



Healthcare's Guide to AI: Security, Compliance, and Governance in the Cloud

→ Contents

Executive summary	02
What's at Stake: AI Risks Healthcare Leaders Are Managing Now	03
AI Readiness: Getting Started and Leveraging AI in the Cloud	04
Amazon Web Services	04
Microsoft Azure	04
Google Cloud Platform (GCP)	04
The Multi-Cloud	05
Governing AI in the Cloud	05
How ClearDATA Helps Accelerate AI Innovation	06
Compliance Reference Architectures	06
Architectural Services: Designing AI-Ready Cloud Environments	06
Security Risk Assessment (SRA)	06
Cloud security posture management platform: Your Control Plane for Compliant AI	06
Cloud-Native Managed Detection and Response (MDR)	07
Accelerate AI Innovation with HITRUST Inheritance	07
How ClearDATA Supports AI Innovation in the Healthcare Cloud	08
Amazon Web Services (AWS)	08
Microsoft Azure	08
Google Cloud Platform (GCP)	08
Build Your Blueprint for Secure AI	09
Partnering for Continuous AI Governance	10
Stay compliant. Stay Secure. Stay in Control.	11

Executive Summary

Artificial intelligence (AI) has rapidly shifted from a conceptual advantage to an operational baseline in healthcare, offering a massive opportunity for business acceleration. For CIOs, CTOs, and InfoSec leaders, the primary challenge isn't building AI models, but operating them securely, compliantly, and under strict governance.

As your teams integrate AI into healthcare workflows for clinical decision support, administrative automation, and predictive analytics, IT and security leaders are left balancing this urgent adoption with stringent regulatory compliance and evolving security threats.

Traditional security controls were not designed for the dynamic nature of AI, leaving many organizations and their sensitive data exposed. This guide provides a framework for architecting, securing, and governing AI workloads in the cloud to help you innovate safely.



What's at Stake: AI Risks Healthcare Leaders Are Managing Now

Before deploying AI at scale, you must familiarize yourself with, and address the unique attack surfaces and vulnerabilities it introduces:

- **Shadow AI & Ungoverned Tool Sprawl:**

[Employees are adopting unapproved AI tools](#)—from consumer LLMs to rapidly proliferating point solutions. This results in Protected Health Information (PHI) flowing outside of IT visibility, policy enforcement, and audit trails. The landscape changes faster than traditional governance can manage.

- **Insecure Third-Party Integrations:**

AI agents and third-party tools linked to PHI data create hidden exposure pathways that are difficult to detect and govern.

- **AI Infrastructure Complexity:**

Configuring compute clusters, storage, and data pipelines in the cloud introduces massive variability. [Misconfiguration](#) risks compound exponentially as workloads scale.

- **Multi-Cloud Sprawl & Configuration Risk:**

AI adoption frequently accelerates [multi-cloud usage](#). Each platform introduces unique configuration requirements, security controls, and operational risks, expanding the surface area you must monitor.

- **Identity & Access Exposure:**

AI workloads introduce new identities, APIs, and service accounts. This expands the attack surface and increases the risk of privilege escalation, often without clear ownership.

- **AI Infrastructure as an Attack Target:**

Compute clusters, data pipelines, and model storage are increasingly targeted by threat actors. While the attack patterns differ from traditional threats, the operational [and regulatory consequences are just as severe.](#)



AI Readiness: Getting Started in the Cloud

AI is rapidly becoming foundational to healthcare, but AI workloads introduce new layers of risk across data, identity, and infrastructure. The cloud environments powering these workloads carry the same regulatory obligations as traditional healthcare systems but with significantly higher operational complexity

Designing an environment that leverages the best native AI tools without falling into single-vendor lock-in is critical. Each cloud provider has a different starting point for healthcare AI, and understanding these nuances is essential for your architecture. Below are examples for each of the major cloud providers and their AI-specific services.

Amazon Web Services (AWS)

AWS is highly mature regarding healthcare compliance documentation, making it a strong choice for regulated workloads.


- **Key Services:**

[Amazon Bedrock](#) provides managed foundation models with built-in Guardrails for PHI filtering and prompt security. AWS [HealthLake](#) acts as a FHIR-native, HIPAA-eligible data store that feeds AI pipelines cleanly. [Amazon Nova](#), the company's proprietary portfolio of generative AI models, is accessible through Amazon Bedrock. It offers several applications for the healthcare industry, including enhancing patient communication and rapidly identifying high-risk findings across millions of clinical records securely.

- **Architectural Consideration:**

Use [AWS PrivateLink](#) to keep all AI traffic off the public internet. AWS's massive breadth can lead to configuration inconsistencies across accounts if guardrails are not continuously managed.

Pause here to hear about [Innovating Securely and Compliantly with AI on AWS](#)



Google Cloud Platform (GCP)

GCP probably features the deepest ML/AI research lineage but requires more hands-on compliance configuration. It often attracts mature data science teams seeking maximum flexibility.

- **Key Services:**

Gemini Enterprise Agent Platform is the unified platform for building and managing ML models. The Cloud Healthcare API handles FHIR and DICOM exceptionally well, making it ideal for imaging AI use cases.

- **Architectural Consideration:**

GCP requires rigorous tagging, logging, and IAM configurations across multiple services to maintain a secure pipeline.

Microsoft Azure AI Services to Foundry Tools

Azure's primary strength is its deep integration with Microsoft's existing compliance stack.

- **Key Services:**

Foundry tools run securely within your own tenant. Microsoft signs a Business Associate Agreement (BAA) covering it, and your PHI does not train the underlying public model.

- **Architectural Consideration:**

Azure Health Data Services is the optimal staging ground for clinical data before it touches your AI workloads, acting as a seamless extension for organizations already running Microsoft 365 or Teams in healthcare settings.

AI and The Multi-Cloud Reality

Most healthcare organizations don't intentionally choose [multi-cloud](#); it chooses them through acquisitions or best-of-breed application requirements. The architectural challenge immediately becomes enforcing consistent security controls and data governance across varied environments.

For years, many technology leaders remained committed to a single-cloud strategy. It was easier to manage, more familiar to internal teams, and often viewed as the more secure option. However, as discussed above, different AI tools and services are frequently optimized for specific cloud providers, prompting organizations to expand beyond a single platform.

To maximize outcomes, healthcare organizations including pharmaceutical companies, HealthTech innovators, payers, and life sciences organizations—are increasingly adopting the best tool for each use case by combining the capabilities of multiple cloud environments. Rather than choosing one platform over another, organizations are embracing an “and” strategy instead of an “or” strategy, leveraging the unique strengths of multiple cloud providers to better support business objectives and innovation initiatives.

As AI tools become more widely adopted across departments often without centralized governance or oversight, the risk of [“Shadow AI”](#) continues to grow. Similar to Shadow IT, Shadow AI occurs when employees implement or use technology without the knowledge, approval, or management of the IT organization.



Governing AI in the Cloud


To secure AI, you must first understand the critical difference between traditional Shadow IT and the more immediate threat of Shadow AI.

- **Shadow IT** is a familiar challenge: employees using applications or cloud resources without IT approval. While this creates visibility gaps and potential cost overruns, it is primarily an access control and billing issue.
- **Shadow AI** poses a far more acute and immediate HIPAA exposure risk. This occurs when employees use consumer-grade AI tools—such as public versions of ChatGPT, Gemini, or Copilot—and input Protected Health Information (PHI) to get their work done faster.

Consumer AI tools do not sign [Business Associate Agreements \(BAAs\)](#) with your organization. Therefore, any PHI entered into these platforms constitutes a data disclosure and a potential HIPAA violation. There is no gray area here.

The solution is not as simple as blocking access to consumer AI endpoints. If you block these tools without providing a sanctioned alternative, your users will inevitably find workarounds to access the capabilities they need. The only effective strategy is to provide a well-governed, enterprise-grade AI environment that your employees can leverage to accelerate your business, not stall it.

This will mitigate the risk by addressing the underlying need for advanced AI tools within a secure, compliant framework. Many organizations need support getting started, and cloud environments provide various tools and services to support AI use.



How ClearDATA Helps Accelerate AI Innovation

ClearDATA is healthcare's dedicated cloud security and operations partner. We apply a shared responsibility model, managing the continuous governance work required for your cloud infrastructure to run securely, so your team can focus on innovation.

Compliance Reference Architectures

[ClearDATA's Compliance Reference Architectures \(CRAs\)](#) provide pre-configured cloud environments for AWS, Azure, and GCP that embed security and compliance controls directly where your AI workloads run. We translate complex regulations like HIPAA, [HITRUST](#), and NIST into enforced technical controls, giving your team pre-governed environments for AI data pipelines, storage, and compute.

This approach significantly reduces the risk of misconfiguration during AI deployment and scaling, as the architectures are continuously updated to reflect evolving cloud services and regulatory requirements.

Our CRAs establish a compliant, secure foundation, so your AI workloads are built correctly from day one, not retrofitted after a risk has already been introduced.

Architectural Services: Designing AI-Ready Cloud Environments

Beyond providing pre-built compliant architecture, our team offers advisory and professional architectural services. We partner with healthcare organizations to design cloud environments that can support demanding AI workloads without introducing unmanaged risk.

Our role is to ensure your AI systems are built with the right controls, data boundaries, and integration patterns before they reach production.

Our architectural services include:

- AI workload and data pipeline architecture
- Secure integration of models, data sources, and third-party services
- Alignment to compliance frameworks and operational best practices
- Outcomes backed by CRA patterns and real-world healthcare experience

Security Risk Assessment (SRA)

To confidently adopt AI, you must first understand your organization's current security and compliance posture. Our Security Risk Assessment (SRA) provides an expert-led evaluation of your cloud environment, specifically designed to identify security, compliance, and operational risks, including those introduced by new AI workloads.

- **Comprehensive Framework Coverage:**

We evaluate your environment against key healthcare standards, including the HIPAA Security Rule, NIST, HHS 405(d), and the [HITRUST CSF](#).

- **Multi-Layered Vulnerability Identification:**

Our experts identify vulnerabilities across your infrastructure, identity, and data layers through a combination of automated scanning, penetration testing, and a thorough compliance gap analysis.

- **Comprehensive Framework Coverage:**

We evaluate your environment against key healthcare standards, including the HIPAA Security Rule, NIST, HHS 405(d), and HITRUST CSF.

Cloud Security Posture Management Platform: Your Control Plane for Compliant AI

AI workloads are only as secure as the cloud environments they run in. ClearDATA's proprietary cloud security posture management (CSPM), the [CyberHealth™ Platform](#), serves as the control plane for healthcare AI environments, delivering continuous visibility, governance, and risk reduction across your cloud infrastructure. As AI workloads evolve, the platform helps ensure they remain secure, compliant, and aligned with organizational and regulatory policies.

Key capabilities include:

- 400+ healthcare-specific safeguards mapped to HIPAA, HITRUST, NIST, and other compliance frameworks
- Continuous monitoring of cloud infrastructure, identities, configurations, and data activity



- **Automated detection** of misconfigurations, policy violations, and compliance drift
- **Integrated remediation workflows** that help resolve issues quickly and consistently

With continuous oversight and automated governance, organizations can innovate with AI confidently while maintaining a strong security and compliance posture.

Cloud-Native Managed Detection and Response (MDR)

Beyond automated safeguards, effective AI governance requires continuous human oversight. [ClearDATA's Cloud-native MDR](#) protects AI workloads at runtime by detecting abnormal behavior, unauthorized access, and emerging threats before they impact operations or compliance.

Key capabilities include:

- Continuous threat detection across infrastructure, identity, and data layers
- Cloud misconfiguration identification and alerting
- CVE and vulnerability scanning across your entire cloud estate
- Incident response and containment led by dedicated healthcare security experts

Accelerate AI Innovation with HITRUST Inheritance

For healthcare organizations, achieving and maintaining HITRUST certification is essential, but often time-consuming and resource intensive. ClearDATA's HITRUST r2-certified platform helps streamline the process by [allowing your organization to inherit](#) a significant portion of required controls, reducing the effort needed to demonstrate compliance while accelerating secure innovation.

- Inherit up to 45% of HITRUST controls through the ClearDATA CyberHealth™ Platform, made exclusively to protect PHI and other sensitive healthcare data in the cloud.
- Reduce the scope, cost, and timeline of HITRUST assessments and audits
- Leverage a proven compliance program embedded within ClearDATA's operational model
- Maintain continuous readiness as cloud and AI environments evolve

By inheriting foundational security and compliance controls, healthcare organizations can establish the trusted baseline needed to develop, deploy, and scale AI solutions with greater confidence—without slowing innovation or increasing regulatory risk.



How ClearDATA Supports AI Innovation in the Healthcare Cloud

ClearDATA helps healthcare organizations understand which AI and machine learning services offered by cloud providers are HIPAA-eligible and appropriate for healthcare workloads.

Our experts provide information and guidance on eligible AI services while helping organizations understand the important distinctions between HIPAA-eligible, compliant, and generally available services. This clarity enables teams to make informed decisions when evaluating and deploying AI technologies in regulated healthcare environments.

Amazon Web Services (AWS)

Healthcare organizations need a secure foundation to scale AI innovation while protecting sensitive patient data and meeting compliance requirements. By combining AWS AI services with healthcare-specific security, governance, and operational expertise, organizations can accelerate AI adoption, reduce risk, and build trusted AI environments that support innovation without compromising security or compliance.

Without continuous monitoring, AWS compliance is just a point-in-time snapshot that drifts immediately.

- **Automated Drift Detection:**
We continuously monitor Bedrock endpoints and SageMaker jobs, immediately catching and remediating configuration drift.
- **IAM Right-Sizing:**
We identify and restrict over-permissioned IAM policies that were set up for a Proof of Concept (POC) and improperly left active for production.



Microsoft Azure

ClearDATA helps organizations build and operate trusted AI environments on Microsoft Azure by combining healthcare-specific security, compliance, and cloud operations expertise with proven governance and risk management frameworks

We help clients enforce rigorous network boundaries and data tracking within Microsoft environments.

- **Policy-as-Code for AI:**
We enforce private endpoints, customer-managed keys (CMK), and block public access on all model deployments
- **Data Lineage and Auditability:**
We connect Microsoft Purview to your AI environments, creating a continuous, auditable trail of what data the AI services touched to satisfy auditor requirements.

Google Cloud Platform (GCP)

ClearDATA helps organizations build and scale trusted AI environments on Google Cloud by combining healthcare-specific security, compliance, and cloud operations expertise with proven governance frameworks. This approach helps organizations securely deploy AI workloads, protect sensitive health data, and maintain compliance while accelerating innovation and business growth.

For organizations leveraging GCP's advanced ML tools, we secure the foundational pipeline.

- **Pipeline Governance:**
We enforce rigorous Vertex AI pipeline governance, including automated tagging, logging, and IAM scoping across services.
- **Pre-Flight Configuration:**
We harden the Healthcare API configuration before any clinical data is ingested or processed.



Build Your Blueprint for Secure AI

Whether you are early in your AI journey or already operating workloads in the cloud, follow these six actionable steps to architect a secure, compliant environment:

01

Classify your data before you touch AI. You cannot make accurate architectural decisions without knowing exactly what data you hold and its sensitivity level.

02

Establish the network boundary in infrastructure-as-code. Define boundaries for AI workloads from day one. Mandate private endpoints and disable public IP assignments.

03

Stand up identity and access before services. Implement managed identities and role-based access control (RBAC). Never rely on static access keys for AI workloads.

04

Instrument logging and monitoring. Configure comprehensive audit logging (e.g., AWS CloudTrail, Azure Monitor) before anyone runs a real workload.

05

Formalize your AI governance policy. Document sanctioned tools, allowable data types for AI processing, and the strict approval process for new use cases.

06

Build continuous compliance validation. Configuration drifts are inevitable. Automated, continuous monitoring is the only way to maintain your security posture at scale.



Partnering for Continuous AI Governance

Deploying AI in healthcare presents significant security and compliance challenges. To innovate safely, organizations must adopt a cloud strategy that automates compliance and secures data pipelines. ClearDATA provides technology-enabled managed services that build a secure cloud foundation, offers continuous monitoring, and resolves issues to free up your team.



Gain Clarity with a Security Risk Assessment

Understand your security posture and identify compliance gaps with a comprehensive assessment. Our experts deliver a prioritized, healthcare-specific roadmap, providing a clear path forward for securing your AI initiatives.



Build a Secure Cloud Foundation

Establish an AI-ready and compliant environment from day one. We scope the controls, reference architectures (CRAs), and infrastructure to your specific workloads, ensuring your cloud foundation is built for secure innovation.



Assurance-Oriented Operations

With continuous monitoring and expert management, we provide assurance that your environment remains secure and compliant. Our team handles the infrastructure risk, enabling your team to focus exclusively on driving AI innovation and achieving better outcomes.



TRUSTED BY



Why Leading Healthcare Organizations Choose ClearDATA

ClearDATA is healthcare's leading trusted partner helping organizations turn threat intelligence into a secure and compliant cloud infrastructure for innovation and growth.

Our shared responsibility model unites expert services and technology to safeguard your public cloud, turning security from a reactive measure into a strategic advantage.

Healthcare's Trusted End-to-End Cloud Security, Compliance, and Operations Partner

- ✓ **Cloud Security:** Protect your infrastructure with continuous risk monitoring, threat detection, and incident response to reduce breaches and protect sensitive information in the public cloud.
- ✓ **Continuous Compliance:** Automatically align cloud resources with compliance standards through always-on monitoring, and maintain clear visibility of your compliance status and trends in our CSPM, CyberHealth Platform.
- ✓ **Cloud Operations & Optimization:** We manage the day-to-day operations in your cloud environments, supporting over 200 cloud services with proactive monitoring and expert troubleshooting.

Stay compliant. Stay Secure. Stay in Control.

[Assess Your AI Readiness Today](#)

MKT-0315 April 2026 © ClearDATA.



© 2026 ClearDATA. All Rights Reserved. This material is proprietary and all rights are reserved, including the right to reproduce this form or parts thereof in any form.