

AI Readiness for Healthcare

Building the governance, data, security, and operational foundation for responsible AI in healthcare organizations

Executive takeaway

AI readiness for healthcare is the ability to use AI while keeping ePHI protected, access controlled, vendors accountable, outputs reviewed, and care running. It depends less on which AI tool you choose and more on ePHI governance, HIPAA-aligned risk management, Microsoft 365 permissions, identity controls, vendor BAAs, audit trails, backup, human review, and patient trust. This guide gives healthcare leaders a seven-domain readiness model, a four-level maturity scale, a regulatory map, a 90-day roadmap, and a ten-question self-check. AdvancedBITS helps healthcare organizations build the secure, governed foundation first, then scale AI with confidence.

Healthcare is one of the most data-intensive and trust-sensitive environments in business. Every workflow may touch protected health information, patient communications, care coordination, billing details, staff credentials, referral records, and insurance data. AI can reduce administrative burden and improve visibility into operations, but it also raises the stakes on every weakness already in the environment.

A Business Associate Agreement (BAA) is not a readiness plan. HIPAA-focused guidance repeatedly makes the point that a BAA is only one piece: the organization must still configure systems, govern access, train users, monitor activity, validate outputs, and document how ePHI is protected. AdvancedBITS treats healthcare AI as a security and operating-model challenge first and a software decision second.

Why healthcare AI readiness matters now

Healthcare is adopting AI faster than it is governing it. Physician use has more than doubled in three years, while independent research shows many organizations still lack controls over the sensitive data AI tools touch.

Table 1. Healthcare AI adoption and governance, by the numbers

Finding	Source
81% of physicians use AI professionally, up from 38% in 2023, and the average physician now reports 2.3 use cases, up from 1.1.	American Medical Association, 2026 Physician Survey on Augmented Intelligence (March 2026; 1,692 physicians)
Patient privacy was the only area where more physicians expect AI to cause harm than help (41% versus 13%).	AMA 2026 survey, as reported by Fierce Healthcare (March 2026)
17% of healthcare professionals surveyed admitted using an unauthorized AI tool and 40% had encountered one. About 1 in 10 had used one for direct patient care.	Wolters Kluwer Health survey of 518 healthcare professionals (December 2025; published January 2026)

Finding	Source
60% of healthcare organizations lack controls governing sensitive data used in AI systems. Only 42% use security measures designed for AI, and just 32% have full visibility into unsanctioned AI use.	Proofpoint and Ponemon Institute, 2026 Cyber Insecurity in Healthcare (October 2026; 705 practitioners)
63% of breached organizations lacked AI governance policies. 1 in 5 reported a breach tied to shadow AI, which added about \$670,000 to breach cost. Of organizations reporting an AI-related breach, 97% lacked proper AI access controls.	IBM Cost of a Data Breach Report 2025 (600 organizations, all industries)
93% of IT leaders were confident their Microsoft 365 governance could support AI, yet 29% said an AI tool had already surfaced sensitive data it should not have reached.	ShareGate, State of Microsoft 365 2026 (851 IT leaders)
Misuse of AI chatbots ranked #1 on ECRI's Top 10 Health Technology Hazards for 2026.	ECRI (January 2026)

Survey figures are self-reported. IBM and ShareGate figures cover all industries or all AI tools, not healthcare or Copilot alone. Healthcare was about 2% of IBM's sample.

The pattern across these sources is consistent. People are adopting AI because it saves time, and Wolters Kluwer found that speed was the top reason staff used unapproved tools. When an organization does not provide approved options and clear rules, staff fill the gap themselves.

The AdvancedBITS healthcare AI readiness model

Healthcare AI readiness has seven domains, and a gap in any one of them can undermine the rest. The model below replaces ad hoc tool-by-tool decisions with a single view leadership can review.

Table 2. Seven readiness domains

Readiness domain	Question to ask	Risk if ignored	Evidence you should have
1. Governance and accountability	Who owns AI decisions, and how are tools approved, monitored, and retired?	Tools spread without an owner, a policy, or a record of decisions.	AI policy, named executive owner, approved tool list, review cadence.
2. ePHI and data readiness	Which workflows touch ePHI, claims, patient messages, or staff data?	Sensitive data enters AI tools without controls.	Data classes, approved and prohibited data rules, sensitivity labels.
3. Microsoft 365 and Copilot readiness	Are SharePoint, Teams, OneDrive, and Outlook permissions ready for AI search?	AI surfaces overshared or forgotten content.	Permission review, site owners, DLP, audit logging, retention.
4. Identity, security, and resilience	Can we prevent account compromise and recover from ransomware or deletion?	AI workflows depend on systems that are not protected or recoverable.	MFA, Conditional Access, endpoint protection, tested restores, incident response plan.

Readiness domain	Question to ask	Risk if ignored	Evidence you should have
5. Vendor and BAA oversight	Which vendors create, receive, maintain, or transmit PHI, and what do contracts say?	Third-party exposure goes undocumented.	Vendor inventory, BAAs, retention and training-use terms, incident commitments.
6. Workflow and human review	Which workflows are mature enough for AI, and who reviews the output?	AI accelerates inconsistent processes and unreviewed drafts reach patients.	Mapped workflows, review owners, escalation paths.
7. Workforce, training, and measurement	Do staff know the rules, and can we show the results?	Policy exists on paper only, and value cannot be proven.	Training records, usage monitoring, pilot metrics reported to leadership.

Four levels of AI readiness

Most healthcare organizations sit at level 1 or 2 today. The maturity scale shows what each level looks like, what it costs you, and the next move.

Table 3. AI readiness maturity model

Level	What it looks like	What it costs you	Next step
1. Ad hoc	Staff use whatever AI tools they find. There is no inventory, policy, or owner.	Unknown ePHI exposure, no evidence for an insurer or auditor, and no safe basis for approving Copilot.	Inventory AI use and name an executive owner.
2. Aware	Leadership knows AI is in use. A draft policy exists, but permissions, labels, and vendor reviews are incomplete.	Policy without enforcement. Staff cannot tell which tools are approved, and overshared content is still searchable.	Clean up permissions and publish approved and prohibited tool lists.
3. Governed	Approved tools, ePHI rules, labels, DLP, review points, and BAAs are in place for pilot workflows.	Manual monitoring effort. Scaling slows if reviews are not routine.	Run controlled pilots with measured outcomes and audit-log review.
4. Managed	Continuous monitoring, periodic access reviews, tested restores, vendor reassessment, and regular reporting to leadership.	Ongoing discipline, which is the price of scale.	Expand to more departments and review against national guidance.

Healthcare AI regulation and standards: what applies

No single law governs healthcare AI, so organizations combine HIPAA, federal nondiscrimination rules, state AI laws, and voluntary standards. The table summarizes what applies today. It is general information, not legal advice, and state rules change quickly, so confirm requirements with counsel.

Table 4. Regulatory and standards map for healthcare AI

Framework	Status	What it means for AI readiness
HIPAA Privacy, Security, and Breach Notification Rules	In force	AI vendors handling PHI on your behalf are typically business associates and need a BAA. You must still perform a risk analysis, manage risk, and protect ePHI. OCR's risk analysis enforcement initiative is expanding to cover risk management.
Proposed HIPAA Security Rule update	Proposed, not final. HHS now projects final action in July 2027.	Would remove "addressable" specifications and require a technology asset inventory and network map. Industry summaries also describe required MFA, encryption, and 72-hour recovery. More than 100 provider groups have asked HHS to withdraw it, so the final text may change.
ACA Section 1557, patient care decision support tools (45 CFR 92.210)	In effect since May 1, 2025, per published guidance	Covered entities must make reasonable efforts to identify decision support tools, including AI, that use protected-class inputs and to mitigate discrimination risk. Tools unrelated to clinical decision-making, such as scheduling, are outside the definition.
Joint Commission and CHAI, Responsible Use of AI in Healthcare	Voluntary guidance released September 17, 2025. A voluntary certification now exists.	Seven recommended elements: AI policy and governance structures, patient privacy and transparency, data security and data use protections, ongoing quality monitoring, voluntary safety-event reporting, risk and bias assessment, and staff training.
Texas SB 1188 and HB 149 (TRAIGA)	SB 1188 effective September 1, 2025. HB 149 effective January 1, 2026.	Practitioners using AI for diagnostic purposes must stay within their license, review AI-generated records under Texas Medical Board standards, and disclose AI use to patients. SB 1188 also requires Texas patients' electronic health records to be stored in the United States. HB 149 adds patient disclosure duties for AI used in treatment.
California AB 3030	Effective January 1, 2025	Clinics, health facilities, and physician offices using generative AI for patient clinical communications must add a disclaimer and instructions for reaching a human, unless a licensed provider reviews the message. Administrative messages such as scheduling and billing are excluded.

Readiness is how you meet all of these at once. An organization with an AI inventory, a risk analysis, human review points, vendor BAAs, and disclosure procedures is positioned for HIPAA, Section 1557, state disclosure laws, and Joint Commission guidance, instead of treating each as a separate project.

Microsoft 365 Copilot readiness for healthcare

Copilot is powerful because it works across the content a user can already access, which is also why permissions come first. It can summarize meetings, draft documents, search

policies, and prepare communications. It can also surface ePHI or sensitive material that was overshared long before AI arrived.

ShareGate's 2026 research shows the gap between confidence and control: 93% of IT leaders were confident their governance could support AI, 29% reported an AI tool had already surfaced data it should not have, and 53% lacked full visibility into what their Copilot agents can access. The findings are self-reported and the question covered AI tools generally, so they show a governance gap and not a Copilot-specific failure rate.

A BAA is available, but it does not do the configuration work. Microsoft states that its HIPAA BAA is available by default through its Data Protection Addendum to covered entities and business associates, and that it applies to in-scope services listed on Microsoft Learn. Confirm that each AI feature your staff will use appears on that list, then handle the work the agreement cannot: permissions, labels, DLP, retention, audit logs, and training.

Table 5. Copilot readiness checklist

Readiness item	Why healthcare organizations should care	AdvancedBITS action
Oversharing review	Copilot can surface old department files, reports, or ePHI-containing documents users technically can open.	Review SharePoint, Teams, OneDrive, group permissions, guest access, and sharing links before deployment.
ePHI sensitivity labels	Sensitive content should not rely on folder names alone.	Define practical labels for ePHI, confidential, internal, restricted, and administrative content.
DLP and access control	AI workflows should not make it easier to copy or send PHI to unauthorized destinations.	Review DLP, Conditional Access, MFA, endpoint posture, and approved sharing paths.
Site ownership and lifecycle	Clinic, department, and project sites stay active without accountable owners.	Assign owners and define create, review, archive, and closure rules.
Audit logging	You need evidence of how AI is used to support investigations and compliance.	Enable audit logs and a review cadence for high-risk workflows.
Controlled pilot	Organization-wide rollout before cleanup increases exposure.	Start with a limited pilot, validate permissions, monitor use, and adjust before scaling.

ePHI-centered SharePoint and knowledge governance

Repeatable workspace standards keep document sprawl from becoming AI exposure.

Without them, AI adoption can expose years of unmanaged files, including templates and prior work that contain patient-specific details.

Table 6. Workspace governance standards

Governance area	Common challenge	AdvancedBITS guidance
Workspace templates	Departments create folders and Teams channels differently.	Create templates for clinic operations, referrals, billing, compliance, quality, HR, and projects.

Governance area	Common challenge	AdvancedBITS guidance
Permission groups	Roles change, and contractors come and go and access drifts.	Use group-based access tied to job roles and the minimum-necessary principle.
Naming and metadata	Policies, forms, and reports are hard to find.	Define simple naming and metadata for department, location, content type, and sensitivity.
Lifecycle and retention	Old project sites and ePHI-containing files remain active.	Use closure checklists, retention rules, archive policies, and periodic access review.
Knowledge reuse	Templates and prior work may contain patient-specific data.	Create curated repositories with approved templates and de-identified examples where appropriate.
Backup and recovery	Retention is not the same as recoverability.	Align governance with Microsoft 365 backup, restore testing, and incident response.

AI governance in practice: rules staff can follow

Healthcare AI governance turns privacy, security, and clinical-safety expectations into rules employees can follow. Staff need to know which tools are approved, what data cannot be entered, when a human must review, and what is off limits. Wolters Kluwer found that administrators were three times as likely as providers to be involved in AI policy development (30% versus 9%), which is a reason to include clinical voices in governance.

Table 7. Governance standards

Governance topic	Question to resolve	Practical standard
Approved tools	Which AI tools may staff use for healthcare work?	Maintain approved, restricted, and prohibited lists by workflow and data type.
Restricted data	What may not be entered into public or unapproved AI tools?	Name ePHI, diagnoses, MRNs, claims data, patient messages, staff records, and sensitive business data.
Human review	Who validates AI output before use?	Require review for patient-facing communication, clinical-support drafting, billing actions, HR materials, and compliance documents.
Patient communication and disclosure	When may AI help draft patient messages, and when must patients be told?	Use approved systems, require review, avoid unsupported clinical advice, and follow state disclosure rules.
Audit and monitoring	Can the organization see how AI is used?	Enable audit logs, usage reviews, and DLP alerts, and report exceptions.
Training and accountability	Do employees understand safe AI use?	Train on approved use, prohibited data, phishing risk, prompt hygiene, and escalation.

Not every AI use case carries the same risk. The tiers below help leaders decide where to start. ECRI ranks misuse of general-purpose chatbots as the top health technology hazard of 2026 because the tools sound confident even when wrong, which is why clinical uses sit in the most controlled tier.

Table 8. AI use case risk tiers

Tier	Examples	Controls required
Lower risk: administrative support	Meeting summaries, policy and procedure search, SOP drafting, executive reporting, cyber and compliance summaries.	Approved tools, clean permissions, human review before distribution.
Moderate risk: operational and patient-adjacent	Referral and intake summaries, patient communication drafts, revenue-cycle and denial analysis, outreach drafting.	BAA, ePHI rules, review owner for every output, state disclosure check, validation against source records.
Restricted: clinical	Diagnostic support, treatment recommendations, patient-specific clinical advice.	Clinical governance, documented clinician review, bias and nondiscrimination review, patient disclosure, and legal review before use.

Vendor security and BAA oversight for healthcare AI

Vendor review should not stop at whether a vendor will sign a BAA. Prompts, uploaded files, outputs, logs, transcripts, and summaries may all contain ePHI. The 2025 Conduent breach, the largest healthcare breach of the year at 62,224,658 individuals, involved a business associate and a vendor environment rather than a provider network.

Table 9. Vendor types and review points

Vendor type	Risk to review	AdvancedBITS guidance
AI platforms	Prompts, files, summaries, and outputs may contain ePHI.	Review retention, training use, BAA status, audit logs, data location, and access controls.
EHR and practice systems	Clinical, scheduling, and billing records are core systems.	Review access, MFA, integrations, backup, audit logs, and vendor security duties.
Telehealth and patient portals	Patient communication occurs through external services.	Review MFA, privacy safeguards, retention, encryption, and incident commitments.
Billing and RCM vendors	Claims data and payer correspondence may contain PHI.	Review BAA, minimum-necessary access, audit logs, retention, and continuity.
Microsoft 365 and collaboration vendors	Email, Teams, SharePoint, OneDrive, and backup may store ePHI.	Review configuration, DLP, labels, backup, restore testing, and access reviews.

Vendor type	Risk to review	AdvancedBITS guidance
Outsourced IT and security providers	Providers may hold privileged access to sensitive systems.	Review admin access, MFA, logging, incident response, contractual duties, and continuity.

CRM, referral, and patient engagement readiness

CRM and outreach data is a distinct AI readiness area because it blends patient inquiries, referral relationships, and marketing activity. AI can help with outreach drafts, campaign summaries, and referral trend analysis, but only if the underlying data is governed and privacy aware.

Table 10. Patient engagement and referral readiness

Area	Why it matters	Readiness practice
Source attribution	Leaders need to know which referral sources and campaigns produce qualified volume.	Track source, service line, location, owner, campaign, status, and outcome.
Follow-up discipline	Referrals and inquiries fall through gaps without ownership.	Define next-action fields, escalation rules, and stale-record alerts.
Patient communication workflow	Outreach must be accurate, privacy-aware, and reviewed.	Use approved templates, reviewed AI drafts, role-based access, and documented standards.
Privacy and consent	CRM records may include sensitive inquiry data.	Define what belongs in the CRM, who can see it, and when added privacy controls apply.

Cyber resilience is an AI prerequisite

AI workflows depend on systems that must stay protected and recoverable. In Proofpoint and Ponemon research, 97% of healthcare organizations experienced at least one cyberattack in the past year and 74% of those reported disruption to patient care. Disruption to normal operations was the largest single cost category at an average of \$1.39 million, and 24% of respondents named AI agents operating without guardrails as a leading cause of data loss. ECRI's 2026 list also ranks unpreparedness for a "digital darkness" event second, behind only AI chatbot misuse.

The Change Healthcare attack shows how a single control gap becomes a national event. Attackers used stolen credentials on a Citrix remote access portal without multi-factor authentication, and the final estimate reached 192.7 million individuals. The same discipline that protects against that scenario also protects an AI program: enforce MFA, segment networks, monitor identity activity, and test restores. For controls, a 90-day security roadmap, and independent benchmarks such as the 99.22% reduction in account-compromise risk reported by Microsoft Research, see our healthcare cybersecurity and HIPAA case study.

A 90-day healthcare AI readiness roadmap

A focused 90-day plan can move most organizations from level 1 or 2 to level 3. Each phase ends with evidence leadership can review.

Table 11. 90-day roadmap

Timeframe	Priority	Evidence at the end of the phase
Days 1 to 30: See	Inventory AI tools and use, Microsoft 365 permissions, SharePoint sites, Teams channels, vendors, BAAs, CRM and outreach workflows, and high-risk repositories. Name an executive owner.	An AI inventory, a permissions findings report, and a vendor and BAA list.
Days 31 to 60: Govern	Publish the AI acceptable use policy, ePHI handling rules, approved and prohibited tools, labels, DLP priorities, human review points, vendor criteria, and state disclosure procedures. Close the highest-risk permission gaps and confirm MFA and backups.	A signed policy, labeled content in priority areas, and a tested Microsoft 365 restore.
Days 61 to 90: Pilot and prove	Pilot two or three lower-risk workflows with human review, audit-log monitoring, and staff training. Report results to leadership and decide what to scale.	Pilot metrics, audit-log findings, and a scale or stop decision with named owners.

Measure the pilot so leadership can decide with evidence. Useful indicators include:

- Minutes saved per task compared with the manual baseline
- Share of AI drafts changed or rejected during human review
- Number of unapproved AI tools detected, then approved or retired
- DLP matches and audit-log findings in AI-assisted workflows
- Time to restore a sample of Microsoft 365 content during a test

Healthcare AI readiness self-check

Answer yes or no to each question and count your yeses. This is an AdvancedBITS planning tool, not a validated or certified assessment.

Table 12. Ten-question readiness self-check

#	Readiness check	Domain
1	We have a written AI acceptable use policy that staff have been trained on.	Governance
2	We know which AI tools staff are using today, approved or not.	Governance
3	An executive owner is accountable for AI decisions, with compliance, IT, and clinical input.	Governance
4	We have defined which data may never be entered into AI tools.	ePHI and data

#	Readiness check	Domain
5	We have reviewed Microsoft 365 permissions and sharing links in the last 12 months.	Microsoft 365
6	ePHI is labeled or otherwise classified, with DLP and audit logging enabled.	Microsoft 365
7	MFA is enforced on email, remote access, and administrator accounts.	Security
8	We have test-restored Microsoft 365 or EHR-related data and know how long it took.	Resilience
9	Every AI vendor that touches PHI has a BAA and a documented review.	Vendors
10	A named person reviews AI output before it reaches a patient, a payer, or a regulator.	Human review

Table 13. Scoring

Yes answers	Where you stand	Recommended move
0 to 3	Level 1 to 2: foundation first	Pause new AI rollouts. Inventory use, publish a policy, and fix permissions and MFA.
4 to 7	Level 2 to 3: partly ready	Close the gaps you answered no to, then pilot lower-risk workflows.
8 to 10	Level 3 and above: ready to pilot	Run measured pilots, review audit logs, and plan scale-up with leadership.

Questions healthcare leaders should ask

- Which AI tools are clinical, administrative, billing, outreach, HR, and leadership teams already using?
- What ePHI, claims data, patient communication, or staff information could enter an AI system?
- Are SharePoint, Teams, OneDrive, and Outlook permissions ready for Copilot or AI-powered search?
- Do AI and SaaS vendors have BAAs, security controls, audit logs, retention terms, and incident commitments?
- Who reviews AI-assisted patient communications, billing drafts, and operational reports?
- Which state AI disclosure rules apply to our patients and locations?
- Can Microsoft 365 content be restored after ransomware, accidental deletion, or account compromise?
- How will we measure AI value while preserving privacy, accuracy, and patient trust?

Where AdvancedBITS fits

AdvancedBITS brings more than 11 years of managed services experience and a SOC 2 Type II examined security practice to healthcare organizations that want a secure, governed foundation for AI. We do not replace clinical judgment, compliance leadership, or patient-care accountability. We strengthen the technology environment around healthcare operations so teams can use AI more safely and effectively, whether they outsource IT entirely or co-manage with an internal lead.

Table 14. How AdvancedBITS supports healthcare AI readiness

AdvancedBITS capability	How it supports AI readiness	Result for the organization
AI readiness assessment	Scores the seven domains, reviews approved tools, ePHI restrictions, use cases, and human review.	A maturity level, a gap list, and a prioritized plan.
Microsoft 365 governance	Reviews SharePoint, Teams, OneDrive, permissions, labels, DLP, retention, audit logs, backup, and Copilot readiness.	Reduced oversharing and a safer AI search foundation.
Vendor and BAA review support	Assesses AI platforms, EHR-related vendors, telehealth, billing, Microsoft 365, and outsourced providers.	Third-party risk made visible before sensitive data is exposed.
Cybersecurity and identity	Strengthens MFA, Conditional Access, endpoint protection, email security, monitoring, and account recovery.	Lower risk of account compromise and ePHI exposure.
Backup and resilience	Reviews Microsoft 365 backup, restore testing, continuity, and incident response.	Better preparation to recover after disruption.
Policy and training support	Drafts AI acceptable use rules, review procedures, and staff training aligned to your state and workflows.	Rules staff can follow and evidence leadership can show.

Closing perspective

AI can help healthcare organizations improve productivity, reduce administrative burden, and strengthen operational visibility, but productivity cannot come at the expense of privacy, security, accuracy, or patient trust. The organizations that use AI well will be the ones that first strengthen the foundations AI depends on.

For a practical illustration of this approach, see our AI in Healthcare case study. For the security side, see our healthcare cybersecurity and HIPAA case study.

Take the next step

Find out where your organization stands. AdvancedBITS offers a no-obligation healthcare AI readiness review covering Microsoft 365 permissions, AI tool use, vendor BAAs, and backup readiness, with a plain-language summary leadership can act on.

Call 561-300-3015 or visit www.advancedbits.com to schedule a conversation.

Frequently asked questions

What is AI readiness for healthcare?

AI readiness for healthcare is the ability to adopt AI while protecting ePHI, controlling access, holding vendors accountable, reviewing outputs, and keeping care running. It covers governance, data, Microsoft 365, security, vendors, workflows, and workforce training.

Does HIPAA allow AI tools to be used with patient data?

HIPAA does not prohibit AI, but vendors that handle PHI for you are typically business associates and need a BAA. You must also complete a risk analysis, limit access to the minimum necessary, and protect ePHI. A BAA alone does not make a tool compliant.

Is Microsoft 365 Copilot HIPAA compliant?

Microsoft offers a BAA through its Data Protection Addendum that covers in-scope services listed on Microsoft Learn. Compliance still depends on your configuration. Confirm that the features you use are in scope, then review permissions, labels, DLP, audit logging, and training before rollout.

What is shadow AI in healthcare?

Shadow AI is the use of AI tools the organization has not approved. In a December 2025 Wolters Kluwer Health survey, 17% of healthcare professionals admitted using an unauthorized AI tool and 40% had seen one. Approved options and a clear policy reduce it.

What state laws apply to healthcare AI?

It depends on where you practice. Texas SB 1188 and HB 149 require patient disclosure and clinician review for AI used in diagnosis and treatment. California AB 3030 requires a disclaimer on unreviewed generative AI clinical messages. Check every state where you treat patients.

Do we need an AI policy before a pilot?

Yes. A short acceptable use policy, a list of approved and prohibited tools, ePHI handling rules, and named review owners should exist before the pilot starts. IBM found 63% of breached organizations lacked an AI governance policy.

Who should own AI governance?

An executive owner should be accountable, supported by compliance, IT and cybersecurity, and clinical or operational leaders. The Joint Commission and CHAI guidance recommends governance structures that include these perspectives.

How long does it take to become AI ready?

A focused 90-day plan can inventory AI use, close the highest-risk permission and security gaps, publish governance, and pilot lower-risk workflows. Scaling then follows from what the pilot shows.

Sources

Statistics and legal summaries were checked against the sources below as of October 22, 2026. This guide is general information and not legal advice.

1. American Medical Association. 2026 Physician Survey on Augmented Intelligence (March 2026). https://www.ama-assn.org/sites/ama-assn.org/files/2026-03/physician-ai-sentiment-report_0.pdf
2. Fierce Healthcare. AMA: Physicians' use of AI doubled from 2023 to 2026 (March 12, 2026). <https://www.fiercehealthcare.com/ai-and-machine-learning/ama-physicians-use-ai-doubled-2023-2026>
3. Wolters Kluwer Health. Shadow AI: Providers are using unapproved tools to improve workflow (January 22, 2026). <https://www.wolterskluwer.com/en/expert-insights/shadow-ai-providers-are-using-unapproved-tools-to-improve-workflow>
4. Fierce Healthcare. Nearly a fifth of healthcare workers use unauthorized AI (January 22, 2026). <https://www.fiercehealthcare.com/digital-health/nearly-fifth-healthcare-professionals-use-unauthorized-ai-tools-work>
5. Proofpoint. Healthcare Sees No Relief from Cyberattacks Disrupting Patient Care (October 13, 2026). <https://www.proofpoint.com/us/newsroom/press-releases/cyber-insecurity-in-healthcare-report-2026>
6. HIPAA Journal. Healthcare Cyberattacks Increasing and Having a Greater Impact on Patient Care (October 16, 2026). <https://www.hipaajournal.com/healthcare-cyberattacks-increasing-greater-impact-patient-care/>
7. IBM. Cost of a Data Breach Report 2025: The AI Oversight Gap. <https://www.ibm.com/think/insights/data-matters/cost-of-a-data-breach>
8. Cybersecurity Dive. 'Shadow AI' increases cost of data breaches, report finds (July 30, 2025). <https://www.cybersecuritydive.com/news/artificial-intelligence-security-shadow-ai-ibm-report/754009/>
9. HIPAA Journal. Average Cost of a Healthcare Data Breach Falls to \$7.42 Million (July 30, 2025). <https://www.hipaajournal.com/average-cost-of-a-healthcare-data-breach-2025/>
10. ShareGate. State of Microsoft 365 2026, Chapter 1: AI confidence is not control. <https://sharegate.com/state-of-microsoft-365/chap-1>
11. WindowsForum. Microsoft 365 Copilot Data Exposure (October 9, 2026), summarizing ShareGate survey methodology. <https://windowsforum.com/news/microsoft-365-copilot-data-exposure-fix-oversharing-and-permission-risks.447824/>
12. ECRI. Misuse of AI chatbots tops annual list of health technology hazards (January 21, 2026). <https://home.ecri.org/blogs/ecri-news/misuse-of-ai-chatbots-tops-annual-list-of-health-technology-hazards>
13. ECRI. Top 10 Health Technology Hazards for 2026: Executive Brief. <https://www.draeger.com/Content/Documents/Content/ECRI-Top-10-Health-Technology-Hazards-for-2026-Executive-Brief.pdf>
14. Joint Commission. Joint Commission and CHAI Release Initial Guidance on Responsible AI Adoption (September 17, 2025). <https://www.jointcommission.org/en-us/knowledge-library/news/2025-09-jc-and-chai-release-initial-guidance-to-support-responsible-ai-adoption>
15. Joint Commission. Responsible Use of AI in Healthcare certification. <https://www.jointcommission.org/en-us/certification/responsible-use-of-ai-in-healthcare>
16. Hooper Lundy and Bookman. Joint Commission and CHAI Release Governance Frameworks to Address AI in Health Care (September 25, 2025). <https://hooperlundy.com/joint-commission-and-the-coalition-for-health-ai-chai-release-governance-frameworks-to-address-ai-in-health-care/>
17. TMLT. New AI disclosure requirements for physicians passed into Texas law. <https://www.tmlt.org/resource/new-ai-disclosure-requirements-for-physicians-passed-into-texas-law>
18. Buchalter. From PHI to AI: What Texas SB 1188 Means for Healthcare Entities and Vendors (October 1, 2025). <https://www.buchalter.com/insights/from-phi-to-ai-what-texas-sb-1188-means-for-healthcare-entities-and-vendors/>

19. Medical Board of California. GenAI Notification Requirements (AB 3030). <https://www.mbc.ca.gov/Resources/Medical-Resources/GenAI-Notification.aspx>
20. McDermott Will & Emery. Section 1557 patient care decision support tools: 12 things to consider (January 16, 2025). <https://www.mcdermottlaw.com/insights/section-1557-patient-care-decision-support-tools-anti-discrimination-compliance-12-things-to-consider/>
21. Meilynx. HHS Section 1557 and Patient Care Decision Support Tools, Explained (reviewed September 5, 2026). <https://www.meilynx.com/reference/hhs-1557-decision-support-tools>
22. Microsoft Learn. HIPAA and HITECH Act compliance offering. <https://learn.microsoft.com/en-us/compliance/regulatory/offering-hipaa-hitech>
23. HHS. HIPAA Security Rule NPRM Fact Sheet. <https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/factsheet/index.html>
24. Fierce Healthcare. Feds push back HIPAA security rule overhaul to July 2027 (July 10, 2026). <https://www.fiercehealthcare.com/health-tech/feds-push-back-hipaa-security-rule-overhaul-july-2027>
25. Holland and Knight. HIPAA Security Rule Amendments Now Projected for July 2027 (July 6, 2026). <https://www.hklaw.com/en/insights/publications/2026/07/hipaa-security-rule-amendments-now-projected-for-july-2027>
26. OGC. Proposed HIPAA Security Rule: 7 Key Changes (updated October 9, 2026). <https://oneguyconsulting.com/blog/new-hipaa-security-rule-changes-2026>
27. The HIPAA Guide. OCR Reports to Congress on HIPAA Compliance and Data Breaches in 2024 (May 29, 2026). <https://www.hipaaguide.net/ocr-reports-congress-hipaa-compliance-data-breaches-2024/>
28. HIPAA Journal. Conduent Business Services Data Breach Affected More Than 62.2 Million Individuals (June 4, 2026). <https://www.hipaajournal.com/conduent-business-solutions-data-breach/>
29. The HIPAA Guide. Change Healthcare Data Breach: 192.7 Million Affected (August 7, 2025). <https://www.hipaaguide.net/change-healthcare-data-breach/>
30. Microsoft Research (May 2023), cited for the 99.22% reduction in account-compromise risk from MFA and security controls.