organizations that host and process health data must verify that they are adhering to all applicable regulatory frameworks and standards. As healthcare organizations evolve and grow, they may either want, or be required, to adhere to multiple regulations or certifications. For example, a European organization may be required to meet GDPR and additional country-specific regulations in each country it operates in.

Map applicable frameworks and controls to AWS controls to align with regulatory frameworks

There are two common approaches to addressing multiple compliance regimes. First, organizations may choose to address each set of requirements from the beginning and develop mappings unique to each. Alternatively, organizations can choose to map to a common security framework, and leverage published controls mappings from that framework to many others in a hub-and-spoke model. AWS recommends the latter approach where possible to avoid duplicating effort.

As an example, here are steps you might take if you use NIST 800-53 as your security framework, and apply it to the HIPAA Security Rule on AWS:

1. Map NIST 800-53 to applicability within the AWS environment, considering the shared responsibility model with AWS and any third parties you may work with.
2. Use prebuilt AWS compliance checks for NIST or other frameworks with AWS config conformance packs, as well as implement any additional custom checks to monitor your AWS environments. Implement immutable logging to archive compliance posture over time.

3. Use NIST Special Publication 800-66 to map controls from NIST 800-53 to the HIPAA Security Rule

In other words, create a crosswalk to map your AWS controls to a common security control framework. Use this crosswalk to connect controls in your cloud environment to the regulation standards as required.

Another example is to create a [responsibility assignment matrix](#) (RACI) that designates roles who are responsible, accountable, consulted, and informed for regulatory controls. A policy with a RACI matrix clarifies accountability and helps affirm that proper actions are taken by designated owners. A consulted party offers guidance related to the control. Finally, the informed party is made aware of the situation and any actions taken. Using RACI matrices can help organizations properly implement plans and procedures when dealing with regulatory controls.

HCL_OPS4. How do you educate employees on access to sensitive data?

Ensure employees who may have access to sensitive healthcare data are trained on the rules and regulations

Organizations that host or process PHI should ensure that employees who have access to healthcare data, either intentionally or accidentally as part of their job function, are trained on the rules and regulations that govern the organization. Employees should have knowledge on what to do when viewing sensitive data. They should know how and where to host or process that data, and how to protect it. Train employees on any other regulation-based requirements, such as breach disclosure. Document all of this in your risk management program.

Create and document a policy and procedure aligned to each control and safeguard

Organizations that are hosting and processing sensitive healthcare data should have a documented policy that aligns with each control or safeguard in place to secure the data. In addition, each policy should have an associated procedure document that outlines how the policy will be implemented. These policy and procedure documents will help educate employees on the safeguards used, and can help demonstrate your compliance posture to your stakeholders. These documents help create a stronger culture of compliance for your organization.

Operate

HCL_OPS5. How do you demonstrate continuous compliance?

Partition workloads involving sensitive data into separate environments

Minimize access to sensitive data by isolating workloads to separate environments requiring additional controls for access. Segmenting can be done by AWS accounts, VPCs, or Amazon Simple Storage Service buckets. Minimize using sensitive data in non-production environments.

Architect and build with the ability to generate evidence that demonstrate continuous compliance

Healthcare organizations must be able to demonstrate their compliance posture. Evidence that includes the safeguards used to protect sensitive healthcare data, as well as the documented policies and procedures, can all be used to demonstrate compliance. The cloud services used to architect a compliant foundation in the cloud, can also be used to gather the necessary evidence to demonstrate compliance posture. For example, using infrastructure as code, coupled with a software development lifecycle, can demonstrate a mature change management process, which is an important compliance control. Being able to demonstrate the full scope of a compliance posture is critical for all stakeholders, whether that is an organizations leadership, shareholders, customers, and patients.

There are several key concepts to consider when building out a continuous compliance posture. While AWS cannot assure compliance for your environment per the shared responsibility model, the following approach will make it easier for your organization to demonstrate compliance on AWS. In general, use managed services from AWS or third-party solutions, such as those available in AWS Marketplace, to simplify your approach.

Identify resources in the cloud environment

An accurate representation of your cloud environments is necessary to demonstrate continuous compliance. Understand what AWS resources exist and how they interact with each other. AWS Config will help you identify these resources and how they are configured. Use distributed tracing solutions, such as AWS X-Ray, to understand how components of your system interact, and to map network accessibility between different resources in your environment.

Restrict resources and applications to pre-defined configurations

Coupling AWS Config with infrastructure as code will allow you to test application configurations before they are deployed in your environment. Apply governance to your AWS deployments using infrastructure as code tools like AWS CloudFormation, AWS Cloud Development Kit (AWS CDK), Terraform, and Service Catalog. Verify that all configurations are *secure-by-default* with best practices around encryption, logging, and least privilege.

Implement compliance-as-code for configuration

For each configuration you specify, test the controls you put in place. Use AWS Config as the central location to evaluate configuration changes. Where possible, use AWS Config managed rules, but also implement custom evaluations with AWS Lambda, fully capturing environment configuration. Configuration triggers will also shorten the time to identify AWS resources that are out of compliance compared to periodic triggers. This helps you demonstrate your compliance posture by automatically building and maintaining a list of resources within your AWS environment. It also allows you to continuously evaluate your compliance posture against the technical controls identified by your organization. For example, you can create an AWS Config rule that marks an Amazon S3 bucket as non-compliant if server-side encryption is not enabled or the Amazon S3 bucket policy allows unencrypted uploads. AWS provides [sample rules bundled into conformance packs](#) that align with many common regulatory frameworks and best practices, allowing you to start creating a compliance monitoring solution.

Centralize security and compliance findings

Many customers will use multiple AWS accounts (such as development, test, and prod, or department-specific accounts). Configuration management, while important, is not the only set of technical controls you may require. For example, you may combine your configuration posture with additional findings, from third-party solutions or AWS security services like Amazon GuardDuty. Technical controls and findings should be grouped together as evidence using a solution such as AWS Security Hub.

Map technical controls to compliance requirements using automation.

Simplify maintaining a complete view of your compliance posture by automatically mapping controls and findings to your internal policies. For example, if you have a compliance policy around encryption at-rest, you may have individual controls on the configuration of each AWS resource to verify encryption is enabled.

AWS Audit Manager helps automate evidence collection from a variety of sources within AWS, including Security Hub and Config. Bundling multiple pieces of evidence together under a single

policy makes it easier to demonstrate compliance with a specific framework or regulation. You can use Audit Manager's prebuilt frameworks, and you can manually specify a list of controls and policies that are important to your organization.

Use up-to-date artifacts.

The creation of artifacts that document the compliance posture of a cloud environment should be automated. Use services such as AWS Config, AWS Audit Manager, and AWS Security Hub to automatically collect and report the compliance state of a cloud environment.

HCL_OPS6. How do you automate remediation of compliance violations?

There are several key concepts to consider when creating an automated remediation solution. While your organization is responsible for compliance for your environment, per the shared responsibility model, the following approach will make it easier to demonstrate compliance on AWS. In general, use managed services either from AWS or a third-party solution, such as one available in AWS Marketplace, to simplify your approach. Similar to the recommendations for demonstrating continuous compliance, define compliance requirements and create associated policies and procedures for remediation before creating the remediation solution.

Automate remediation actions for non-compliant resources

Automate remediation of configurations that are out of compliance with your technical controls for rapid, consistent application of your policies. Event-driven architectures improve remediation times. Not everything can be predicted ahead of time. Certain remediations may be manual at first, but investigated when they occur and automated when possible in future occurrences.

In developing automated remediations, there are several steps you can follow:

1. **Specify controls:** Define the evidence and configuration you want to track. Use AWS Config and Security Hub to identify and surface findings.
2. **Identify when configuration changes happen:** Use AWS services that support event-driven architectures for identification. For example, AWS Config can monitor resource changes and Amazon EventBridge can serve as an event bus for additional resource changes.
3. **Implement remediation:** Services such as AWS Lambda and AWS Systems Manager can implement configuration changes.

4. **Rerun evaluation:** Verify remediation was implemented and the environment is back in compliance.

For example, you can create an AWS Config rule that marks an Amazon S3 bucket as non-compliant if the server-side encryption is not enabled. That rule can invoke a corresponding remediation Lambda function that configures server-side encryption on the bucket, bringing the bucket to a compliant state. For more information, refer to [Remediating Noncompliant AWS Resources by AWS Config Rules](#). AWS also provides sample AWS Config rules with remediation actions for [Amazon DynamoDB](#) and [Amazon S3](#).

Evolve

There are no operational excellence best practices for Evolve specific to the Healthcare Industry Lens.

Key AWS services

[AWS CloudFormation](#), [Amazon CloudWatch](#), and [AWS Config](#) comprise three services that can drive operational excellence. AWS CloudFormation can be used to create infrastructure templates based on best practices, and provision resources in an orderly and repeatable fashion. Amazon CloudWatch can be used for monitoring metrics, collecting logs, generating alerts, and triggering responses. AWS Config can assess, audit, and evaluate the configurations within AWS environments and trigger remediations when necessary. Following from the [Shared Security Model](#), it is critically important to properly configure cloud services. Configuration guidance can be found in the [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) whitepaper.

Other services and features that support the three areas of operational excellence are as follows:

Prepare:

- [AWS CDK](#)
- [AWS Control Tower](#)
- [AWS Organizations](#)

Operate:

- [Amazon GuardDuty](#)
- [Amazon Inspector](#)

- [Service Catalog](#)
- [AWS CloudTrail](#)
- [AWS Systems Manager](#)
- [AWS X-Ray](#)
- [AWS Security Hub](#)
- [AWS Audit Manager](#)

Evolve:

- [Amazon CodeGuru](#)
- [AWS Network Firewall](#)
- [Amazon Macie](#)
- [Amazon Detective](#)
- [AWS Elastic Disaster Recovery](#)

Resources

Refer to the following resources to learn more about our best practices related to operational excellence.

Videos

- [AWS re:Invent 2021 - Cloud compliance, assurance, and auditing](#)
- [Enforce compliance with AWS Config](#)

Documentation and blogs

- [Amazon Web Services: Compliance Resources](#)
- [Healthcare compliance in the cloud](#)
- [Remediate noncompliant AWS Config rules with AWS Systems Manager automation runbooks](#)

Whitepapers

- Amazon Web Services: [Risk and Compliance](#)

Security pillar

The security pillar provides AWS recommendations that help you meet your business and regulatory requirements. The practices in this pillar inform architectures that protect health data, control access, and respond automatically to security events.

Healthcare is a highly regulated industry in most geographies. This is in part because health data contains sensitive information about individuals' medical history, health behaviors, socioeconomic status, and financial information. Healthcare entities have a responsibility to keep health data safe and highly available for appropriate use. This responsibility of healthcare entities is formalized by governmental regulations across the globe that stipulate how health data must be protected. Therefore, it is important that a Well-Architected security and compliance program is used by healthcare entities when they move workloads to the cloud. A Well-Architected security and compliance program helps organizations design and operate workloads while meeting the necessary regulations, aligning with the proper frameworks, and keeping health data safe.

Design principles

The security pillar of the AWS Well-Architected Framework sets out principles that can help strengthen the security of your workload:

- **Implement a strong identity foundation:** Implementing the principle of least privilege is foundational to the security and compliance of healthcare workloads. Centralize identity management, and aim to avoid reliance on long-term static credentials.
- **Enable traceability:** Monitor, alert, and audit actions and changes to your environment in real time. Integrate log and metric collection with systems to investigate and remediate issues automatically.
- **Apply security at all layers:** Apply a defense in depth approach with multiple security controls. Security should apply to all layers, from the edge of the network to the application and code.
- **Automate security best practices:** Automated software-based security mechanisms improve your ability to scale more securely, rapidly, and cost-effectively.
- **Encrypt data in transit and at rest:** Classify your data to identify health data and other sensitive data. Use encryption, tokenization, and de-identification to decrease the sensitivity of data, and implement access controls.
- **Keep people away from data:** Use mechanisms and tools to reduce the need for direct access or manual processing of health data, consistent with the principle of least privilege.

- **Prepare for security events:** Prepare for an incident by having incident management and investigation policy and processes that align to your organizational requirements and applicable regulatory frameworks.

Best practices

Definition

Before you architect any system, you must put in place practices that integrate security as a foundation, and control who can do what. You also want to identify security incidents, protect your systems, and maintain the confidentiality and integrity of data through data protection. You should have well-defined and practiced processes for responding to security incidents. These tools and techniques are important because they support objectives such as protecting health data and complying with regulatory obligations.

The AWS shared responsibility model enables organizations that adopt the cloud to achieve their security and compliance goals. Security and compliance is a shared responsibility between you and AWS. This shared model can help relieve your operational burden as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of service facilities. Your organization assumes responsibility and management of the guest operating system (including updates and security patches) and other associated application software, as well as the configuration of the AWS provided security group firewall. Carefully consider the cloud services you adopt, as your responsibilities vary depending on the services used, and applicable laws and regulations. This differentiation of responsibility is commonly described as AWS assumes security of the cloud, and your organization assumes security in the cloud.

The following are best practice areas for security in the cloud:

- [Identity and access management](#)
- [Detective controls](#)
- [Infrastructure protection](#)
- [Data protection](#)
- [Incident response](#)

Identity and access management

HCL_SEC1. How do you identify where health data is in your environment?

Determine applicable regulatory frameworks and controls as it pertains to data classification

It is critical that organizations understand what types of data are being hosted and processed, and where that data resides. This understanding is a basis for ensuring that the right controls are in place for aligning with relevant regulatory frameworks and standards. Data classification also aids in traceability and access monitoring of sensitive data.

You can start by creating policies and procedures that align to the relevant regulatory frameworks. The policies and procedures should outline a data classification strategy that fits your business and regulatory requirements.

Create and document a data classification strategy

Based on the business requirements, and any applicable regulatory frameworks, implement a data classification policy. This policy should extend beyond simply marking health data, but should include other sensitive or confidential data, as well as public data. The [Data Classification: Secure Cloud Adoption](#) whitepaper provides examples of how to categorize data, and how to implement a data classification strategy that implements the appropriate controls based on the data category.

Make sure that health data is classified in accordance with the proper regulatory frameworks that your business aligns to.

Select the appropriate cloud deployment model according to your specific needs, the type of data you handle, and the assessed risk

As outlined in the [Data Classification: Secure Cloud Adoption](#) whitepaper, select the appropriate cloud deployment model according to your specific needs, the type of data you handle, and the assessed risk. Depending on the classification of the data, apply the relevant security controls (such as encryption) within your cloud environment. AWS also recommends that health data be classified and labeled as such, simplifying audits and ensuring that the proper technical controls can be implemented.

If your environment uses multiple AWS accounts, designate specific accounts to host and process health data to simplify managing where health data is located. For example, if your account structure mirrors your software development lifecycle with accounts designated for development,

testing, staging, and production, the production and staging accounts may be designated as “health data” accounts and are therefore documented as containing health data. Then, implement procedures and controls in the development and testing accounts to prevent health data from being stored there.

You can also assign [tags](#) to your AWS resources, which consist of a user-defined key and value. Tags help you manage, identify, organize, search for, and filter resources. Create tags to categorize resources by purpose, owner, environment, or other criteria. Use tags to help identify and document resources and objects that contain sensitive health data in accordance you’re your data classification strategy. Do not store sensitive health data in tags, as they are not intended to be used for private or sensitive data. Finally, access to resources can be [controlled through tags](#).

Implement automated data classification

[Amazon Macie](#) is a fully managed data security service that can help you identify sensitive data residing in Amazon S3. Macie automates the discovery of sensitive data, such as personally identifiable information (PII), to provide you with a better understanding of the data that your organization stores in Amazon S3. Macie also provides you with an inventory of your Amazon S3 buckets, and it automatically evaluates and monitors those buckets for security and access control.

You can use Amazon Comprehend (PII) and Amazon Comprehend Medical (PHI) to evaluate unstructured text data in your environment. Amazon Comprehend will provide a confidence score to measure the confidence that the data contains PHI as defined by HIPAA. This score can help you determine the sensitivity of the data reviewed.

HCL_SEC2. How are you implementing least privilege access to health data?

The ability to access health data should be limited to the people or systems who require the access to perform specific tasks. This covers access to the data itself, and access the systems that host health data.

Use identity and access management to control access to systems, resources, and data

Use AWS Identity and Access Management to control access to AWS services and resources. Use IAM to control who is authenticated to the environment and who is authorized to use services and resources. As outlined in the [IAM grant least privilege](#) documentation, start with a minimum set of permissions, and grant additional permissions as necessary. This approach exposes you to less risk than starting with permissions that are too lenient and then trying to tighten them later.

Health data on the cloud is typically stored in databases, file systems, and object storage services. The optimal storage service is determined by the data type (for example, structured vs. unstructured) and access patterns required by the workload. For each data store, use a combination of IAM permissions and any additional authorization methods to secure stored health data.

For object storage on Amazon S3, use access policies attached to your resources (buckets and objects) to implement additional authorization if necessary. More information can be found at [identity and access management in Amazon S3](#). Health data residing in a data lake based in Amazon S3, including those managed by AWS Lake Formation, should consider implementing column, row, and cell-level authorization controls where appropriate.

Use operating system file system permissions to limit access to health data stored on instance storage, including when using managed storage services such as Amazon Elastic File System and Amazon FSx for Lustre. Additionally, use resource and condition statements within IAM policies to limit IAM principal access to file systems when using managed storage services.

Control access to managed file systems through narrowly scoped security groups to prevent unauthorized resources from connecting to the file system.

Sensitive data stored in managed database services, such as Amazon Aurora, Amazon Relational Database Service, Amazon Redshift, and Amazon DynamoDB, implement authorization rules using a combination of IAM permissions and any additional authorization mechanisms available in the AWS service. For example, Amazon Redshift supports access controls as the column-level to limit users access to columns that may contain sensitive data. The AWS documentation for each managed database service contains a section titled Identity and Access Management which documents the access configuration options.

Detective controls

HCL_SEC3. How are you logging access to health data?

Log access to systems, resources, and data in accordance with your policies and procedures

If your workload hosts health data, then under the [Architecting for HIPAA Security and Compliance on Amazon Web Services whitepaper](#) you must implement and maintain logging of access to that data in accordance with the regulatory frameworks applicable to your workload. AWS makes it easy to log access to health data stored in many services with AWS CloudWatch and AWS CloudTrail.

AWS also provides service-specific mechanism to audit access to health data and health data systems. For audit logging details, see the [Architecting for HIPAA Security and Compliance on Amazon Web Services whitepaper](#).

Configure audit logs to be centralized and immutable

Environments that host and process health data should record and audit any person or system that accesses the data. Such logging provides evidence that the proper people and systems are accessing health data, and can be helpful in investigating a security incident. Configure logging to save to a centralized location and the logs made immutable to verify their integrity in the event of a forensic requirement. Prevent modification of log data by creating an AWS account in your organization that is designated to host audit logs and implement strict authorization rules. AWS audit and logging services, such as CloudWatch and CloudTrail, can save logs to a central location, yielding one set of logs that encompass an entire IT environment.

Use CloudTrail to log actions taken by users, roles, and AWS services across your AWS infrastructure. Enable CloudTrail log file integrity validation to prevent modification, deletion, or forgery of CloudTrail log files without detection.

Enable AWS Config in all AWS accounts to assess, audit, and evaluate resources within your AWS environment. AWS Config maintains a database of resources, and their associated configurations. This provides an audit record of AWS resource configurations over time.

Capture network layer logs to track the transport layer activity going to and from network interfaces in your VPC using VPC Flow Logs. When using Elastic Load Balancing, enable access logs to capture detailed information about the requests received and processed by one or more load balancers, including client IP addresses, request paths, and server responses. Similar approaches should be employed for other AWS services, such as Amazon API Gateway, which offer similar functionality.

Configure operating system and application logs, including managed compute services like AWS Lambda, Amazon Elastic Container Service, and Amazon Elastic Kubernetes Service, to send logs to CloudWatch log groups. CloudWatch log groups can be configured to forward logs to a centralized account for long-term retention. Develop processes and coding standards to avoid putting sensitive information into logs. Additionally, use AWS encryption services to encrypt log data. When using managed database services to store health data, such as Amazon RDS and Amazon Redshift, enable database level audit logging to collect information about connections and user activity within the database. You can use service features to publish database logs to CloudWatch, simplifying centralized log management.

Enable and configure Amazon S3 access logging for any Amazon S3 buckets that may contain sensitive health data. Amazon S3 Access Logs record every upload, download, and modification to stored objects.

Refer to the AWS documentation for each AWS service to find the supported service-specific logging options.

HCL_SEC4. How often do you review audit logs?

Create, document, and follow a policy and procedure to regularly review audit logs

In addition to the creation and documentation of an audit log review policy and procedure, organizations who are auditing access to health data should also have systems and procedures in place to review the audit logs on a regular basis. Facilitate audits by collecting all logs in a centralized location. For example, AWS CloudTrail can be configured to deliver logs from multiple accounts to a single Amazon S3 bucket. This provides both an easier location allowing regular review of the logs, while limiting the scope of access required for the reviewer by limiting them to a single location rather than multiple accounts.

Enable [CloudTrail Insights](#) to identify unusual activity in CloudTrail logs in order to help improve the audit log review process.

Automate alerts for potential anomalies detected in logs

Additionally, use automated systems that will generate alerts if anomalies are detected in logs. For example, create [CloudWatch alarms based on anomaly detection](#) that uses previously recorded metrics to create a model of expected results. You can also use [the Amazon OpenSearch Service to detect anomalies in logs](#). Enable CloudTrail Insights to detect unusual operational activity that is recorded in your CloudTrail audit logs. Review all applicable regulatory frameworks and standards and ensuring the specific requirements are being met. Configure all alarms to be received by an identified owner, ensuring that the alarm is acknowledged, triaged, and actioned. Finally, create and follow a procedure that outlines a regular cadence to review all automation configurations for continued accuracy, sufficiency, and relevance of the alerts.

Infrastructure protection

HCL_SEC5. How does your organization protect critical systems?

Follow Well-Architected best practices for [infrastructure protection](#) when designing and managing your transactional systems of record.

Implement security controls necessary to protect the infrastructure within the AWS account

Migrated healthcare workloads may have dependencies on technologies, older software applications, and host operating systems. For such systems, limit network access to sensitive hosts, apply the latest available security patches, and the employ monitoring practices described above.

Enable [Amazon GuardDuty](#) in accounts that host and process PHI to add intelligent threat detection to your environment. GuardDuty continuously monitors your AWS accounts and workloads for malicious activity and provides detailed security findings. You can also create custom, automated [responses to GuardDuty findings using Amazon CloudWatch Events](#).

For details on workload protection, see the [security pillar of the AWS Well-Architected Framework](#).

Data protection

HCL_SEC6. How do you determine and enforce data residency requirements?

Determine applicable regulatory frameworks and controls as it pertains to data locality

Many healthcare organizations fall under data locality requirements or regulations on where data may be physically located. Begin by reviewing these requirements within any applicable regulatory frameworks. To determine applicable regulatory frameworks, start with local regulations and frameworks for the country where your sensitive healthcare data is generated, hosted, and processed. Engage with legal counsel who can help you define the scope of the local regulations, as well as any additional regulation frameworks that may apply to you.

Enforce data locality requirements by implementing controls

Once the determination of requirements has been made and documented, technical controls can be put in place to enforce them.

The AWS Cloud spans many Availability Zones and geographic Regions around the world. Each AWS Region is fully isolated, and comprised of multiple Availability Zones. You can choose to use one or many AWS Regions in your environment. AWS stores and processes your content in the AWS Regions you select, using the services you select. AWS will not move your content without your consent, except as legally required. This allows you to establish environments in a location or locations of their choice. For example, AWS customers in Germany can choose to deploy their AWS services exclusively in one AWS Region, such as the Europe (Frankfurt) Region, and store their content onshore in Germany.

AWS also provides mechanisms for customers to allow or deny access to specific Regions. When using AWS Organizations, implement service control policies (SCP) to limit access to specific AWS services or resources. SCPs can be used with [AWS Control Tower data residency controls](#) to create additional guardrails to enforce data locality requirements. Continuing the example above, you can implement a service control policy that allows only access to the Europe (Frankfurt) Region and denies access to all others. A sample service control policy, or SCP, can be found [here](#). Additionally, if you are using AWS Control Tower, implement the [Region deny guardrail](#) to deny access to specific Regions.

HCL_SEC7. How are you protecting health data at rest and in transit?

Encrypt sensitive health data at rest and in transit at all times

Protect all sensitive data stored and transmitted within a cloud environment with AWS encryption services. The AWS Business Associate Addendum (BAA), applicable to customers who align with the Health Insurance Portability and Accountability Act (HIPAA), requires the encryption of protected health information (PHI) as defined by HIPAA at rest and in transit. Encryption at rest and in transit may be required by other applicable frameworks.

As documented in the [data encryption](#) section of the [Introduction to AWS Security](#) whitepaper and the [Encrypting Data-at-Rest and -in-Transit](#) section of the [Logical Separation on AWS](#) whitepaper, AWS allows you to add a layer of security to your data at rest in the cloud, providing scalable and efficient encryption features. These include:

- Data at rest encryption capabilities available in most AWS services, such as Amazon Elastic Block Store (Amazon EBS), Amazon S3, Amazon RDS, Amazon Redshift, Amazon ElastiCache, AWS Lambda, and Amazon SageMaker

- Flexible key management options, including AWS Key Management Service (AWS KMS), that allow you to choose whether to have AWS manage the encryption keys or enable you to keep complete control over your own keys
- Dedicated, hardware-based cryptographic key storage using AWS CloudHSM, allowing you to help satisfy your compliance requirements
- Encrypted message queues for the transmission of sensitive data using server-side encryption (SSE) for Amazon Simple Queue Service (Amazon SQS)

In addition, AWS provides APIs for you to integrate encryption and data protection with any of the services you develop or deploy in an AWS environment.

Implement encryption controls as part of the infrastructure architecture

Use infrastructure as code to declare encryption as a configuration when creating an environment template. Use alerts and automated remediation, where possible, with AWS Config that can detect when a resource is not configured to use encryption. Where automated remediation is not available, verify that alerts are generated and sent to the appropriate parties.

Amazon EBS, block-storage for use with Amazon EC2 Auto Scaling, uses AWS KMS keys to encrypt storage volumes and snapshots. Encryption operations occur on the servers that host Amazon EC2 instances, ensuring the security of data at rest and data in transit between an instance and attached Amazon EBS volumes. Consider configuring your AWS accounts for Amazon EBS [encryption by default](#), enforcing that any created Amazon EBS volumes are encrypted automatically.

Amazon RDS uses configuration policies that can enforce encrypted connections to a hosted database. The Amazon RDS documentation contains information on [using SSL/TLS to encrypt a connection to the database](#), including options for enforcing the connection through either parameter groups (Amazon RDS for Postgres, Aurora for Postgres, Amazon RDS for MariaDB, or Amazon RDS for Microsoft SQL Server) or option groups (Amazon RDS for Oracle).

Amazon S3 buckets that may contain health data should be configured to require secure connections with the use of a bucket policy, and require encryption of data at rest. Below is an example of a bucket policy that explicitly denies access when a secure transport connection is not used:

```
{
  "Id": "ExamplePolicy",
```

```
"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "AllowSSLRequestsOnly",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    ],
    "Condition": {
      "Bool": {
        "aws:SecureTransport": "false"
      }
    },
    "Principal": "*"
  }
]
```

Encryption in transit between systems must be identified and enforced where possible. AWS services provide HTTPS endpoints using TLS for communication, providing encryption in transit when communicating with the AWS APIs. Insecure protocols, such as HTTP, can be blocked in a VPC through the use of security groups. HTTP requests can also be [automatically redirected to HTTPS](#) in Amazon CloudFront or on an [Application Load Balancer](#).

When hosting applications on Amazon EC2 instances, use open standard transport encryption mechanisms such as Transport Layer Security (TLS) to encrypt data during transit between instances and endpoints. Certain Amazon EC2 instance types can offload encrypting traffic between instances to the underlying Nitro System hardware, using Authenticated Encryption with Associated Data (AEAD) algorithms with 256-bit encryption. For more detail and a list of supported instance types, see [encryption in transit](#).

Microservice architectures should also consider controls to enforce encryption of data between services when hosting and processing health data. AWS App Mesh, a service mesh that makes it easy to monitor and control service, features Transport Layer Security (TLS) encrypts communication between the Envoy proxies deployed on compute resources that are represented in App Mesh. When the proxy is deployed with an application, your application code is not responsible for negotiating a TLS session. The proxy negotiates TLS on your application's behalf. For more detail, see [Transport Layer Security \(TLS\)](#).

HCL_SEC8. How do you isolate sensitive data?

Isolate health data from non-health data

Organizations working with health data should take steps to isolate and segment health data from non-health data. In conjunction with the recommendations around data discovery and classification, it is important to separate health data so the organization can implement the necessary technical and administrative controls.

Isolation can be accomplished through a variety of methods depending on the cloud environment. AWS recommends beginning with using multiple AWS accounts and designating specific accounts as containing health data. AWS accounts provide a level of segmentation that allows strict controls to be put in place for workloads that host and process health data. [AWS Organizations](#), an account management service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage, provides management capabilities that allow you to group like accounts into organizational units (OUs) and apply policies at the OU level. This verifies that all accounts within that OU are using a standardized set of policies and controls, which can help organizations align to their specific compliance needs.

Furthermore, as outlined in the [Organizing Your AWS Environment Using Multiple Accounts whitepaper](#), when you limit sensitive data stores to an account that is built to manage it, you can more easily constrain the number of people and processes that can access and manage the data. This approach simplifies the process of achieving least privilege access. Limiting access at the coarse-grained level of an account helps contain exposure to highly sensitive data.

Limit access to health data

It is also important to limit access to health data within an account. Use resource isolation, such as designated Amazon S3 buckets, to separate health data from non-health data. Resource isolation can also be used to isolate tenants and tenant-specific data. Resource isolation and tenant isolation reinforce the benefits to account isolation, limiting access to sensitive data to only the people and systems that require it, without unnecessarily blocking access to less sensitive data. Refer to the Security pillar section of the [Well-Architected Framework SaaS Lens](#) for additional recommendations on tenant isolation.

Tagging allows limiting access to specific resources through the use of [conditional statements](#) in IAM policies. By adding a specific resource tag to an IAM policy, such as `data type: health`,

organizations can allow or deny access to resources with that tag. This approach adds an additional layer of authorization to resources that host and process health data.

Incident response

HCL_SEC9. What is your disaster recovery for critical systems?

Mitigate and respond to potential incidents by creating policies, procedures, and playbooks

Healthcare, and health data, are valuable targets for malicious actors. Create policies, procedures, and playbooks designed to respond to and mitigate the potential impact of a security event or natural disaster. This includes exercises that practice the response to a simulated incident using the defined policies, procedures, and playbooks to prepare your organization.

As malicious actors continue to target healthcare and health data owners with attacks such as ransomware, implement a data availability strategy to help reduce the potential impact. This can include backups that are stored in a separate AWS account with authorization controls in place to prevent modifying the backup (such as setting the backup as read only) or a [pilot light disaster recovery environment](#). Create specific policies, procedures, and playbooks for ransomware to prepare your organization.

The [incident response section of the security pillar in the AWS Well-Architected Framework](#) contains further details on preparing for and responding to security incidents in the cloud.

Key AWS services

The AWS service that is essential to security is AWS Identity and Access Management, which allows you to securely control access to AWS services and resources for your users. The following services and features support the four areas of security:

Identity and access management:

- [AWS Identity and Access Management](#);
- [AWS IAM Identity Center](#)
- [AWS Directory Service](#)
- [Amazon Cognito](#)

Detection:

- [Amazon CloudWatch Logs](#)
- [Amazon Detective](#)
- [Amazon GuardDuty](#)
- [Amazon Inspector](#)
- [AWS CloudTrail](#)
- [AWS Config](#)
- [AWS Security Hub](#)

Infrastructure protection:

- [AWS Key Management Service](#)
- [AWS CloudHSM](#)
- [AWS Systems Manager](#)

Data protection:

- [Amazon Macie](#)
- [Amazon CloudFront](#)
- [Application Load Balancer](#)
- [AWS Config](#)
- [AWS Network Firewall](#)
- [AWS Virtual Private Network](#)

Incident response:

- [Amazon GuardDuty](#)
- [Amazon Detective](#)
- [Amazon Inspector](#)
- [Amazon EventBridge](#)
- [AWS Security Hub](#)

Resources

Refer to the following resources to learn more about our best practices for security.

Documentation and blogs:

- [Encryption-in-transit for public sector workloads with AWS Nitro Enclaves and AWS Certificate Manager](#)

Whitepapers:

- [Architecting for HIPAA Security and Compliance on Amazon Web Services](#)
- [Introduction to AWS Security](#)
- [Data residency: AWS policy perspectives](#)
- [Logical Separation on AWS](#)

Videos:

- [AWS Security Webinar: The Key to Effective Cloud Encryption](#)

Reliability pillar

The reliability pillar encompasses the ability of a workload to perform its intended function correctly and consistently when it's expected to. This includes the ability to operate and test the workload through its total lifecycle.

Design principles

In the cloud, there are a number of principles that can help you increase reliability:

- **Evaluate and understand availability and latency implications:** Many healthcare applications are latency-sensitive. Healthcare workloads may also be life-critical, meaning that service interruptions may lead to patient harm. Consider latency, connectivity, and availability requirements in defining the cloud architecture of your workload. Choose your deployment option to meet those requirements based on modeling real-world scenarios.
- **Address potential implications by defining reliability and availability requirements:** Understand the availability requirements for your solutions and the consequences of service

disruptions before designing your architecture. From that understanding, define the specific requirements for reliability and availability.

- **Understand the end user setting:** Document the settings your end users will use for your solution. Some healthcare solutions are only used within hospitals, while others may be used by providers working remotely. Use your understanding of the settings by which end users access your solution to meet reliability requirements.

Best practices

Definition

To achieve reliability, a system must have a well-planned [foundation](#) in place, with mechanisms for handling changes in demand or requirements. The system should be designed to detect failure and automatically repair itself.

The following are best practice areas for reliability in the cloud:

- [Foundations](#)
- [Workload architecture](#)
- [Change management](#)
- [Failure management](#)

Foundations

HCL_REL1. How does your system adapt to changes in demand?

Architect systems for elasticity

Healthcare applications often have time-based peaks in demand. For example, clinical systems may need to respond to periods of high demand using either time-based or usage-based metrics to define automatic scaling rules. An example of periods of high demand could be daytime business hours or a known event such as open-enrollment for an insurance provider.

Where possible, implement architectures that automatically adapt to changes in demand.

Embracing elasticity enables healthcare organizations to *right-size* performance during all hours of

the day while minimizing excess cost. General distributed systems recommendations apply. Where possible, leverage AWS services that allow you to scale with demand.

AWS Auto Scaling simplifies how you adapt to demand. In all cases, you should be collecting metrics to inform your scaling actions and detect events that could impact reliability. You may either use default CloudWatch metrics or define custom metrics that track specific aspects of utilization. For example, software as a service (SaaS) applications may monitor demand with custom CloudWatch metrics that capture the number of concurrent active users, or open sessions supported by the application.

Healthcare solutions implemented as distributed systems may use multiple services that can scale elastically. In these cases, implement automatic scaling for all relevant compute and database services, using Amazon EC2 Auto Scaling and Application Auto Scaling for [supported AWS services](#). You can leverage scaling plans to simplify configuration of scaling rules, and follow auto scaling [best practices](#) to maximize reliability.

Other managed services, such as Amazon Data Firehose, allow you to operate at scale without worrying about capacity or managing infrastructure.

Workload architecture

HCL_REL2. How do you ensure acceptable network availability for your healthcare workloads?

Architect redundant and reliable network connections to ensure care continuity

Many healthcare applications, such as those in a hospital, require secure connectivity between the cloud and on-premises resources and users. When evaluating your network setup, consider your business continuity and disaster recovery requirements. Certain healthcare applications will be more accepting of shifts in latency or availability compared to others. For example, medical imaging or EHR systems may require more consistent latency and connectivity compared to other systems in a hospital.

Following the [Well-Architected Reliability Pillar](#), redundant, encrypted connections are critical to verify service continuity and, more importantly, consistent patient care. Many best practices can be found in the [AWS Well-Architected Framework Hybrid Networking Lens](#).

AWS Direct Connect is key to establishing consistent, redundant connections with your on-premises data sources. Direct Connect establishes a dedicated network connection to AWS. It is possible to create multiple connections to AWS from a single location.

There are [two common approaches](#) to establishing redundancy across your network connection to AWS:

- Redundant Direct Connect connections: Use the [AWS Direct Connect resiliency toolkit](#) to enable resilient applications and achieve an SLA of 99.99%.
- [Failover to public internet connection with VPN routing](#). Customers can either connect to Amazon VPCs using VPNs or through AWS Transit Gateway.

In both cases, check that encryption is enabled for all connections into AWS as well as within AWS. Encryption can be enabled at multiple layers, such as at the network or application layers.

Change management

There are no reliability best practices for change management specific to the Healthcare Industry Lens.

Failure management

There are no reliability best practices for failure management specific to the Healthcare Industry Lens.

Key AWS services

The AWS service that is key to ensuring reliability is Amazon CloudWatch, which monitors runtime metrics. Other services and features that support the four areas of reliability are as follows:

Foundations:

- [AWS Direct Connect](#)

Workload architecture:

- N/A

Change management:

- [AWS Config](#)
- [Amazon EC2 Auto Scaling](#), and [Application Auto Scaling](#)

Failure management:

- [AWS Backup](#)
- [AWS Transit Gateway](#)

Resources

Refer to the following resources to learn more about our best practices related to reliability.

Video and analyst report:

- [Connectivity to AWS and hybrid AWS network architectures](#)

Documentation and blogs:

- [AWS Direct Connect User Guide](#)

Whitepapers:

- [Building a scalable and secure multi-VPC AWS network infrastructure](#)

Performance efficiency pillar

The performance efficiency pillar focuses on the efficient use of computing resources to meet requirements and maintaining that efficiency as demand changes and technologies evolve.

Design principles

The following design principles can help you achieve and maintain efficient workloads in the cloud.

- **Democratize advanced technologies:** Make advanced technology implementation easier for your team by delegating complex tasks to your cloud vendor. Rather than asking your IT team to learn about hosting and running a new technology, consider consuming the technology as a service. For example, NoSQL databases, media transcoding, and machine learning are all

technologies that require specialized expertise. In the cloud, these technologies become services that your team can consume, allowing your team to focus on product development rather than resource provisioning and management.

- **Go global in minutes:** Deploying your workload in multiple AWS Regions around the world allows you to provide lower latency and a better experience for your customers at minimal cost.
- **Use serverless architectures:** Serverless architectures remove the need for you to run and maintain physical servers for traditional compute activities. For example, serverless storage services can act as static websites (removing the need for web servers) and event services can host code. This removes the operational burden of managing physical servers, and can lower transactional costs because managed services operate at cloud scale.
- **Experiment more often:** With virtual and automatable resources, you can quickly carry out comparative testing using different types of instances, storage, or configurations.
- **Consider mechanical sympathy:** Use the technology approach that aligns best with your goals. For example, consider data access patterns when you select database or storage approaches.

Best practices

Definition

Take a data-driven approach to selecting a high-performance architecture. Gather data on all aspects of the architecture, from the high-level design to the selection and configuration of resource types. By reviewing your choices on a cyclical basis, you verify that you are taking advantage of the continually evolving AWS Cloud. Monitoring keeps you aware of any deviance from expected performance and can act on it. Finally, your architecture can make tradeoffs to improve performance, such as using compression or caching, or relaxing consistency requirements.

The following are best practice areas for performance efficiency in the cloud:

- [Selection](#)
- [Review](#)
- [Monitoring](#)
- [Trade-offs](#)

Selection

The optimal solution for a particular workload varies, and solutions often combine multiple approaches. Well-Architected workloads use multiple solutions and enable different features to improve performance.

Performance architecture

HCL_PERF1. How do you encrypt data while ensuring performance?

Offload encryption to hardware

Certain encryption approaches, such as VPN tunnels or IPsec meshes, can impact performance when implemented at scale. Where possible, offload encryption to hardware to maintain security while improving performance.

The [AWS Nitro System](#) provides hardware components that allow for easy offloading of encryption services to the hardware. For example, some instance types use the hardware capabilities of the Nitro System hardware to encrypt in-transit traffic between instances with no impact to network performance. This allows healthcare organizations to enable encryption in-transit for sensitive healthcare data. Use instance types that support the Nitro System where possible.

Compute selection

HCL_PERF2. How do you select your compute solution?

Select compute services that meet regulatory and performance requirements

Healthcare requirements for compute are generally consistent with other industries. Guidance from the [Well-Architected Framework Compute Architecture Selection](#) still applies. Healthcare applications can take advantage of virtual machines, containers, or serverless technologies.

Healthcare applications should enable encryption in-transit at one of the OSI layers. Some legacy communication protocols, such as Minimal Lower Layer Protocol (MLLP) for healthcare interoperability, may not natively support encryption in-transit. A common industry solution has been to overlay a VPN or create an IPsec mesh on top of virtual machines in a VPC to encrypt

sensitive data in transit; however, such approaches can create performance penalties. Instead, where possible, use [Amazon EC2 instances with encryption in-transit](#) handled by the underlying Amazon EC2 Nitro System to reduce any performance penalties associated with inter-Amazon EC2 communication.

You can get a full list of Amazon EC2 instance types that support this feature with the following CLI command:

```
aws ec2 describe-instance-types --filters
    Name=network-info.encryption-in-transit-supported,Values=true
    --query "InstanceTypes[*].[InstanceType]" --output text
```

Network architecture

HCL_PERF3. How do you define and test network performance requirements?

Healthcare requirements for compute are generally consistent with other industries. Guidance from the [Well-Architected Framework Network Architecture Selection](#) still applies.

Storage architecture

HCL_PERF4. How do you define and test storage performance requirements?

Healthcare requirements for compute are generally consistent with other industries. Guidance from the [Well-Architected Framework Storage Architecture Selection](#) still applies.

AWS offers storage with extreme durability and performance. For example, [Amazon S3](#) provides 99.999999999% (11 nines) of data durability of objects over a given year. [Amazon EBS io2](#) block storage offers not only 99.999% durability, but up to 500 IOPS per GiB, enabling high performance and durability for healthcare workloads that require higher performance, such as transactional databases.

Review

There are no performance efficiency best practices for review specific to the Healthcare Industry Lens.

Monitoring

There are no performance efficiency best practices for monitoring specific to the Healthcare Industry Lens.

Trade-offs

There are no performance efficiency best practices for trade-offs specific to the Healthcare Industry Lens.

Key AWS services

The key AWS service for performance efficiency is Amazon CloudWatch, which monitors your resources and systems, providing visibility into your overall performance and operational health. The following services are important in the areas of performance efficiency:

- [Amazon CloudWatch](#)
- [Amazon EC2](#)
- [AWS Nitro System](#)
- [AWS Client VPN](#)
- [AWS Site-to-Site VPN](#)

Resources

Refer to the following resources to learn more about our best practices related to performance efficiency.

Videos

- [Powering next-gen Amazon EC2: Deep dive on the Nitro System](#)

Documentation

- [AWS Well-Architected Framework: Performance Efficiency Pillar: Performance Architecture Selection](#)
- [Data protection in Amazon EC2: Encryption in-transit](#)

Cost optimization pillar

The cost optimization pillar includes the continual process of refinement and improvement of a system over its entire lifecycle, helping you build and operate cost-aware systems that achieve business outcomes and minimize costs.

Best practices

Definition

As with the other pillars, there are tradeoffs to consider. For example, do you want to optimize for speed to market or for cost? In some cases, it's best to optimize for speed—going to market quickly, shipping new features, or simply meeting a deadline—rather than investing in upfront cost optimization. Design decisions are sometimes guided by haste as opposed to empirical data, as the temptation always exists to overcompensate in case of issues rather than spend time benchmarking for the most cost-optimal deployment. This often leads to drastically over-provisioned and under-optimized deployments. The following sections provide techniques and strategic guidance for the initial and ongoing cost optimization of your deployment.

The following are best practice areas for cost optimization in the cloud:

- [Practice cloud financial management](#)
- [Expenditure and usage awareness](#)
- [Cost-effective resources](#)
- [Manage demand and supply resources](#)
- [Optimize over time](#)

Practice cloud financial management

There are no cost optimization best practices for practicing cloud financial management specific to the Healthcare Industry Lens.

Expenditure and usage awareness

Understanding your organization's costs and drivers is critical for managing your cost and usage effectively, and identifying cost-reduction opportunities. Organizations typically operate multiple workloads run by multiple teams. These teams can be in different organization units, each with its own revenue stream. The capability to attribute resource costs to the workloads, individual

organization, or product owners drives efficient usage behavior and helps reduce waste. Accurate cost and usage monitoring allows you to understand how profitable organization units and products are, and allows you to make more informed decisions about where to allocate resources within your organization. Awareness of usage at all levels in the organization is key to driving change, as change in usage drives changes in cost.

Data retention

HCL_COST1. How do you define and enforce data retention policies?

Determine applicable regulatory frameworks and controls as it pertains to data retention

Healthcare organizations may be subject to regulatory and company requirements dictating how long they must store both health data as well as logs detailing access to that data. Over time, storage requirements can reach petabyte scale for an individual organization. New imaging modalities, such as digital pathology, have the potential to push data volumes even higher. Developing strategies for data retention is imperative to maintain compliance while minimizing cost.

Healthcare organizations should establish a data retention policy identifying the types and duration that data should be retained in accordance with internal and external requirements.

Implement data lifecycle policies

As healthcare data ages, its access frequency declines. Implement lifecycle policies that transition infrequently-accessed data to lower-cost storage tiers. When designing your policies for each type of data, be sure to factor the size of the data, the frequency of access, and expectations for retrieval time; these are the three predominant cost-drivers. For example, use Amazon S3 Lifecycle configurations to archive infrequently accessed data after some period of time, automatically reducing storage costs. Alternatively, use Amazon S3 Intelligent-Tiering to shift the archival policies to focus on time of last access instead of time the object has been in Amazon S3.

Centralize automated policy enforcement

Adopting infrastructure as code, as discussed in the operational excellence pillar, enables you to define and test your data retention policies before they make it into production. For example, if regulatory requirements specify that you must retain certain medical images for 10 years, you can verify that no Amazon S3 lifecycle policy expires an object before that time. Additionally, consider

AWS Backup for centralized backup management, which enables you to back up application data in a consistent and compliant manner.

Validate lifecycle policies are enforced

Because the cloud is API-driven, you can monitor changes to your environment as described in the operational excellence and security tiers. Set up alerts for when an API action alters a data retention policy so you can quickly review the change to make sure it was authorized and operating correctly. If using AWS Backup, use AWS Backup Audit Manager to automatically detect when your AWS Backup policies violate your data retention requirements.

Cost-effective resources

There are no cost optimization best practices for cost-effective resources specific to the Healthcare Industry Lens.

Manage demand and supply resources

There are no cost optimization best practices for managing demand and supplying resources specific to the Healthcare Industry Lens.

Optimize over time

There are no cost optimization best practices for optimizing over time specific to the Healthcare Industry Lens.

Key AWS services

The key AWS feature that supports cost optimization is cost allocation tags, which help you to understand the costs of a system. The following services and features are important in the four areas of cost optimization:

Cost-effective resources:

- [AWS Compute Optimizer](#)
- [AWS Trusted Advisor](#)

Expenditure and usage awareness:

- [Amazon CloudWatch](#) – For monitoring access frequency statistics for objects stored on Amazon S3, like medical images.

- [AWS Cost Explorer](#)
- [AWS Budgets](#)

Resources

Refer to the following resources to learn more about our best practices related to cost optimization.

Videos

- [Simplify Your Data Lifecycle and Optimize Storage Costs With Amazon S3 Lifecycle](#)
- [Simplify Backup Auditing and Compliance with AWS Backup Audit Manager - AWS Online Tech Talk](#)

Documentation and blogs

- [How to manage retention periods in bulk using Amazon S3 Batch Operations](#)
- [Amazon S3 cost optimization for predictable and dynamic access patterns](#)
- [Amazon S3: Managing your storage lifecycle](#)
- [Audit backups and create reports with AWS Backup Audit Manager](#)

Whitepapers

- [Storage Best Practices for Data and Analytics Applications](#)
- [Data Classification](#)

Sustainability pillar

The sustainability pillar provides guidance on how to understand the environmental impact of cloud workloads, how to quantify impacts through the workload lifecycle, and how to apply design principles that help minimize these impacts.

Health expenditures account for approximately 10% of global economic output, and consequently, healthcare impacts the environment through emissions, pollutants, and water consumption. For example, [the healthcare sector is estimated to generate between 7.9% and 9.8% of all US greenhouse gas emissions](#). As a result, it is important for healthcare organizations to take steps to

minimize this impact. For example, the UK has included their national health services in their plans to meet Paris Agreement commitments on climate change mitigation. In the US, independent, forward-looking healthcare delivery systems are making [investments to achieve carbon neutrality](#).

A central mission of healthcare is to [improve population health, with consideration of the patient experience](#). The environmental impacts of healthcare have downstream impacts on the health and experience of communities, making considerations of sustainability aligned with that central mission. Underserved communities may be more susceptible to environmental impacts and climate change more broadly. Thus, the goals of decreasing health inequities and improving sustainability are linked.

Many organizations in the healthcare vertical operate on thin profit margins (or operate at a loss), which limits their capacity to make sustainability investments. Fortunately, the Well-Architected best practices for sustainability can also help to lower total cost of ownership for healthcare workloads. As presented below, organizations can decrease costs as they mitigate downstream environmental impacts, supporting better health across our communities.

The [Well-Architected Sustainability pillar](#) offers an improvement process to guide efforts to minimize unfavorable environmental impacts for all cloud workloads. For healthcare workloads, the following considerations and best practices should also be considered.

Best practices

Definition

The following are best practice areas for sustainability in the cloud:

- [Region selection](#)
- [User behavior patterns](#)
- [Software and architecture patterns](#)
- [Measure results](#)
- [Data patterns](#)

Region selection

HCL_SUS1. How do you identify targets for sustainability improvement?

Prioritize targets for improvement by reviewing your workloads against the [sustainability principles](#)

Migrating workloads from on-premises data centers to the cloud can reduce the workload's carbon footprint by 88%. Regular review of cloud architecture for optimization opportunities can reduce carbon footprints as well. Both mean that healthcare workload migrations can deliver the benefits of the cloud and simultaneously improve sustainability. The [Sustainability pillar of the AWS Well-Architected Framework](#) contains a number of [best practices for sustainability in the cloud](#).

User behavior patterns

HCL_SUS2. How do you match workload infrastructure to user behavior patterns?

Scale infrastructure to continually match user demand and performance requirements

Many healthcare workloads are life-critical, and have steady demand 24 x 7. However, other workloads (such as those supporting ambulatory care delivery or revenue cycle workflows) exhibit cyclical utilization patterns with peak demand during business hours. [Minimize the amount of hardware used](#) by scaling workloads down during periods of low demand.

Legacy solutions may use statically provisioned infrastructure, with redundancy for high availability. Consider cloud-native ways to meet business requirements with an elastic, efficient architecture and disaster recovery strategy (like a pilot light architecture rather than active-active).

Software and architecture patterns

HCL_SUS3. Does your organization monitor workload activity and remove or refactor components that are no longer necessary?

Analyze demand on workloads to identify components that can be removed or refactored. Then, engage component owners and stakeholders to redesign clinical workflows, and decrease workload infrastructure

Some healthcare delivery systems and large independent software vendors (ISV) have sprawling IT footprints with numerous siloed systems. Identifying and [removing or refactoring components](#)

with little or no use can simplify workflows, decrease cost, and improve sustainability. Cloud archives can minimize the cost of retaining data from retired components.

HCL_SUS4. How do you optimize the impact of applications and the equipment that run them?

Evaluate the overall impact of applications, devices, and equipment

As documented in the [Sustainability pillar of the AWS Well-Architected Framework](#), it is recommended to [optimize impact on customer devices and equipment](#). For example, as new features are released for a healthcare application, build those features as backward compatible, minimizing the need for new hardware. Additionally, evaluate the potential impact of new or upgraded hardware requirements to minimize the overall impact when architecting new workloads or features.

Measure results

HCL_SUS5. How does your organization measure the effectiveness of sustainability efforts?

Quantify and report results to drive continuous improvement processes

The healthcare vertical extensively uses metrics and measures to quantify care quality, effectiveness, and patient experience. Adding metrics to quantify sustainability improvement can better align business interests with sustainability goals. Further, analysis of such reporting can help identify repeatable processes for achieving sustainability improvements.

Data patterns

HCL_SUS6. How does your organization remove unneeded or redundant health data?

Automate data retention processes that retain the minimum amount of health data required to meet regulatory and business requirements

Regulatory requirements may impose data retention periods on healthcare providers and ISVs.

However, it is common for health data to be retained in perpetuity, well beyond its useful life. Begin by reviewing and classifying data in line with your business and regulatory requirements, such as how long health data records must be retained. Review data assets with consideration of regulatory compliance, care delivery needs, and secondary use goals. Optimize storage costs and environmental impact by taking advantage of storage classes within cloud services and aligning with access patterns. Remove [unnneeded or redundant data](#). Use automated lifecycle policies wherever possible to automate the deletion or archival of unnecessary data.

Key AWS services

There are a number of key AWS features that support sustainability in the cloud. The following services and features are an important part of the sustainability journey:

Workload management:

- [AWS Compute Optimizer](#)
- [AWS Trusted Advisor](#)
- [Amazon EC2 Auto Scaling](#)
- [AWS Customer Carbon Footprint Tool](#)
- [AWS Data Exchange](#)

Resources

Refer to the following resources to learn more about our best practices related to sustainability.

Documentation and blogs

- [Introduction of the Sustainability Pillar for the AWS Well-Architected Framework](#)
- [Reducing carbon by moving to AWS](#)
- [AWS and sustainability in the cloud](#)
- [AWS enables sustainability solutions](#)

Scenarios

In this section, we cover five key healthcare scenarios and how they influence the design and architecture of your healthcare application workloads on AWS. We will present characteristics of each scenario and answers to Well-Architected Framework questions pertaining to the scenario, as well as reference architectures where applicable.

Key scenarios

- [Electronic healthcare record and revenue cycle systems](#)
- [Healthcare interoperability](#)
- [Medical imaging](#)
- [Healthcare analytics](#)
- [Machine learning for healthcare](#)

Electronic healthcare record and revenue cycle systems

Healthcare providers, payers, and SaaS ISVs operate transactional systems of record, such as electronic healthcare record and claims adjudication systems, that are critical to the parent organization's operation.

Characteristics of these solutions include:

- High-scale transactional databases that host patient, care delivery, and payment records. These databases vary by ISV and application, but commonly support high levels of concurrent transactions and employ complex data schemas.
- Workflow specific user interfaces, often catering to highly skilled end users. Healthcare customers may also customize commercial solutions to meet their specific needs.
- Embedded analytics and AI/ML algorithms that support end user workflows.
- Stringent availability and resilience requirements, as these systems are often critical to the operation of the parent organization.
- Integration with third party applications that make up the electronic medical record.

Questions

HCL_REL3. How does your organization define and design for availability and reliability requirements?

Cloud based transactional systems of record should have well-defined availability (how often the system can be unavailable) and reliability (how quickly the system can respond to an issue) requirements based on business objectives. These requirements should be created with input from the IT organization, as well as the clinical support organization, in order to create requirements that satisfy all stakeholders.

Once the business has established the requirements, use cloud native features such as architecting across multiple Availability Zones to make sure your system meets the requirements. The AWS global footprint of Regions and Availability Zones provides a significant number of geographic options when architecting for availability, reliability, and disaster recovery.

AWS products and services contain features that help customers meet these availability and reliability requirements. These can include the Amazon RDS Multi-AZ feature that automatically creates a replicated copy of your RDS instance in a second Availability Zone, as well as AWS Backup, which retains backups of your data across multiple Availability Zones to reduce the risk of data loss.

Electronic health records may have hybrid requirements that include on-premises data center infrastructure. AWS Direct Connect provisions redundant connectivity between private networks in the cloud and on-premises environments.

For more information on designing systems that meet the business availability and reliability requirements, see the [Reliability Pillar whitepaper](#).

HCL_PERF3. Does your organization meet IOPS and other performance requirements?

Cloud based transactional systems of record are considered commercial off the shelf (COTS) software. COTS vendors should provide the minimum and recommended compute, memory, storage, and other performance requirements. These requirements can be mapped to the associated cloud virtual compute and storage options, which provides a well-defined starting point.

After installing COTS software, performance metrics should be monitored during all aspects of the lifecycle, from testing through staging and production. Use [Amazon CloudWatch metrics](#) to track your systems over time, including the configuration of alarms for high usage. Once performance metrics are collected, through each lifecycle stage, the compute and storage requirements can be adjusted to optimize performance.

For more detail on monitoring, reviewing, and optimizing your workload, see the [Performance Efficiency Pillar whitepaper](#).

Healthcare interoperability

Healthcare interoperability refers to health data exchange between information systems like electronic healthcare records (EHR), pharmacy systems, diagnostic imaging systems, laboratory systems, and claims systems. With interoperability, health data can be communicated between electronic systems, between organizations, and across geographical boundaries with standardized protocols.

Enabling interoperability requires that data be captured in electronic systems, with standards for the content, transport, vocabulary/terminology, privacy and security, and identifiers. The Healthcare Information and Management Systems Society (HIMSS), a globally recognized thought leader on healthcare interoperability, [defines four levels of interoperability](#) as:

- **Foundational:** Establishing interconnectivity so that systems can securely exchange data.
- **Structural:** Defining and adopting standards for the format, syntax, and organization of data.
- **Semantic:** Using standardized coding vocabularies so that the parties exchanging data can have a shared understanding of its meaning.
- **Organization:** Managing the governance and legal aspects of exchanging data between organizations and individuals. This level also covers consent for how an individual's data will be shared.

Characteristics of functional interoperability solutions include:

- Using a published standard to structure and transmit health data. Many standards are currently used across healthcare, such as [Health Level Seven Version 2](#) (HL7 v2), HL7 Version 3 (HL7 v3), [HL7 Fast Healthcare Interoperability Resources](#) Version 4 (FHIR), [X12 Electronic Data Interchange \(EDI\) transactions](#) used in healthcare benefits coordination and claims processing, and the Digital

Imaging and Communications in Medicine (DICOM) standard for storing and transferring medical images.

- Exposing either an API or an integration server to other systems seeking to exchange data.
- Checking received data for conformance to the interoperability standard being used.
- Being able to apply transformations to data received or sent and mapping the internal formats of systems of record, like EHRs, to a structural and semantic form dictated by the standard.
- Facilitating secure authorization and launch of user-facing applications, often with protocols like [SMART on FHIR](#).

Interoperability reference architecture

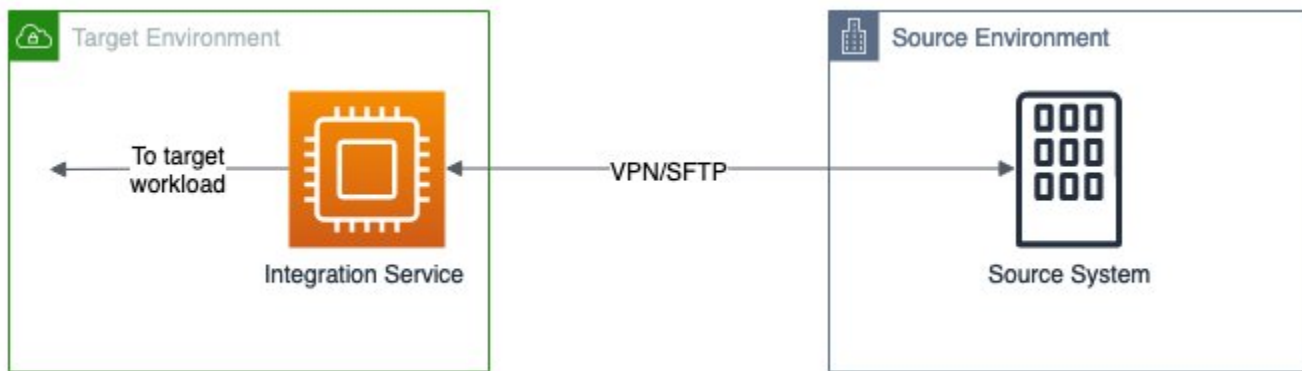
HL7 v2, HL7 v3, and HL7 FHIR are three of the most common healthcare messaging standards used today. This section highlights architectures for each.

HL7 v2

In practice, HL7 v2 tends to have reasonable structural standardization, but weak semantic interoperability. Structurally, HL7 v2 messages consist of one or more *segments* delineated as separate lines. Within each segment, *fields* are pipe-delimited. Semantically, the contents of HL7 v2 messages may use any number of standard vocabularies or lack externally linkable terms or codes entirely. If using HL7 v2, consider automated ways to map to standard ontologies, such as using NLP algorithms trained on medical text.

The Minimum Lower Layer Protocol (MLLP) is a common transport standard used for HL7 v2 interoperability. MLLP streams bytes using TCP/IP and uses special header and trailer (such as a footer) wrappers to signify the beginnings and ends of messages. Notably, the MLLP does not natively provide encryption. However, in recent years, many interfaces have started supporting MLLP over TLS for added security. For interfaces that do not support TLS, additional measures must be taken to verify encryption of data in transit.

An HL7 v2 interoperability architecture may leverage the HL7 v2 APIs of a source system, like an EHR. To provide end-to-end encryption between the source system and target system, it's a best practice to use a VPN tunnel or secure protocols like SSH File Transfer Protocol (SFTP), as MLLP may not provide TLS. The source system may asynchronously send batches of messages to an encrypted Amazon S3 bucket, or may synchronously exchange messages with a target integration service, as shown in the following figure.



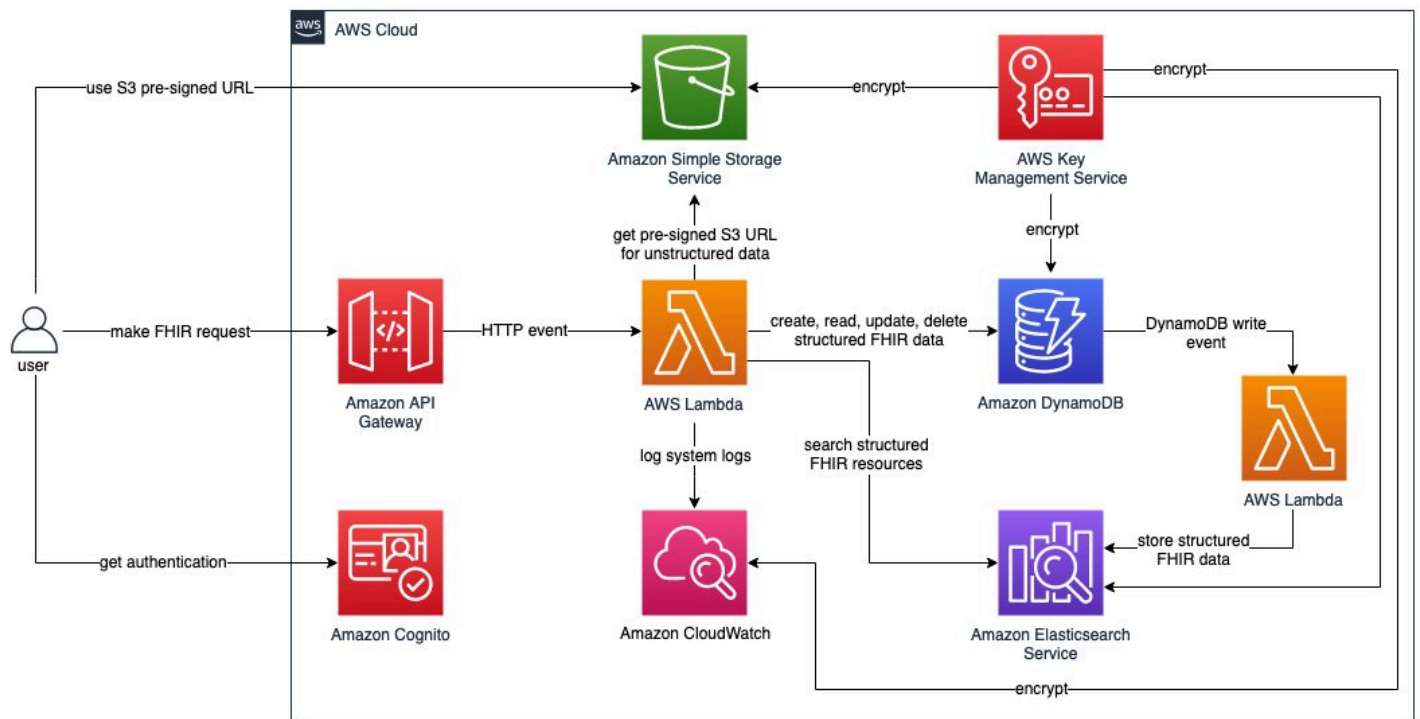
A common HL7v2 interoperability architecture where the integration may exchange HL7 v2 messages through a VPN tunnel or SFTP. MLLP over TLS can help keep messages encrypted in transit.

HL7 v3

The HL7 v3 standard is commonly used for exchanging Continuity of Care Document (CCD) and Consolidated Clinical Document Architecture (C-CDA) messages. These messages are often exchanged using Cross-Enterprise Document Sharing (XDS) transactions, which are based on web services. Amazon API Gateway can be used to create a web service to support these XDS transactions such as ITI-41. TLS with [certificated-based authentication](#) can be implemented at API Gateway for data security. [Resource policy](#) can be used to restrict access to the API Gateway by certain source IP addresses.

FHIR

The following diagram illustrates a representative FHIR interoperability architecture, which presents an integration point in front of one or more systems of record, typically within a provider or payer organization.



Example FHIR enabled interoperability architecture.

- An API or server endpoint is made available for integration with other systems or users. A managed API service, like Amazon API Gateway, ensures scalability and high availability. Restrict network access to the API if possible and use a web application firewall to filter malicious requests.
- Inbound requests are authenticated with [OAuth](#) or a service like Amazon Cognito to verify that sensitive health data is only exchanged with appropriate parties.
- Inbound payloads are checked for conformance to relevant structural interoperability standards, ensuring the safety and integrity of data exchange.
- Data may be written to or pulled from systems of record within the organization hosting the interoperability architecture.
- Many systems of record host health data in proprietary formats and schemas. Therefore, interoperability architectures perform some transformations to map data elements to and from the interoperability standard.
- The architecture may store the exchanged health data using a database service like DynamoDB, or it may exchange data without persisting it.

- [AWS HealthLake](#) enables use cases that require enrichment of health data in FHIR format or downstream analytics and machine learning. AWS HealthLake can improve semantic interoperability by applying natural language processing (NLP) to link concepts in unstructured data to terms in standard health ontologies, like ICD-10-CM, SNOMED CT, and RxNorm.

Questions

HCL_OPS7. How does your organization identify and prioritize which interoperability standards to adopt?

Leverage thought leadership organizations to understand emerging trends, and use the capabilities of your customers' systems. Many EHR and clinical systems have interoperability APIs. Older standards, like HL7 v2, are most prevalent. However, more modern standards can reduce the complexity of integration and enable new use-cases.

HCL_OPS8. Do you require unidirectional or bidirectional interoperability?

Older source systems may share data, but not support writes through interoperability APIs. Understand what interoperability is needed to enable your use-cases, and then how to achieve it in practice. Determine the importance of unidirectional transfer (sending data in only one direction between systems) or bidirectional transfer (exchanging data between two systems), as they will have different requirements and use cases. Custom integration work, or support from AWS Partner Network (APN) partners or ISVs, may be required to achieve bidirectional interoperability.

HCL_OPS9. How do you standardize terminology to support semantic interoperability?

Work with your customer to understand what vocabularies they employ and how you can link to standards. Where possible, leverage modern standards that provide semantic interoperability, such as Consolidated-Clinical Document Architecture (C-CDA). AI-based services, like [Amazon Comprehend Medical](#), can link concepts to standard ontologies in cases that deterministic linking is not possible.

HCL_SEC11. How do you protect integration endpoints or APIs?

Use end-to-end encryption of health data exchanged over the network. Older standards, like the Minimum Lower Layer Protocol (MLLP) used with HL7 v2, may not natively support TLS. In such cases, protect data with a VPN tunnel or additional encryption solution. For MLLP interfaces that support TLS, use certificates for authentication and encryption. Security groups can be used to protect integration endpoints to only allow traffic from specific IP addresses. Within AWS, simplify encrypted communication between instances by using Amazon EC2 instances that support offloading encryption to Nitro System hardware.

Restrict access to integration endpoints to allow-listed IP addresses and through secure VPN tunnels when possible. Employ [AWS WAF](#) to protect RESTful APIs, and the suite of [AWS security services](#) to add additional security measures.

HCL_PERF4. How do you determine the volume of data you need to exchange, and can it all be exchanged with one approach?

Some use-cases require large historical datasets, like training machine learning models. Exchanging large datasets may place an unacceptable load on the integration APIs of source systems. Quantify how much data you must exchange across the phases of an interoperability project. If necessary, perform an initial exchange of records in bulk format (for example, by using [AWS Transfer Family](#)), and then use on-going real-time integration methods for updates.

HCL_PERF5. How do you improve availability of integration endpoints or APIs?

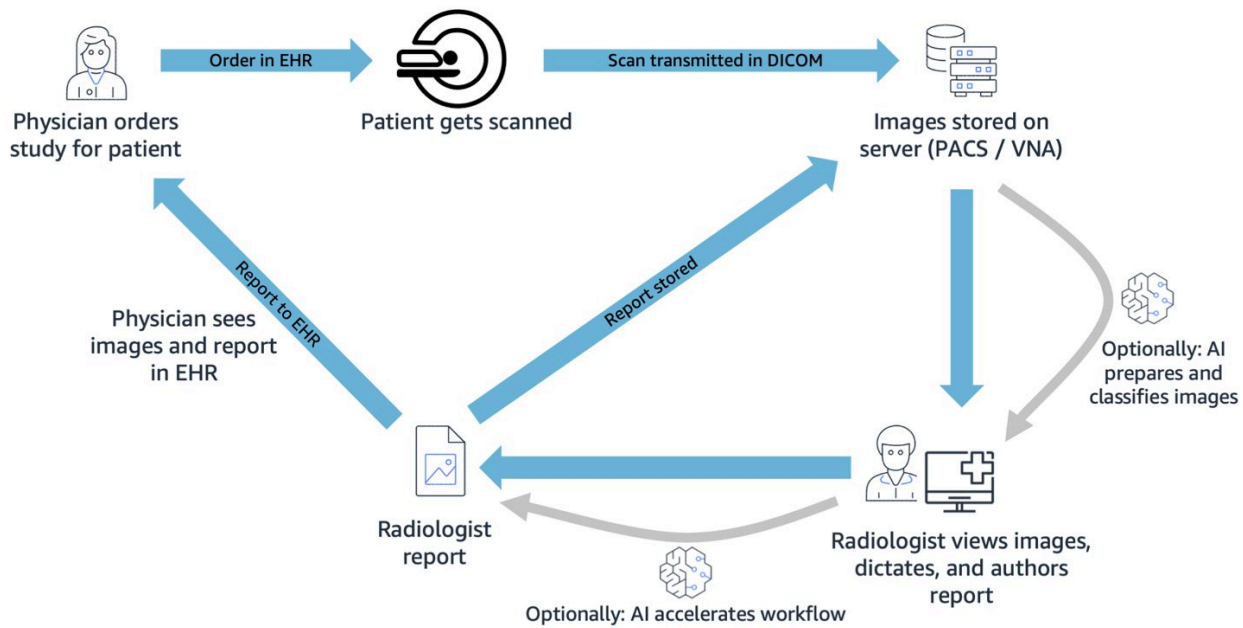
Use Network Load Balancer to route HL7 v2 traffic to your integration endpoints. Network Load Balancer can also monitor the health of your integration endpoints and only route traffic to healthy targets. Deploy your integration services in multiple Availability Zones and enable multiple Availability Zones for the load balancer to increase fault tolerance. APIs hosted by API Gateway are automatically resilient by using multiple Availability Zones in the deployed Region.

Medical imaging

Medical imaging in healthcare spans diagnostic medical imaging, digital pathology, dental imaging, and related applications. Medical imaging systems enable workflows anchored in the generation, storage, and analysis of medical imaging study data. The data is commonly generated by medical imaging hardware, like [CT scanners](#), magnetic resonance imaging ([MRI](#)) machines, and [ultrasound](#) devices, and may be stored in picture archiving and communication systems (PACS) or vendor neutral archives (VNA). There is also secondary use of medical imaging data for non-clinical workloads, such as medical research or development of new medical devices.

Characteristics of medical imaging architectures include:

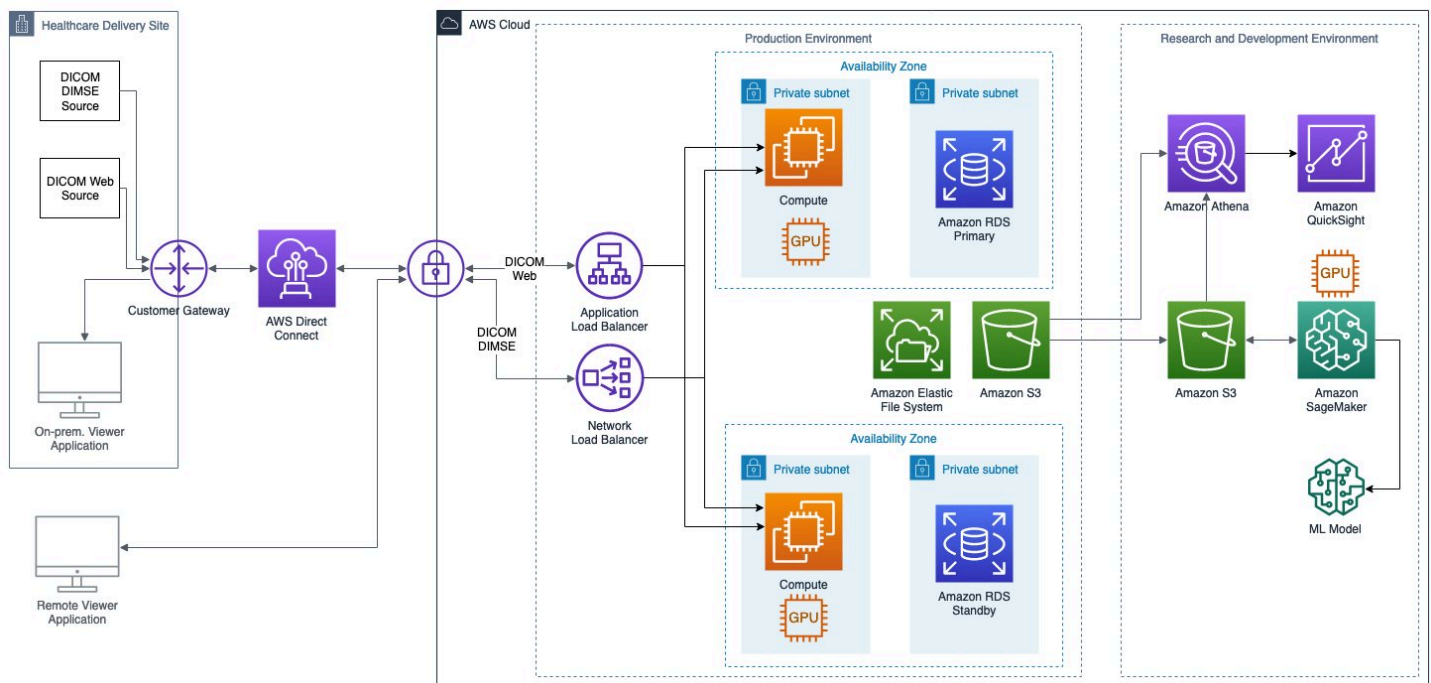
- Imaging study data is generated by on-premises modalities, durably stored by PACS and VNA solutions, and maintained highly available for immediate retrieval by radiologists and other end users. Solution end users, like radiologists, likely expect low latency retrievals of new and old imaging studies.
- An imaging study will likely be accessed several times shortly after ingestion, and tends to be accessed less frequently with age. Regulatory requirements may dictate that studies be retained for several years, and operators may decide to keep studies in perpetuity.
- The imaging scanners, physicians, technicians, and radiologists may be located at the same physical site. Or, cross-site collaboration and outsourced radiology services may extend workflows across multiple care settings.
- Medical imaging systems commonly interoperate with other provider systems, like EHRs, though data integrations.
- Medical imaging study data is often stored and transmitted using the [DICOM standard](#). This standard covers both a network protocol and file format definitions.
- Some medical imaging solutions, such as PACS, may fall under regulatory control, such as [GxP regulation](#) in the US. Machine learning algorithms that aid in diagnoses may also fall under regulatory control. Determine which, if any, regulatory controls are applicable to your solutions.
- Data processing, provider workflows, and end user interfaces may leverage AI to improve care quality and boost productivity. Example applications include computer vision applied to detect disease in medical images, and NLP applied to support report authoring.
- A high-level overview of the medical imaging workflow is shown in figure 3.



A representative medical imaging workflow.

Medical imaging system reference architecture

This section describes key aspects of medical imaging systems, such as PACS and VNA solutions.



A cloud-based medical imaging system, such as a PACS or VNA, on AWS. High availability and low-latency study retrieval for medical imaging solutions.

- The solution should be highly available and deployed across multiple AWS Availability Zones.
- Medical imaging systems often consist of front-end viewers, application servers, databases, and storage for the imaging data. Where possible, each tier of the solution should be able to auto scale independently. Containerization or serverless can simplify operations. Auto Scaling based on load provides performance during peak demand and minimize costs during periods of low demand.
- Images may be programmatically retrieved using the DICOM Message Service Element (DIMSE) or DICOMweb protocol. A Network Load Balancer may be used to route traffic on ports used by DIMSE.
- End users likely demand low latency for retrieving and displaying medical images. Consequently, the data must be durably stored and highly available for immediate retrieval. Users may expect immediate retrieval of medical images that are several years old.
- Recently ingested studies may be cached on [SAN storage](#), EBS volumes, or high-performance file systems like Amazon FSx. Cost optimized solutions provision the minimum volume sizes needed to meet performance requirements, and maximize the use of cost-effective object storage like Amazon S3.

- Medical image data tends to be accessed less frequently as it ages, so newly ingested data should land on Amazon S3 Standard, and then move to lower-cost [tiers](#), such as Amazon S3 Glacier Instant Retrieval, as access frequency declines over time. Amazon S3 Intelligent-Tiering can automatically move data to the most cost-effective access tier based on access frequency.
- Metadata for medical image objects and associated clinical data is commonly stored in a database. These databases may require high-performance storage for the requisite latency and IOPS. In-memory caches, like ElastiCache, may also be used to improve performance. Leverage fully managed database services to attain high availability with minimal operational complexity.
- The data acquired by some medical imaging scanners — like MRI, CT, C-Arm — must be processed in *image reconstruction* to yield readable images. Image reconstruction can be thought of as a [high performance computing \(HPC\) workload](#). Cloud based compute provides elasticity, reducing the time required to perform image reconstruction for emergency procedures.
- Front-end viewers can leverage protocols like [HTTP/2](#) to minimize image download times. Applications may also pre-fetch, cache, or prioritize transmitting the images that are likely to be opened by the end user.
- On-premises caches can provide low-latency hot storage. [AWS Local Zones](#) and [AWS Outposts](#) may help meet hybrid architecture, latency, and data sovereignty concerns.
- Redundant network connections between care settings and cloud services are recommended when a loss of connectivity can impact patient health. [AWS Direct Connect](#) should be used by customer sites with high study volume. Hybrid architectures may help meet stringent latency and business continuity requirements.
- [Data lakes](#) can enable both operations and research and development. Datasets for the development of machine learning algorithms and AI features can be stored in data lakes. [AWS AI services](#) and [Amazon SageMaker](#) can help ISVs rapidly develop AI-based features drawing from a data lake. SageMaker Ground Truth can streamline the process of labeling data for model training.

Questions

HCL_PERF6. How does your organization benchmark the performance of a medical imaging solution?

Quantitatively benchmark the performance of systems in retrieving, analyzing, and reporting on images to meet customer expectations. Tests should realistically capture the network bandwidth, annual volume of data ingested (estimate the types and number of imaging modalities generating data), and the number of concurrent users that the solution will support.

HCL_PERF7. Does your organization perform tests to quantify medical imaging system performance and quantify end user experience under realistic conditions?

Collect quantitative performance metrics under simulated loads and representative network speeds to provide an acceptable end user experience. Test your solution under accurately simulated loads (such as number of concurrent users and realistic last mile network bandwidth). Radiologists demand high performance from PACS and VNA solutions, and are sensitive to image retrieval latencies. For example, collect metrics like time to first image display that quantify how long it takes for an end user to see a medical image after requesting it.

Once performance requirements are met, systems can be cost optimized by balancing image retrieval performance with use of cost-effective storage services (such as right sizing cache volumes and [throughput provisioning](#)).

HCL_PERF8. Does your organization leverage high-performance network protocols and compression of data in transit for medical imaging systems?

Optimize the throughput of medical image data between backends and viewer applications with high-performance, parallelized network protocols like HTTP/2. Also, leverage compression algorithms to reduce the volume of data transferred.

HCL_PERF9. How do you optimize end user experience with algorithms that prioritize the sequence of data transmitted from the backends to front ends?

Medical imaging study data may consist of multiple images or a single large image with regions of varying importance. End user experience can be optimized by prioritizing transmission of the

images or [areas](#) that will be of highest initial interest. In this way, the end user can begin their work while data of less interest is transmitted.

HCL_COST2. How does your organization determine the appropriate storage medium to collect, process, and store medical images?

Quantify the overall user-experience using metrics (like first image display time). Provision cloud services that are appropriate for each component's performance requirements. For example, high IOPS EBS volumes may increase cost but may not improve user experience if the overall solution performance is limited by network connectivity.

Further, collect data on medical image access frequency and use it to optimize the cost of storage through each image's lifecycle. As noted above, medical images tend to be accessed frequently when created and then see less access with time. Decrease storage costs by moving images to lower-cost cloud storage tiers as access becomes less frequent.

Healthcare analytics

Healthcare delivery systems, payors, and service providers use analytics for a range of purposes, such as revenue cycle management, quality management, and process improvement. These entities generate, analyze, and exchange large volumes of data. The health data processed spans a diverse set of domains, including clinical, finance, supply chain, human resources, research and more. The volume and variety of data processed by analytics is steadily increasing, making extensible, elastic, cloud-based architectures increasingly attractive.

Healthcare analytics architectures have the following characteristics:

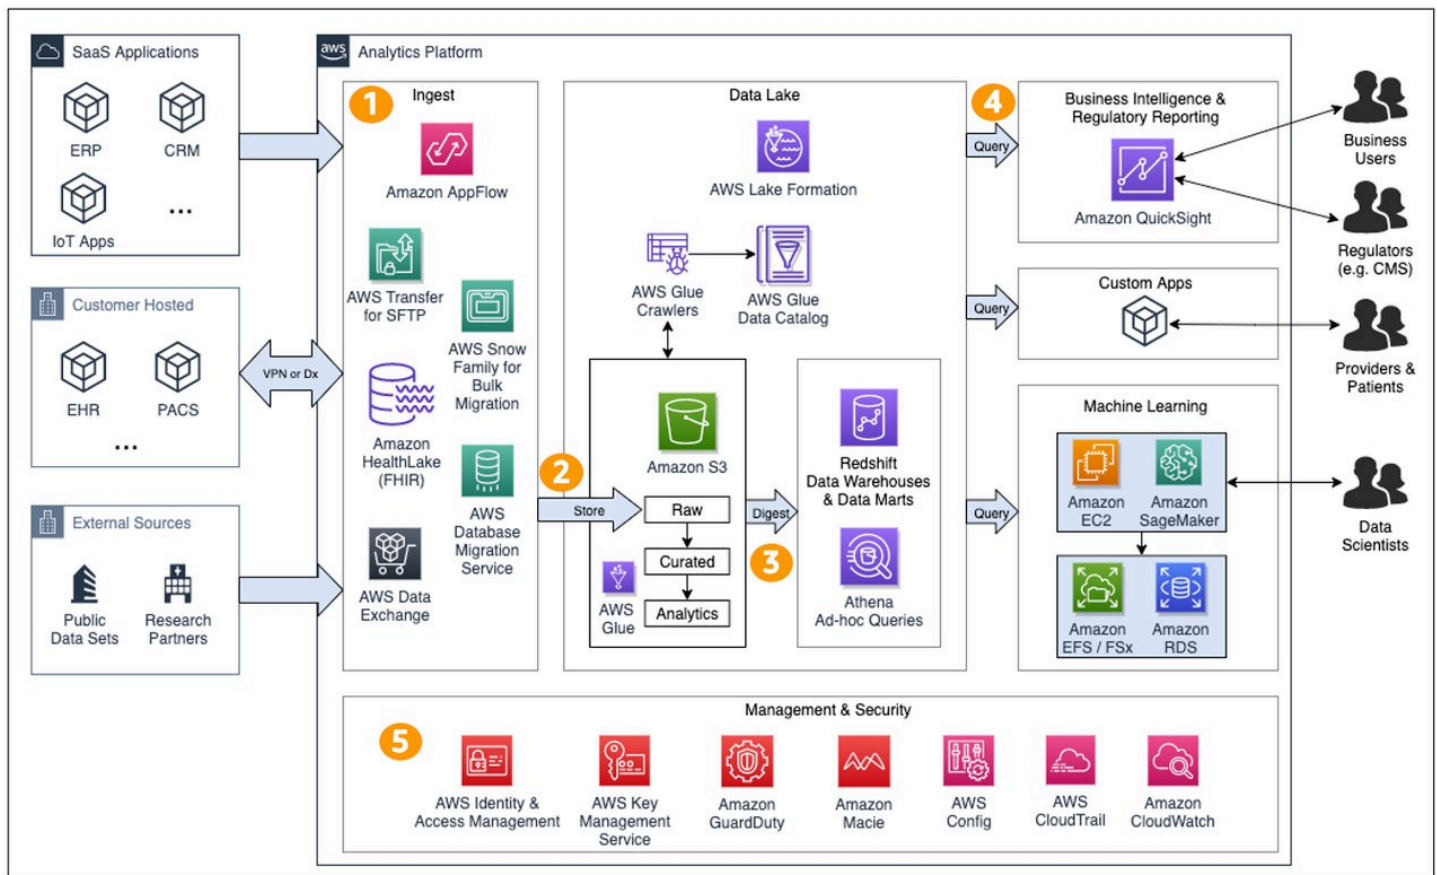
- Source data is ingested from upstream sources such as EHRs using bulk and streaming protocols. The raw data is oftentimes persisted for data lineage and reprocessing. For example, raw HL7 v2 messages, EDI transactions for claims, genomic variant data, medical images, faxed or scanned documents, and so on are ingested and stored for downstream processing, oftentimes for multiple, disparate use-cases (such as input data for AI/ML inferencing, dashboards, forecasting, and sharing).
- Data transformations create derived datasets by applying structural and semantic transforms and linking data from different tables, systems, or domains. During this process, the data can

be mapped to standard clinical terminologies (such as SNOMED, ICD, CPT, and NDC) to provide consistency for downstream consumers of the data. Transformation jobs may include machine learning steps such as OCR of scanned or faxed documents, transcription of voice recording, or natural language processing of clinical text. Additionally, these transformations may include steps to reduce the fidelity of the data for specific use cases, such as generating de-identified or limited datasets for research use cases.

- Business logic runs on the derived data to generate actionable insights. Late binding approaches for the data schema enable a more agile approach and avoids large, upfront investments in heavy ETL. Results are sent to one or more purpose-built data stores.
- Various stakeholders consume the data to glean insights. Data consumers are often diverse, ranging from non-technical reviewers of business intelligence dashboards, to researchers running deep learning algorithms.
- Data remains encrypted at-rest and in-transit throughout the entire process. Identity and access controls are enforced so that access to sensitive health data is limited appropriately. All data assets can be recorded in a centralized Data Catalog to promote data discovery and reuse.

Healthcare analytics reference architecture

This section covers a reference implementation of a healthcare analytics platform using native AWS services. Refer to the [Architecture Best Practices for Analytics and Big Data](#) to browse best practices for data management and analytics. The components in this architecture are building blocks that can be used as-is or substituted with third party components to meet business requirements.



A representative healthcare analytics environment.

- The analytics platform must support the wide variety of communication protocols used by healthcare systems including bulk data feeds and real-time data streams. Examples include bulk data transfers using secure FTP, HL7v2 over MLLP and standard FHIR web services. Legacy protocols that don't support encryption must run over an encrypted channel such as a Site-to-Site VPN.
- Store raw data in a durable, highly available, and secure object store such as Amazon S3. Enable default encryption to verify that all objects are encrypted at rest. Lifecycle policies can be set up to reduce costs based on your access requirements. Many AWS and third party services provide direct integrations with Amazon S3 for data integration and backup. AWS Lake Formation provides a framework to organize and secure the data within the Amazon S3 data lake.
- For high volume message ingestion, batch messages through services such as Amazon Kinesis to reduce the number of actions taken to store the data. This can reduce the overall cost of data ingestion. Prevent data integrity issues by ensuring the batching process aligns with the requirements of the data pipeline.

- Use AWS Glue Crawlers to automatically discover and catalog schemas for the raw datasets. AWS Glue ETL processing workflows transform and normalize the data through serverless and horizontally scalable jobs. Track data lineage to establish traceability and reproducibility for compliance. Use Amazon Redshift for data warehousing and Amazon Athena for SQL queries against cataloged datasets.
- End users interact with the data and insights across all the normalized healthcare data through a number of ways. For example:
 - Business users and regulators perform analysis, view dashboards, and receive reports using business intelligence tools like Amazon QuickSight.
 - Custom application integrations use the data to surface insights to end users, including to the point of care. Data can be accessed using a variety of AWS services such as Lambda functions, containers running in Amazon ECS, Amazon EKS, or AWS AppSync. Verify that the AWS services being used are eligible for the healthcare compliance framework applicable to your workload (such as the [HIPAA Eligible AWS Services](#)).
 - Machine learning (ML) experts can pull standardized datasets and combine them with datasets using custom data preparation processes.
- Use IAM and Lake Formation to narrowly scope permissions. Access controls should be enforced across all AWS environments. Use Amazon CloudWatch to monitor your solution's metrics, logs, and alarms. Use AWS CloudTrail to monitor access to AWS APIs along with GuardDuty to alert on unusual activity. Use Amazon Simple Notification Service (SNS) for sending notifications to on-call engineers and other data consumers. Amazon Macie can automatically discover and categorize sensitive data such as personally identifiable information (PII) and protected health information (PHI). An audit log must be used to capture all sensitive data access (create, read, update, and delete) for regulatory compliance purposes.

Questions

HCL_OE1. Does your organization use master data management processes to unambiguously identify patients or individuals and link clinical concepts and measures to standard healthcare vocabularies (such as SNOMED, ICD, CPT, or NDC)?

Datasets should be cleansed, tracked, and versioned to maintain their quality. Use a central Data Catalog to register and discover these datasets. Automate master data management processes to

reduce burden and improve adoption. Label sensitive datasets with authorization enforced. Track licenses for applicable datasets. Standard healthcare ontologies are preferred over proprietary to promote reuse and data sharing.

HCL_OE2. Does your organization have a process for updating the vocabularies used in master data management processes?

Various public and private organizations maintain healthcare ontologies. The frequency of publication and methods of distribution vary widely. New concepts are added to these ontologies over time and others are deprecated. Automated processes for keeping the content updated reduce manual error-prone efforts and improve data quality. The publication of new content may be used to trigger this update process; otherwise, set up an automated schedule aligned with the content owner.

HCL_OE3. Are open data formats being used for health data?

Health data lakes should store a copy of data in open standard formats (such as HL7, DICOM, standard genomic files like Binary Alignment Map (BAM), Compressed Reference-oriented Alignment Map (CRAM) and (g)VCF, XML, CSV, JSON, Parquet, or JPEG). Derivative copies of data may use non-standard formats to support downstream analytics. Storing health data in open formats expands the software that can be used to process the data, and may decrease the data transformations needed for interoperability.

HCL_SEC13. Does your organization employ role-based and attribute-based access control, as well as fine-grained secure data access at the table, column, or row level to ensure least privilege access?

Healthcare data access should follow least privilege. Access to the data should be granted to only those individuals who need it. Maintaining this access is greatly simplified using role-based and attribute-based mechanisms built on a common identity provider. Likewise, revoke access when no longer needed. Control access to individual datasets, but also verify that the user's ability to combine datasets does not expose unintended risks.

HCL_SEC14. Are comprehensive audit logs capturing all data access (create, read, update, and delete) and show compliance with centrally defined policies?

Audit logs for access to all protected information must be maintained in a central immutable data store. Lock down permissions to audit logs. The logs must be maintained in accordance with your regulatory compliance. The logs can be migrated to lower-cost storage tiers automatically over time to reduce costs.

HCL_SEC15. Is sensitive data (for example, PII or health data) being deidentified or redacted when possible?

Many healthcare analytics use cases don't require personally identifiable information. By deidentifying or redacting data, storing and accessing data becomes less risky. Use irreversible de-identification mechanisms to scrub sensitive information if it's not needed by downstream consumers of the data.

Machine learning for healthcare

Artificial intelligence/machine learning (AI/ML) is being applied to a growing set of problems across healthcare, such as prioritizing treatments, predicting health outcomes, guiding provider workflows, and streamlining revenue cycle operations. A key strength of AI/ML technology is the ability to continually learn from real-world data and improve performance over time. However, healthcare applications of AI/ML pose unique problems, including regulatory oversight, design control obligations, and interpretability requirements imposed by stakeholders. Machine learning development is often performed in concert with traditional analytics and can leverage elements of the infrastructure described in the [Healthcare analytics](#) scenario.

Characteristics of machine learning for healthcare architectures include:

- Data for training models is commonly extracted from production systems within payer and provider organizations, such as EHRs, medical imaging systems, claims and revenue cycle solutions, scanned or faxed documents, biobanks, and genomics data stores. Healthcare tends to be high dimensional, multi-modal, and often suffers from missingness due to the episodic nature of care resource consumption.

- Data lakes are well suited to landing, preparing, and storing health datasets for ML. The traditional data silos of healthcare data can hamper training models that perform well. Extracting data from multiple systems and data modalities can improve model performance and transferability between settings.
- Healthcare analytics tools may be used for exploratory data analysis before machine learning development.
- Data is commonly extracted in bulk, cleaned, and prepared for use training models. Data for inferencing may also be processed in batches, or in real-time using streaming data integrations such as HL7 v2, FHIR, or other interoperability standards. Data lineage should be tracked, providing a map of various data sources and transformation steps that the data flows through.
- Organizations may use identifiable health data for ML development. In such cases, organizations must adhere to the principle of least access, protecting health data from inappropriate access while enabling data scientists to run ML workflows. De-identifying health data before making it accessible to ML teams can mitigate privacy and security concerns.
- Federated Learning (FL) may be adopted to facilitate multi-centric, collaborative ML modeling when data privacy and anonymity is required. Federated paradigm has been proven to be domain-agnostic and framework independent for distributed modeling training. The heterogeneous datasets from diverse sources and patient populations are desired to build a robust and accurate ML model. To overcome the privacy concern and complex data sharing process, FL in healthcare can be used to rapidly and securely develop new ML models without data ownership hurdles.
- Stakeholders may require explainability and repeatability of models employed in healthcare, especially in care delivery and revenue cycle use cases. Validating a model may require that structure and output of the model have *face validity* and align with medical knowledge.
- Thresholds of acceptable model performance may follow from the risks and benefits of the outcome modeled. For example, if the outcome being modeled is high risk to patients or high cost, then stakeholders may demand that models display high [precision](#) (confidence in a positive prediction). Alternatively, if the model is used for a low risk, low cost, yet high benefit, stakeholders may prioritize higher [recall](#) from models.
- Model deployments should be automated with pipelines to minimize human touchpoints, deploy data integration workloads consistently and repeatedly, and formalize how code is promoted from development to production.
- Healthcare AI/ML models often need to be integrated in the workflows of healthcare providers, patients, and other actors to have utility. This may require integrating systems (such as an EHR),

updating workflows, and retraining providers. It may also require integration with medical equipment and require regulatory oversight.

- Model performance should be monitored after deployment in production. The complexity and variability of health data makes it especially important to monitor the health of inferencing pipelines.

Machine learning resources

Refer to the following resources to learn more about our best practices for machine learning.

Documentation and blogs

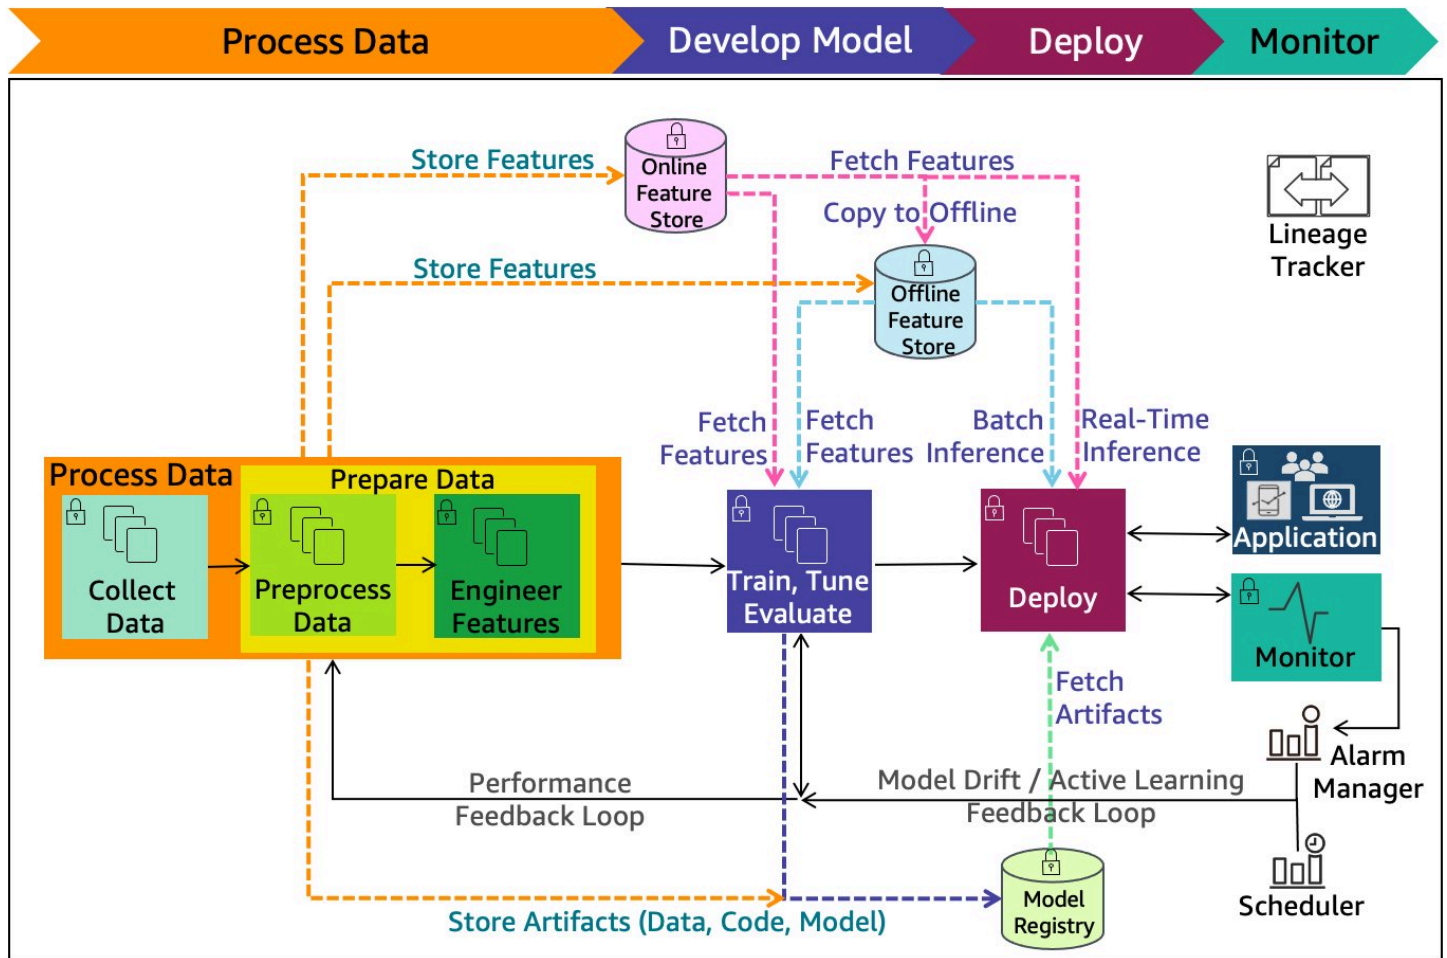
- [AWS Config and Best Practices for AI and ML](#)
- [Model Explainability with AWS Artificial Intelligence and Machine Learning Solutions](#)
- [Detect machine learning \(ML\) model drift in production \(video\)](#)

Whitepapers

- [Machine Learning Best Practices in Healthcare and Life Sciences](#)
- [Model Explainability with AWS Artificial Intelligence and Machine Learning Solutions](#)

Machine learning reference architecture

This section depicts a typical machine learning lifecycle and data flow.



ML lifecycle with detailed phases and expanded components

The following steps detail the end-to-end data flow for machine learning:

1. Data is collected and pre-processed using a data lake, as described in the preceding *healthcare analytics* scenario.
2. Features representing clinically valid events, concepts, and processes of care are extracted from raw data and stored in feature stores for model training.
3. The ground truth of data labels is populated and reviewed by humans, which can be used to build supervised classification or regression models.
4. Standard ML training, tuning, and evaluation workflows are used to develop models.
5. Models are reviewed by cross-functional stakeholders, such as clinical leaders and regulatory reviewers. Models are evaluated based on performance and explainability requirements.
6. Accepted models may be integrated with IT systems used for care delivery, such as EHRs and medical devices.

7. Model inferences are incorporated in clinical workflows. Providers may be trained on how to use the models as they deliver care.
8. Model inferencing pipelines are monitored, and the performance of deployed models is periodically checked.

Questions

HCL_OE4: How do you track model revisions and ensure traceability of your ML artifacts?

Employ version control for source code, data, and ML artifacts to ensure traceability and reliability of production ML deployments. Version control and traceability may be required by applicable regulatory frameworks, if for example models are deployed in support of medical devices.

HCL_SEC16. How does your organization review, accept, and manage the licenses of open-source software dependencies?

Data science in healthcare often depends on open-source libraries for data processing, model development, training, and hosting. Establish a process to review the privacy and license agreements for all software and ML libraries needed throughout the ML lifecycle. Verify that these agreements comply with your organization's legal, privacy, and security requirements.

HCL_SEC17. Does your organization deidentify health data used for machine learning, or otherwise limit access to sensitive, identifiable health data?

Many ML workflows do not require identified health data. Applying ML to deidentified data is one way to develop AI-powered applications without compromising privacy or data security. Cloud services like the [Amazon Comprehend Medical DetectPHI API](#) can streamline generating deidentified datasets.

HCL_REL4: How does your organization identify and limit biases in training data and statistical models?

Statistical models trained on real-world health data are susceptible to biases. Health data may inadvertently be collected from populations of individuals with similar characteristics, such as median household income, social determinants of health, and access to care. Care setting and health insurance coverage may also impart biases. For example, treatment cohorts may exhibit higher household income because such individuals may have greater access to advanced care.

Trained models may be misleading if biases are not quantified and mitigated. Also, models may be inaccurate when trained on biased data and applied to settings with different distributions. Examine data distributions and [perform analyses](#) to quantify and mitigate biases before training models.

HCL_PERF10: What processes do you use to monitor model performance after deployment and protect against drift?

Health data is often complex, and subject to temporal variations in quality and concept expression. Model performance may degrade over time due to data quality, model quality, and concept drift. Create a baseline for data quality, and [automate monitoring performance](#) in production. Automate alerts for changes in data quality or distributions, such as age deciles and prevalence of relevant chronic diseases. [SageMaker Model Monitor](#) provides an end-to-end framework model monitoring and lifecycle management.

Conclusion

The AWS Well-Architected Framework provides architectural best practices across the pillars for designing and operating reliable, secure, efficient, and cost-effective healthcare workloads in the cloud. The Framework provides a set of questions that allows you to review an existing or proposed architecture, and also a set of AWS best practices for each pillar. Using the Framework in your architecture will help you produce stable and efficient systems, which allows you to focus on your functional requirements.

Contributors

Healthcare Industry Lens authors and contributors (in alphabetical order):

- Sanford Coker – Principal Solutions Architect, WW HCLS – Private Equity, Amazon Web Services
- Conor Colgan – Senior Startup Solutions Architect, Healthcare, Amazon Web Services
- Patrick Combes – Head of Solutions Strategy & Technology, Amazon Web Services
- Kevin Cox – Senior Product Solutions Architect – AWS Industry Products Engineering, Amazon Web Services
- Aaron Friedman – Principal Product Manager, HealthAI, Amazon Web Services
- Razvan Ionasec - Healthcare Technical Leader, Europe, Middle East, and Africa, Amazon Web Services
- Nelson Lee – Senior Healthcare Technology Strategist, Healthcare, Amazon Web Services
- Bakha Nurzhanov – Interoperability Solutions Architect, Healthcare, Amazon Web Services
- Dmitry Pavlov – Senior Solutions Architect, Healthcare, Amazon Web Services
- Charlie Peterson – Principal Program Manager, WWPS Tech Team, Amazon Web Services
- Ujjwal Ratan - Principal AI/ML Solutions Architect, Healthcare & Life Sciences, Amazon Web Services
- Bruce Ross – Well-Architected Lens Leader, Amazon Web Services
- Andy Schuetz – Principal Product Manager, HealthAI, Amazon Web Services
- Justin Stanley – Principal Solutions Architect, Healthcare, Amazon Web Services
- Donny Wilson – Security Solutions Architect, Healthcare, Amazon Web Services

Document revisions

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Initial publication	Healthcare Industry Lens published.	November 17, 2022

Notices

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

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

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.