INDIUM

Build HIPAA-Ready AI Solutions in 2027

What Healthcare Leaders Must Know

E-BOOK

Table of Contents

03
The new operating reality of AI in healthcare

04
Understanding the compliance stack

07
Trust boundaries around PHI

10
Risk profiles differ among 7 AI capabilities

14
Enterprise standards for healthcare AI governance

16
6 decision pathways for enterprise leaders

17
Build AI solutions for long-term governance

Preface

A Note to Healthcare Leaders

Physician AI use just jumped from 47% to 63%[1] in under a year, according to Doximity's 2026 survey of 3,151 U.S. physicians. 94% are either using AI in practice or actively want to. The adoption wasn't driven by a hospital policy. 80%[2] of respondents found it useful in a professional context and faster than formal guidance could keep up with.

The Health Insurance Portability and Accountability Act (HIPAA) of 1996 remains the legal foundation for protecting Protected Health Information (PHI). It defines who should access patient data, and how it should be protected. Large language models and AI agents retrieve information by connecting to multiple systems and complete tasks with minimal human involvement.

Healthcare leaders now need to make new governance decisions around data access, auditability, accountability, vendor management, and oversight–areas that extend beyond the foundation HIPAA provides.

The pages ahead focus on AI governance through the lens of healthcare and help healthcare leaders make the decisions required to build HIPAA compliant AI.



Chapter 1

The New Operating Reality of AI in Healthcare

AI adoption in healthcare has outpaced the assumptions on which HIPAA was built. Healthcare leaders need to treat **AI governance as a continuous process.**

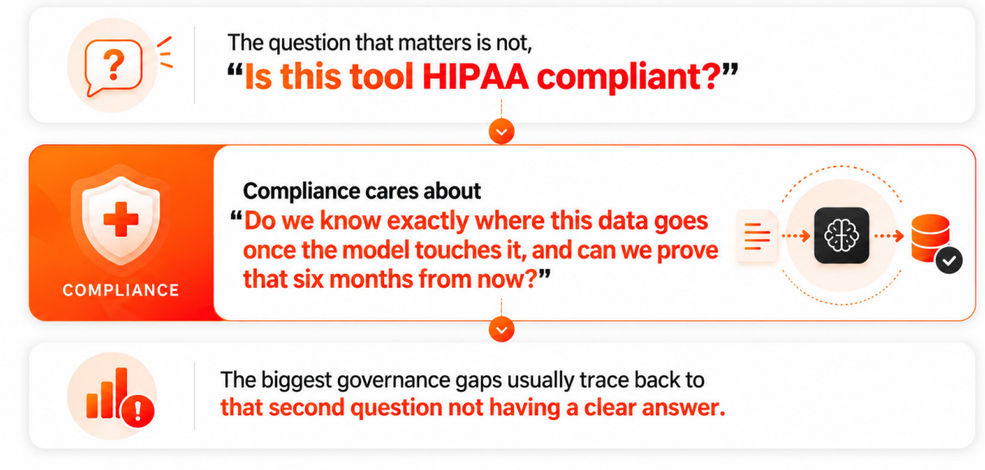
The reality of AI in healthcare shows how quickly AI solutions proved useful, even before formal policy caught up. Clinicians began using AI in their daily workflows before the organization had a clear framework for governing its use.

HIPAA was not written with LLMs or AI agents in mind. It was established in 1996 for a world of paper charts and fax machines, then updated in 2009 and 2013 for the electronic health record era. The rules do not directly address systems that can read a patient chart, decide what information they need, and act across departments in seconds.

Healthcare organizations have a clear picture of where AI can create value. Their attention is now shifting to how it can be deployed responsibly.

AI performance needs regular monitoring to confirm they're still making dependable recommendations and behaving the way they were intended to.

Proactive compliance and AI governance frameworks are replacing the old audit-first mindset. Automated controls are becoming part of day-to-day operations, monitoring AI activity, detecting potential risks early, and preventing PHI from leaving approved environments.



Chapter 2

Understanding the Compliance Stack

Healthcare AI compliance works as a cascade.
Skipping a level in that chain can create compliance gaps.

Federal law sets out broad requirements. HIPAA translates that requirement into specific rules. Those rules translate into security controls. Security controls translate into engineering decisions that shape how an AI system gets deployed.

Building HIPAA compliant AI depends on getting each layer of that chain right.

We have divided compliance into four layers to help you understand how data moves and where AI tools access it.

Layer 1: Privacy

- The privacy rule is fundamentally about controlling access to patient data. AI makes those decisions more granular.
- AI engineers need to define which records a system can access and assign permissions based on the tool's purpose.
- They also need to limit developer access to production data and determine whether patient identifiers can be removed before the model uses the information.

Layer 2: Security

Security rule focuses specifically on protecting electronic PHI, and it organizes that protection into three categories.

Category	What it Covers
Administrative	Policies, staff training, risk assessments, access approvals, incident response, vendor management.
Physical	Server access, office security, laptops, workstations, data center management.
Technical	Encryption, authentication, logging, access control, audit trails, role-based access, data integrity, secure APIs, backups.

Layer 3: Breach Notification

If patient information is exposed, accessed by someone who shouldn't have seen it, or disclosed without permission, the organization can't simply fix the issue and move on.

HIPAA requires them to:

- Investigate what happened
- Document the incident
- Determine the severity
- Notify the right people if the breach meets certain criteria

Depending on how large the breach is, that could mean notifying:

- The affected patients
- The U.S. Department of Health and Human Services (HHS)
- Media for very large breaches

If a breach occurred while an AI tool was actively in use, logs support AI compliance by providing the evidence needed to reconstruct what happened.

For example:

- **Access logs** tell you who or what accessed patient data.
- **Decision logs** show what action the AI took or recommended.
- **Prompt history** helps explain what information the AI received and what the user asked it to do.
- **Version history** identifies the exact version of the model that was running at the time.

Layer 4: Business Associate Agreements

Healthcare companies generally choose not to build AI solutions in-house. External vendors often provide cloud infrastructure, AI models, implementation support, or ongoing operations.

Whenever one of those partners handles Protected Health Information (PHI), a signed Business Associate Agreement (BAA) must already be in place before the data is shared.

Other Regulations That Stack on Top of HIPAA

HIPAA and the Health Information Technology for Economic and Clinical Health (HITECH) Act are only part of the compliance picture.

AI influencing diagnosis or treatment can fall under the U.S. Food and Drug Administration (FDA) oversight as Software as a Medical Device (SaMD). AI used in clinical trials must meet requirements such as Title 21 of the Code of Federal Regulations, Part 11 and Good Practice (GxP) guidelines. If the system processes data belonging to EU patients, General Data Protection Regulation (GDPR) and the EU AI Act may also apply.

An AI governance framework should account for these requirements alongside HIPAA to support compliant AI deployments.

Layer	Regulation	When It Applies
FDA software regulation	U.S. Food and Drug Administration (FDA)	AI functions as Software as a Medical Device (SaMD) or supports diagnosis or treatment decisions.
Nondiscrimination (Section 1557, Affordable Care Act)	HHS Office for Civil Rights (OCR)	AI influences clinical decisions, patient access, or care within covered health programs.
State health and AI laws	State regulators and enforcement agencies	The organization operates in states with healthcare, privacy, or AI laws that extend beyond federal requirements.
International regulations (European Union Artificial Intelligence Act and General Data Protection Regulation)	European Union regulators	The organization processes personal data of individuals in the European Union or provides services within the region.

Chapter 3

Trust Boundaries Around PHI

The principles behind HIPAA haven't changed. The difficulty lies in applying those same principles when an AI agent, rather than a person, performs the work. The shift brings five common friction points into focus.

HIPAA does not carve out an exception for artificial intelligence, and it does not need one. HHS OCR has been clear that legacy HIPAA frameworks apply fully to any AI or LLM that accesses, stores, or transmits electronic PHI.

In an agentic AI workflow, agents need to move across several systems to complete a single task. Every new connection expands the trust boundary around PHI, introducing new governance challenges.

Understanding these trust boundaries is a prerequisite for designing healthcare

5 Friction Points to Check During Deployment

1. **Minimum Necessary Standard**

AI agents retrieve information from multiple systems and without clear boundaries, they may access more patient information than what's required.

Things to check:

- **Limit AI access** to only the data required for the task.
- **Define permissions** based on the workflow rather than broad user roles.
- **Review access regularly** as AI capabilities expand.

2. Auditability

It creates a record of an AI system's behavior over time. Instead of treating AI outputs as a black box, organizations need a clear way to retrace the events that produced them and verify the system acted as expected.

Things to check:

- Maintain logs for data access and AI-generated actions.
- Record prompts, model versions, and system responses where appropriate.
- Make AI activity easy to review during investigations or audits.

3. Training Data and Memorization

Using patient data to train or improve AI models requires additional care. Organizations should avoid exposing identifiable patient information during training unless appropriate safeguards are already in place.

Things to check:

- Exclude patient identifiers before training or fine-tuning models.
- Separate production data from training environments whenever possible.
- Review data retention policies for every AI service being used.

4. Numerical Representation of Patient Data

Patient information doesn't become less sensitive simply because it has been converted into a machine-readable format for AI retrieval or semantic search. The same governance controls should follow that data throughout its lifecycle.

Things to check:

- Protect vector databases using the same security controls as other patient data stores.
- Apply encryption and access controls to protect data integrity.
- Monitor how embeddings are created and stored.

5. Human Oversight

A designated reviewer should be responsible for reviewing AI recommendations before they influence patient care or critical operational decisions.

Things to check:

- Identify workflows that require human approval before an action is taken.
- Clearly assign ownership for AI-generated recommendations and decisions.
- Review high-impact AI actions regularly to validate performance and identify unexpected behavior.

Run every AI use case through these five checks before go-live. A tool that fails even one of them isn't ready for PHI, regardless of how good its output looks in a demo.

The Big Rule

Never assume an AI provider is HIPAA compliant by default. Compliance depends on the deployment architecture, security configuration, contractual agreements, and how patient data is handled—not on the vendor's reputation or popularity.

Chapter 4

Risk Profiles Differ Among 7 AI Capabilities

The friction points raised in Chapter 3 aren't solved by a single fix. Each of the seven capabilities addresses a different aspect of AI governance. Together, they help leaders identify existing strengths and gaps that still need attention.

Why Each AI Capability Introduces Different Governance Risks

AI, when it's being used in each capability, solves different problems across the enterprise. For instance, the foundations and processes involved in modernizing an Electronic Medical Record (EMR) system are different from deploying a large language model or automating claims processing.

As a result, the governance decisions behind each one are different too.

Building one AI governance framework as a common capability can't be applied everywhere. You can see that each capability introduces its own operational decisions around patient data.

Our goal is to understand how governance priorities change when different AI capabilities come together during everyday healthcare operations and how each one contributes to HIPAA compliant AI.

1. Modernization

During migration, the main challenge is moving patient data into new environments. When a modernization process takes place, PHI passes through migration tools, integration platforms, cloud storage, and temporary repositories that weren't part of the original architecture.

This new data path should follow the same governance standards as the production system; otherwise, the organization creates new exposure while trying to reduce technical debt.

Validation Checkpoints:

- Encryption and access controls remain consistent throughout migration.
- Data integrity is verified after every migration stage.
- New integrations follow the same governance standards as existing clinical systems.

2. Data Engineering

Clinical records, claims, laboratory systems, and imaging platforms contribute information that eventually supports analytics and AI. An AI system can only be as reliable as the data supporting it.

Missing lineage, inconsistent metadata, duplicate records, or unclear ownership make it difficult to understand where information originated or whether it should be used at all. Governance focuses on creating traceable data before AI begins consuming it.

Validation Checkpoints:

- Metadata and ownership are documented across datasets.
- De-identification and consent policies are consistently enforced.
- Data quality controls are built into ingestion and transformation pipelines.

3. Generative AI Deployment

Once you deploy Gen AI, any new interaction introduces a new way for patient information to enter an AI system. Gen AI governance includes how prompts are processed, where inference takes place, whether responses are stored, and how data moves across supporting infrastructure. The overall architecture finally determines how PHI is handled.

Validation Checkpoints:

- Gen AI deployment architecture is reviewed before the model is selected.
- Contractual agreements, including a Business Associate Agreement (BAA) where applicable, are in place.
- Prompt handling, retrieval pipelines, vector databases, and generated outputs follow approved data handling practices
- Patient information is retained only according to organizational policies.

4. Agentic AI

Rather than individual prompts, AI agents retrieve information, coordinate tasks across multiple systems, and complete workflows[3] with limited human intervention. Autonomous execution changes how governance works. Risk includes both AI-generated outputs and the actions AI agents perform. When agents access EHRs, organizations should be able to explain what decisions were made, why, and what systems were affected.

Validation Checkpoints:

- AI agents operate within clearly defined task boundaries.
- Permissions are limited to the systems required for each workflow.
- Every action remains traceable across connected applications.
- Human oversight exists for actions with clinical, financial, or regulatory impact.

5. Intelligent Automation

Healthcare information continues to circulate across applications where decisions are made autonomously, and thousands of transactions can be completed without someone reviewing each one.

The governance challenge is maintaining visibility into how patient information moves through the workflow. Organizations should be able to trace that movement and identify where exceptions were escalated for human review.

Validation Checkpoints:

- Automated workflows access only the information required for each task.
- Audit trails capture every automated action involving patient data.
- Business rules are reviewed as workflows evolve.
- Sensitive exceptions are routed to human reviewers.

6. Quality Engineering

Testing environments often contain patient information, making them just as important to govern as production systems. Organizations should carefully decide if production data is needed for testing, should be masked, or if synthetic data is a better option.

Test environments, logs, screenshots, and defect reports can all expose Protected Health Information (PHI) if they are not handled properly. The goal is to apply the same level of AI compliance protection throughout the testing process.

Validation Checkpoints:

- Use masked or synthetic data whenever practical.
- Restrict access to testing environments and test data.
- Protect screenshots, logs, and defect reports that may contain PHI.
- Apply the same security controls in testing as in production.

7. Testing AI Systems

Traditional software testing confirms whether an application behaves as expected. AI testing requires additional evaluation because model outputs can change as models evolve and new data influences behavior.

Outputs should be traceable to the reasoning or data that produced them. Clinicians and auditors should be able to understand why the system responded the way it did.

Validation Checkpoints:

- Hallucinations are evaluated using realistic healthcare scenarios.
- Model drift is monitored throughout the model lifecycle.
- Prompt security testing identifies unsafe or unintended behavior.
- AI performance is reviewed continuously after deployment.
- Apply the same security controls in testing as in production.

Chapter 5

Enterprise Standards for Healthcare AI Governance

Governance should be part of the AI development process from the beginning. These five practices help organizations do that consistently.

Governance Is Part of the Development Process

Legal and risk review as a final checkpoint before launch creates unnecessary friction. It should influence the decisions made throughout development, from data and model selection to testing and release.

Resolving data sourcing and model behavior during development reduces rework and creates a stronger foundation for compliant AI systems.

Treated this way, governance stops being a barrier to speed and becomes the reason new AI initiatives can launch without starting compliance from zero each time.

Governance Design Patterns for Long-Term AI Success

Governance practice	In action
Classify data & block shadow AI	Define what data may and may not flow into AI systems. Monitor API calls and block unsanctioned external AI tools from the network.
Establish guardrails & escalations	Establish clear limits for AI autonomy and define where human approval or intervention is required.
Ensure traceability & auditability	Log the reasoning behind each decision, along with the specific service account, user, or agent responsible for it. Maintain a clear chain of custody for every AI-driven action.
Adopt recognized frameworks	Anchor internal standards to the National Institute of Standards and Technology AI Risk Management Framework (NIST), the Health Information Trust Alliance Common Security Framework (HITRUST), or another recognized framework instead of building a bespoke standard from scratch.
Automate compliance tasks	Use Governance, Risk, and Compliance (GRC) platforms to monitor for model drift and collect audit evidence continuously, instead of assembling it by hand before each review.

Additional Governance Standards for AI

HIPAA sets the foundation. Other frameworks help organizations apply compliance requirements throughout the AI lifecycle.

Framework / Regulation	Issued by	Key purpose / How it helps
NIST AI Risk Management Framework (AI RMF)	National Institute of Standards and Technology	Helps identify and manage AI risks across the full system lifecycle.
NIST Cybersecurity Framework (CSF)		Widely used cybersecurity framework that complements HIPAA's Security Rule by strengthening cybersecurity risk management.
NIST Privacy Framework		Helps organizations incorporate privacy risk management into system design and operations.
ONC HTI-1 Final Rule	Office of the National Coordinator for Health Information Technology: Health Data, Technology, and Interoperability	Introduces transparency requirements for certain predictive decision support interventions (DSIs) used in certified health IT.
FDA guidance	U.S. Food and Drug Administration Guidance	Applies when AI functions as Software as a Medical Device (SaMD) or is incorporated into a regulated medical device.
21st Century Cures Act	21st Century Cures Act	Promotes interoperability, and prohibits information blocking.
HITRUST CSF	Health Information Trust Alliance Common Security Framework	A certifiable framework widely used by healthcare organizations to demonstrate comprehensive security and privacy controls.

Pick one framework as your anchor. Retrofitting governance onto live systems costs far more than designing it in from the first deployment.

Chapter 6

6 Decision Pathways for Enterprise Leaders

Healthcare AI programs reach a point where leadership needs to make six key decisions. Defining them early brings clarity to planning.

The Six Decision Pathways at a Glance

Pathway	Strategic Priority	What's at Stake
Build vs. Buy	Decide which governance capabilities belong in-house and which come from trusted partners.	Long-term control, implementation speed, and investment.
Internal Development vs. Vendor Reliance	Define ownership across internal teams and third-party providers.	Accountability when something goes wrong.
Legacy Systems vs. Modernization	Evaluate whether current systems are ready for enterprise AI.	Delayed innovation and growing security exposure.
AI Capability vs. Data Exposure	Limit AI access to only the data it needs.	Better AI outcomes without unnecessary PHI exposure.
Local Innovation vs. Cross-Border Compliance	Prepare AI for different regulatory environments.	Expansion limited by regional compliance requirements.
Periodic Reviews vs. Continuous Governance	Move governance into everyday AI operations.	Falling behind systems that change every day.

Final Takeaway

Don't try to solve all six pathways at once. Start with the first one. The choices you make around governance capabilities and technology partners will influence every decision that follows. Those same decisions stay with you throughout AI transformation.

Conclusion

Build AI Solutions for Long-Term Governance

We stated early in the eBook that a model recommendation is not our goal. The point to carry forward from this resource is to offer a way to think through an AI governance framework, so the model choice becomes straightforward. The thread runs through everything covered so far.

Your compliance decisions need to follow a clear sequence. Classify the data before selecting the tool. Understand which capabilities a project depends on before assigning oversight. Anchor it to a recognized standard that fits the organization – before there are several systems in production with no shared reference point across them.

Data Governance for Sustainable AI

The starting point is different for each healthcare organization. Indium spends far more time working on the decisions behind AI than the model itself. Most of that work happens around legacy modernization, data readiness, AI guardrails, testing, and continuous maintenance.

We see that governance isn't missing. Healthcare organizations already have review boards, and compliance checkpoints. The issue starts with where governance takes place. When a project reaches formal review, the architecture is already complete.

The next logical step is to seek approval rather than consider how it could have been designed differently.

The impact often becomes clear only when an audit or security review uncovers something unexpected. The earlier governance becomes part of the build, the easier it becomes to create AI that is ready for what comes next.

Checklist

Is Your Organization Ready for HIPAA-Compliant AI?

Use this checklist as a gut check before your next AI deployment. It draws directly on the chapters above – if a question doesn't have a confident answer, that's the place to start.

01	Can you name every AI tool currently touching PHI in your organization, including ones adopted informally?	☐
02	Does every vendor in that list have a signed BAA on file, and does it explicitly address AI and machine learning use?	☐
03	Do your AI tools operate on a Zero Data Retention basis, or otherwise avoid using PHI to train a vendor's general model?	☐
04	Can you scope what data an AI system can access to the minimum necessary for its specific task, rather than broad standing access?	☐
05	Do you maintain access logs, decision logs, prompt history, and version history for every AI-driven action?	☐
06	Have you classified which of the seven capabilities in Chapter 4 each AI initiative depends on, with a clear owner for each?	☐
07	Have you anchored your AI governance program to a recognized framework, such as the NIST AI RMF or HITRUST CSF?	☐
08	Do your testing environments avoid using raw production PHI, relying instead on masked or synthetic data?	☐
09	Is your audit process continuous, or does it still rely on a periodic, point-in-time review?	☐
10	If regulators asked today, could you produce documentation showing who accessed what data, when, and why?	☐

References

1. Doximity (2026). Doximity 2026 State of AI in Medicine Report. [online] Doximity.com. Available at: https://www.doximity.com/reports/state-of-ai-medicine-report/2026 [Accessed 31 July 2026].

2. November 2023 AMA Augmented Intelligence Research Physician sentiments around the use of AI in heath care: motivations, opportunities, risks, and use cases 1 Background and Objectives. (2025). [online] Available at: https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf [Accessed 31 July 2026].

3. Introducing OpenAI for Healthcare. (2026). [online] Openai.com. Available at: https://openai.com/index/openai-for-healthcare/ [Accessed 31 July 2026].

About Indium

We are an AI services company specializing in Agentic AI, Data & Analytics, Application Engineering, and Quality Engineering.

By combining autonomous AI systems, modern engineering, and data intelligence, we build innovative AI solutions that drive measurable business outcomes and long-term value.

With 5,000+ associates worldwide, Indium partners with Fortune 500 companies, Global 2000 enterprises, and leading technology firms in Financial Services, Healthcare, Manufacturing, Retail, and Technology. Our teams operate in North America, India, the UK, Singapore, Australia, and Japan, helping businesses stay ahead in an AI-first world.

USA
Cupertino | Princeton
Toll-free: +1-888-207-5969

INDIA
Chennai | Bengaluru | Mumbai | Hyderabad | Pune
Toll-free: 1800-123-1191

UK
London
Ph: +44 1420 300014

SINGAPORE
Singapore
Ph: +65 6812 7888

www.indium.tech

For Sales Inquiries
sales@indium.tech

For General Inquiries
info@indium.tech