HIPAA-Compliant AI Tools for Healthcare Compared (2026)

Published July 20, 2026 39 min read

HIPAA-Compliant AI Tools for Healthcare Compared (2026)

Healthcare organizations evaluating artificial intelligence (AI) tools in 2026 should not treat "HIPAA compliant" as a product-wide badge. Eligibility depends on the exact service, an applicable Business Associate Agreement (BAA), configuration, and how the organization uses protected health information (PHI). HHS defines a business associate by the function performed and the relationship to a covered entity, including whether the vendor creates, receives, maintains, or transmits PHI on the covered entity's behalf (HHS. Current vendor documentation reflects that service-level approach: OpenAI lists ChatGPT for Healthcare, ChatGPT for Enterprise with Regulated Workspace, ChatGPT FedRAMP, ChatGPT for Clinicians, and API services provisioned with Modified Retention as HIPAA-eligible when covered by its BAA (OpenAI; Anthropic limits coverage to the eligible services and settings described in its BAA documentation (Anthropic; Microsoft says properly configured Microsoft 365 Copilot implementations can support HIPAA compliance but excludes web-search queries from its BAA (Microsoft; and Google describes HIPAA as a shared responsibility limited to services covered by its BAA (Google Cloud.

Beyond the general-purpose chatbots, a distinct category of healthcare-native AI has emerged, led by ambient clinical documentation ("AI scribe") tools such as **Abridge**, **Nabla**, **Suki**, and **Microsoft Dragon Copilot**, which KLAS identified as the most widely adopted clinical AI use case in its survey. A separate Microsoft-commissioned IDC study reported that 79% of surveyed healthcare organizations were using AI technology generally—not ambient documentation specifically (Microsoft. Abridge was named the number one **Best in KLAS** vendor for ambient AI for the second consecutive year in 2026 (Source: abridge.com). A Mass General Brigham-led study published in *JAMA Network Open* found ambient documentation tools produced a 21.2% absolute reduction in physician burnout prevalence at 84 days (Source: massgeneralbrigham.org), and the Permanente Medical Group rolled Nabla Copilot out to 10,000 physicians in Northern California after a successful pilot.

Compliance failures carry a real financial cost that predates generative AI but applies squarely to it: missing or deficient BAAs have contributed to OCR settlements ranging from $31,000 to $1,550,000, although the $1,550,000 North Memorial settlement also involved failure to conduct an organization-wide risk analysis, and as of January 2026 the maximum civil penalty per violation category reaches $2,190,294 annually (Source: hipaajournal.com). Physician adoption of AI tools broadly has more than doubled since 2023: the American Medical Association's 2026 survey of

nearly 1,700 physicians found 81% now use AI professionally, up from 38% in 2023 (Source: ama-assn.org). The global AI-in-healthcare market was valued at $36.7 billion in 2025 and is projected to reach $505.6 billion by 2033, a 38.9% compound annual growth rate (CAGR), according to Grand View Research (Source: grandviewresearch.com).

This report compares the HIPAA posture of the leading general-purpose AI assistants (ChatGPT, Claude, Microsoft Copilot, and Gemini) against purpose-built clinical AI vendors, details what a BAA does and does not cover, and walks through named deployments at Mass General Brigham, the Permanente Medical Group, and UCSF Health. No AI tool is HIPAA compliant "out of the box"; compliance depends on the specific plan tier, an executed BAA, and configuration choices the covered entity controls, and organizations that skip this due diligence face the same enforcement exposure that has applied to cloud vendors, billing services, and record-storage contractors for two decades.

Introduction and Background

The question "**is this AI tool HIPAA compliant**" is, strictly speaking, a category error. HIPAA does not certify products; it governs relationships and conduct. The Health Insurance Portability and Accountability Act of 1996 established, through subsequent HHS rulemaking, three governing rules: the **Privacy Rule**, which limits how protected health information (PHI) may be used and disclosed; the **Security Rule**, which requires administrative, physical, and technical safeguards for electronic PHI (ePHI) (Source: hhs.gov); and the Breach Notification Rule, which mandates disclosure of unsecured PHI incidents. These rules bind two types of actors: **covered entities** (health plans, healthcare clearinghouses, and providers who transmit health information electronically) and their **business associates**, defined as any person or entity that creates, receives, maintains, or transmits PHI on a covered entity's behalf, including "claims processing or administration; data analysis, processing or administration; utilization review; quality assurance; billing" (Source: hhs.gov); notably, "a member of the covered entity's workforce is not a business associate" (HHS. That workforce exception applies to the employee, not automatically to an outside AI provider: if a worker sends PHI to a cloud or AI service that creates, receives, maintains, or transmits the PHI on the organization's behalf, the provider may still be a business associate and an applicable BAA and permitted disclosure are still required (HHS cloud guidance. An AI vendor that processes clinical notes, transcribes patient encounters, or answers questions using PHI fits squarely within the business-associate definition, and "it's important to note that not all AI systems are designed with healthcare regulations in mind" (Source: blog.doximity.com). Independent guidance for clinicians evaluating these tools converges on the same checklist: HIPAA compliance safeguards must include "secure data encryption and storage, controlled access to patient information, audit trails to track how data is used, and business associate agreements" between the technology provider and the healthcare organization (Source: blog.doximity.com); as Doximity's own medical director notes, "these platforms vary in their feature set, ease of use, and price point," and the rise of free AI tools raises a fair question of whether "free" always means "non-compliant" (Source: blog.doximity.com).

The Privacy Rule requires that a covered entity obtain satisfactory assurances from its business associate, in writing, that PHI will be safeguarded appropriately: HHS states the rule "requires that a covered entity obtain satisfactory assurances from its business associate that the business associate will appropriately safeguard" the protected health information it receives (Source: hhs.gov). When an AI provider creates, receives, maintains, or transmits PHI on behalf of a covered entity or business associate, HIPAA generally requires an applicable written BAA before the service processes that PHI, alongside compliance with the remaining HIPAA requirements. This distinction, between a product's technical security posture and its contractual HIPAA coverage, is the single most common point of confusion in healthcare AI procurement, and it explains why "is ChatGPT HIPAA compliant" or "is Claude HIPAA compliant" do not have single yes-or-no answers: the correct answer always depends on which specific plan, which specific features, and whether a BAA has actually been executed and configured. HIPAA also requires business associates to flow the same restrictions down to their own subcontractors, since "the BAA between the business associate and its subcontractor must also contain the foregoing terms" (Source: hollandhart.com), meaning an AI vendor that itself relies on a cloud host or sub-processor must have its own BAA in place with that sub-processor before PHI can lawfully move through the stack.

This report examines that landscape as of July 2026 in two layers. The first layer covers the four dominant general-purpose AI assistants, OpenAI's ChatGPT, Anthropic's Claude, Microsoft's Copilot family, and Google's Gemini, each of which now offers some form of healthcare-eligible tier but with materially different scope, exclusions, and administrative burden. The second layer covers the purpose-built clinical AI category: ambient documentation ("AI scribe") platforms, clinical decision-support assistants, and administrative automation tools built specifically for regulated healthcare workflows, where HIPAA-readiness is typically a baseline product requirement rather than an enterprise add-on. Throughout, the report distinguishes marketing claims of "HIPAA compliance" from the actual contractual and technical facts: whether a BAA is available at all, which product surfaces it covers, what configuration a covered entity must complete, and what independent evidence (KLAS rankings, peer-reviewed outcomes studies, OCR enforcement history) exists to support adoption and safety claims. Advisory firms working in life sciences and healthcare IT, including consultancies such as IntuitionLabs that support regulated organizations building compliant data architectures, generally frame HIPAA-compliant AI adoption as a shared-responsibility exercise built around the covered entity's own encryption, access-control, and audit-logging practices, since the same rules "apply to covered entities...and their business associates" who "handle protected data on their behalf," regardless of whether the technology involved is a database, an API, or a generative AI model (Source: intuitionlabs.ai).

ChatGPT (OpenAI): Capabilities, Adoption, and HIPAA Posture

Capabilities

OpenAI's current HIPAA-eligibility documentation lists **ChatGPT for Healthcare**, **ChatGPT for Enterprise with Regulated Workspace**, **ChatGPT FedRAMP**, **ChatGPT for Clinicians**, and the **API with Modified Retention** as products available under an OpenAI BAA. Coverage is feature- and configuration-specific: the same page identifies covered workspace functionality, notes that administrators can enable additional features that are not covered by the BAA for non-PHI uses, and states that API eligibility depends on Modified Retention unless OpenAI specifies otherwise (OpenAI HIPAA Eligible Products and Functionality. OpenAI separately describes ChatGPT for Healthcare as an enterprise product for clinicians, administrators, and researchers with healthcare-specific capabilities and controls (OpenAI. Buyers should use the current eligibility list and their executed agreement, rather than infer coverage from a plan name.

Adoption

OpenAI is one of the three vendors, alongside Microsoft and Epic, that KLAS identified as "the most frequently used or considered" across all healthcare AI categories in its 2026 healthcare IT survey, which "draws on responses from 3,370 individuals across 1,742 unique healthcare organizations" (Source: techtarget.com) and separately found "Microsoft, Epic and OpenAI were the most frequently used or considered vendors across AI categories" (Source: techtarget.com). Separately, the AMA's 2026 physician survey found that among general medical AI use cases, OpenEvidence, a platform built atop foundation models that licenses JAMA content, "estimates that 40% of U.S. doctors use the health AI tool daily" (Source: ama-assn.org), illustrating how much clinical AI usage now flows through vertical products built on top of general-purpose models rather than through ChatGPT directly.

Strengths and Limitations

ChatGPT's strengths are breadth of general reasoning, a mature enterprise administration console, and healthcare-specific offerings with documented BAA pathways. Its principal limitation for healthcare buyers is tier fragmentation: PHI must never enter Business, Plus, Go, or Free tiers, and even within a BAA-covered workspace, "most API services are covered, with a few exceptions" (Source: help.openai.com), requiring administrators to actively restrict feature access rather than assume blanket coverage. Enterprise pricing is also opaque: OpenAI publishes no list price, and third-party procurement research reports that "2026 figures cluster around $60/user/month, a 150-seat minimum, and an annual prepaid commitment" (Source: beam.cloud), a cost structure smaller practices may find prohibitive relative to a purpose-built scribe.

Claude (Anthropic): Capabilities, Adoption, and HIPAA Posture

Capabilities

Anthropic's BAA covers what it calls "HIPAA-ready services": the first-party API and Claude Enterprise, but only after the organization's Primary Owner explicitly activates HIPAA compliance in the "Data and privacy" organization settings, since "Anthropic provides a BAA covering our HIPAA-ready services, such as use of our first-party API or Enterprise plans" (Source: privacy.claude.com), and Anthropic states plainly that "standard Claude Enterprise plans do not include BAA coverage without action from a Primary Owner" (Source: privacy.claude.com). Within a covered Enterprise workspace, chat, Projects, Artifacts, file creation and code execution, Voice, web search, Research, and Skills are covered as "Eligible Services," while Cowork is explicitly flagged as "available to use but feature is not covered under Anthropic's BAA" (Source: privacy.claude.com), and MCP connectors and enterprise search carry the same caveat that "sending data to 3rd parties via this feature isn't covered under Anthropic's BAA" (Source: privacy.claude.com). Claude Code is only covered when zero data retention (ZDR) is enabled; without it, the desktop remote mode is "available to use without ZDR, but this feature is not covered under Anthropic's BAA" and is structurally incompatible with ZDR (Source: privacy.claude.com). Consumer plans, Claude Free, Pro, and Max, are entirely outside Anthropic's BAA structure, as confirmed by Anthropic's separate documentation for consumer products.

Adoption

Anthropic does not publish healthcare-sector-specific seat counts, and independent healthcare-IT survey data (KLAS, HIMSS) treats Claude as a less frequently cited vendor in clinical AI relative to Microsoft, Epic, and OpenAI. Anthropic's pricing structure, Team seats at $20 per seat per month if billed annually (Source: claude.com) rising to $100 per seat for Premium seats bundling Claude Code and Cowork (Source: claude.com), positions it as a general-purpose knowledge-work tool that healthcare organizations are more likely to deploy for administrative, research, and software-

engineering functions than for direct clinical documentation, where purpose-built scribes dominate. At the API layer, Anthropic markets Opus 4.8 as "ideal for complex agentic coding and enterprise work" (Source: claude.com), priced separately from the seat-based products, a cost structure organizations must factor into any custom, PHI-handling application built directly on the API rather than on a seat-based product.

Strengths and Limitations

Anthropic's model card, transparency reporting, and Trust Center compliance artifacts are frequently cited by security teams as unusually detailed for a foundation model vendor, and its API-first BAA structure suits organizations building custom clinical applications atop Claude models. The chief limitation, as with ChatGPT, is scope: Cowork is excluded, while data sent to third parties through connectors and enterprise search is not covered by Anthropic's BAA. Administrators must therefore configure these features carefully rather than treat the entire Enterprise product surface as covered, and Claude Code can handle PHI only in qualified configurations with ZDR enabled. Anthropic's Claude Team Standard seat pricing has itself shifted over time; independent pricing trackers note that Team Standard now runs at the lower of its historical range, while Enterprise has moved toward a hybrid seat-plus-usage model (Source: beam.cloud), underscoring that healthcare buyers should confirm current pricing directly with Anthropic rather than relying on older published figures.

Microsoft Copilot: Capabilities, Adoption, and HIPAA Posture

Capabilities

Microsoft documents HIPAA coverage at the specific service and configuration level. Copilot Studio is covered under Microsoft's HIPAA BAA, while Microsoft 365 Copilot and Copilot Chat can support HIPAA compliance in properly configured implementations (Copilot Studio; Microsoft 365 Copilot data protection. A material exception is public web search: Microsoft states that generated web-search queries are not covered by its Data Protection Addendum or BAA. Organizations using Copilot with PHI should disable web search for those workflows or otherwise ensure PHI cannot be included in generated web queries, and should verify that every licensed workload and connected agent remains within the executed agreement's scope. Separately, **Dragon Copilot**, Microsoft's ambient clinical documentation product built on its 2022 Nuance acquisition, "streamlines clinical documentation with AI, automates healthcare workflows, and empowers clinicians to focus on what matters most" (Source: microsoft.com), marketed directly to health systems as a clinical workflow tool distinct from general Microsoft 365 Copilot.

Adoption

KLAS's 2026 survey found that "in ambient speech, Microsoft and Abridge were the most frequently considered third-party vendors" (Source: techtarget.com), while Microsoft, Epic, and OpenAI were named the three most frequently used or considered vendors across all healthcare AI categories combined. Microsoft's long-standing Azure OpenAI Service partnership with **Epic Systems** was announced when "Microsoft Corp. and Epic on Monday announced they are expanding their long-standing strategic collaboration to develop and integrate generative AI into healthcare" in April 2023 (Source: news.microsoft.com), giving Microsoft a distribution channel that neither OpenAI nor Anthropic can match independently, since Epic's own coverage spans the largest share of U.S. hospital EHR installations.

Strengths and Limitations

Microsoft's core strength is its broad enterprise security and administration ecosystem, but HIPAA coverage must be verified for each specific service and configuration rather than inferred from an organization's use of Microsoft 365 or Azure. Microsoft Copilot Studio is also "compliant with or covered by" HIPAA alongside the Health Information Trust Alliance (HITRUST) Common Security Framework and the Federal Risk and Authorization Management Program (FedRAMP) (Source: learn.microsoft.com), giving buyers multiple overlapping assurance frameworks rather than a single certification to evaluate. Its limitation is complexity: Copilot's HIPAA eligibility depends on which specific Microsoft 365, Copilot Studio, or Azure service an organization licenses and how each is configured, a pattern shared across all major cloud BAAs, per Google's own guidance below, meaning IT teams must maintain an accurate, current inventory of covered versus uncovered services rather than treating "we have a Microsoft BAA" as a blanket assurance.

Google Gemini: Capabilities, Adoption, and HIPAA Posture

Capabilities

Google's HIPAA compliance guide states unambiguously that "Google will enter into Business Associate Agreements with customers as necessary under HIPAA" (Source: cloud.google.com), and this BAA covers Google Cloud's entire infrastructure for the products it lists, explicitly naming **Gemini Enterprise**, **Gemini Code Assist**, **Gemini in BigQuery**, **Gemini in Colab Enterprise**, and **Gemini Notebook Enterprise** among the covered items (Source: cloud.google.com). Critically, Google states that "HIPAA compliance for Google Workspace is covered separately" from the Google Cloud guide (Source: cloud.google.com), and Google's own overview page states plainly that "there is no certification recognized by the US HHS for HIPAA compliance" (Source: cloud.google.com), a shared-responsibility framing that applies equally to every vendor in this report. Google Workspace's own compliance documentation states that customers "who are subject to HIPAA and wish to use certain Google Workspace or Cloud Identity services listed on the HIPAA Included Functionality list must enter a Business Associate Amendment" with Google (workspace.google.com, and separately that "customers who have not signed a BAA with Google must not use PHI in Google Workspace or Cloud Identity services" (workspace.google.com, and specifies that "the help me write, contextual smart replies, and side-panel features are covered as part of Google Workspace with Gemini in HIPAA Included Functionality" (workspace.google.com, meaning Workspace Gemini and Cloud Gemini are governed by two separate compliance determinations even though both carry the same brand name. Google notes that this covered-products list is not static: "Google continues to evaluate the scope of the Included Functionality" and may add products in the future (workspace.google.com, meaning administrators should periodically re-check the list rather than assume it is fixed at the time of contracting. Independent pricing analysis also notes that "Google discontinued the separate Gemini add-on in 2025 and now bundles Gemini into Workspace at no extra per-seat charge," folding what was once a distinct line item into existing Workspace tiers (Source: beam.cloud).

Adoption

Google's healthcare AI presence is smaller in the KLAS 2026 survey's named-vendor rankings than Microsoft, OpenAI, or Epic, though Google Cloud maintains dedicated healthcare and life sciences go-to-market teams and has publicized deployments of its MedLM and Gemini-based clinical search tools with individual health systems on a case-by-case basis, generally under negotiated Google Cloud BAAs rather than a single standardized healthcare product tier comparable to ChatGPT for Clinicians or Dragon Copilot.

Strengths and Limitations

Google's strength is the breadth of its covered-products list once a Cloud BAA is executed, spanning data warehousing (BigQuery), notebook and coding environments, and its enterprise agent platform under a single infrastructure-wide agreement, plus a pricing model: Google states that "other public clouds charge more money for their HIPAA cloud, we do not" (Source: cloud.google.com). Its limitation, shared with every hyperscaler, is that customers must individually disable or otherwise avoid any product not explicitly listed as covered, placing significant operational burden on the customer's own cloud architecture team; Google's guide runs to dozens of product-specific configuration bullet points (Cloud Storage caching, Cloud Speech-to-Text data logging opt-outs, Cloud CDN caching of PHI, and more) that a covered entity must individually verify rather than a single toggle.

Feature Comparison

Table 1 below compares BAA availability, product scope, and pricing across the four general-purpose AI assistants and three widely deployed healthcare-native AI scribe platforms. Figures reflect publicly listed pricing or vendor-stated figures as of July 2026; enterprise clinical AI contracts are frequently negotiated and not publicly disclosed, which is itself noted where applicable.

VENDOR / PRODUCT	BAA AVAILABILITY	WHAT'S COVERED	NOTABLE EXCLUSIONS	INDICATIVE PRICING (AS OF JULY 2026)
ChatGPT (OpenAI)	BAA available for OpenAI's currently listed HIPAA-eligible products	ChatGPT for Healthcare; ChatGPT for Enterprise with Regulated Workspace; ChatGPT FedRAMP; ChatGPT for Clinicians; API with Modified Retention	Administrators may enable additional functionality not covered by the BAA only for uses that do not involve PHI; API eligibility is configuration-specific	Enterprise pricing is not publicly standardized; verify scope and price with OpenAI (official eligibility list
Claude (Anthropic)	Enterprise (Primary Owner must activate) and first-party API (sales-assisted)	Chat, Projects, Artifacts, Voice, web search, Research, Skills within covered Enterprise org	Cowork, Console, and Workbench are not covered; third-party data flows through MCP connectors, Enterprise Search, and Claude in Chrome are not covered; Free/Pro/Max are not covered	Enterprise pricing: contact Anthropic; Team seats $20 to $25/user/month (Source: claude.com) and Premium seats $100 to $125/user/month (Source: claude.com) are not HIPAA-eligible
Microsoft Copilot	Inherits Microsoft 365/Azure enterprise BAA; Copilot Studio explicitly covered (Source: learn.microsoft.com)	Microsoft 365 Copilot, Copilot Studio agents, within covered Microsoft 365/Azure workloads	Coverage depends on the licensed service and configuration; generated web-search queries are not covered by Microsoft's DPA or BAA	Microsoft 365 Copilot $30.00/user/month paid yearly, add-on to eligible M365 subscription (Source: microsoft.com)
Gemini / Google Cloud	Google Cloud BAA covers entire infrastructure for listed products	Gemini Enterprise, Gemini Code Assist, Gemini in BigQuery/Colab Enterprise	Consumer Gemini and Google Workspace Gemini governed separately (workspace.google.com; Pre-GA offerings excluded	HIPAA customers pay same list pricing as non-regulated customers
Abridge (ambient scribe)	Confirm BAA availability and exact eligible-service scope directly in the proposed contract	Ambient clinical documentation across specialties; Best in KLAS for Ambient Speech, 2025 and 2026 (KLAS	Confirm retention, subprocessors, integrations, and any excluded modules	Enterprise pricing is negotiated; verify directly
Nabla Copilot (ambient scribe)	Confirm BAA availability and exact eligible-service scope directly in the proposed contract	Ambient documentation and EHR workflow features; verify current capabilities with Nabla	Confirm retention, subprocessors, integrations, and tier-specific exclusions	Verify current pricing directly with Nabla
Microsoft Dragon Copilot (ambient scribe)	Confirm the product, licensed modules, and eligible-service scope under the executed Microsoft agreement	Clinical documentation and workflow automation for health systems (Microsoft	Confirm connected services, data flows, and exclusions for the planned deployment	Enterprise pricing is negotiated; verify directly

The table illustrates a consistent pattern: every general-purpose vendor gates HIPAA eligibility behind a specific enterprise or API tier and excludes at least one popular consumer-adjacent product from BAA coverage, while purpose-built clinical AI vendors build BAA-readiness into their core offering because PHI handling is their entire business model rather than an enterprise add-on. Pricing for the general-purpose assistants is comparatively transparent and per-seat, while ambient scribe pricing is more fragmented between small-practice subscription tiers (Nabla's $120 per provider per month after a 30-consultation free allotment) and large health-system enterprise agreements that are individually negotiated and not publicly listed, reflecting the fact that scribe vendors sell primarily to health-system procurement teams rather than self-serve individual buyers.

Table 2 below summarizes the OCR enforcement history most relevant to AI procurement decisions, since these are the closest available proxies for what happens when a covered entity treats a vendor as compliant without verifying BAA status and configuration.

CASE	YEAR	ISSUE	PENALTY
Center for Children's Digestive Health (CCDH)	2017	No signed BAA existed with records-storage vendor FileFax from 2003 to 2015 (Source: hhs.gov)	$31,000 settlement (Source: hhs.gov)
Illustrative BAA-related OCR matters	2016–2017	CCDH lacked a signed BAA; North Memorial lacked a BAA and also failed to conduct an organization-wide risk analysis	$31,000 for CCDH; $1,550,000 for North Memorial (HHS: CCDH; HHS: North Memorial
Solara Medical Supplies, LLC	2025	Phishing cybersecurity investigation	$3,000,000 settlement (Source: hhs.gov)
Warby Parker	2025	Hacking/cybersecurity investigation	$1,500,000 civil money penalty (Source: hhs.gov)
Gulf Coast Pain Consultants	2024	Security Rule violations	$1,190,000 civil money penalty (Source: hhs.gov)
Children's Hospital Colorado	2024	Privacy and Security Rule violations	$548,265 civil money penalty (Source: hhs.gov)

Since April 2003, OCR "has received over 374,321 HIPAA complaints and has initiated over 1,193 compliance reviews" (Source: hhs.gov) and "settled or imposed a civil money penalty in 152 cases resulting in a total dollar amount of $144,878,972.00" (Source: hhs.gov). None of these cited cases specifically involve a generative AI vendor as the breached business associate, because OCR's public resolution-agreement log has not yet surfaced a named generative AI case as of this writing; the relevant lesson from the historical record is structural, not AI-specific: covered entities that skip BAA due diligence with any category of technology vendor, cloud storage, billing software, records management, or now AI, face comparable enforcement exposure, and there is no reason to expect OCR's approach to change for generative AI providers.

Performance and Benchmarks

Independent, standardized benchmarking of general-purpose LLMs specifically on HIPAA compliance does not exist because HIPAA compliance is not a model-capability property that a benchmark can score; it is a contractual and configuration state. What does exist is adoption and outcomes data for the clinical-documentation category, which is both the most heavily benchmarked and the most mature commercial segment of healthcare AI. KLAS's 2026 healthcare IT survey found ambient speech the single most widely adopted clinical AI use case, and separately, agentic AI adoption remained negligible: "of more than 3,000 respondents interviewed, only 17 specifically mentioned agentic AI, and just one organization reported active use" (Source: techtarget.com). This gap between marketing enthusiasm for agentic and autonomous AI and actual measured deployment is one of the clearest signals available for procurement teams: the category with the most mature HIPAA-compliant tooling (ambient documentation) is also the category with the most independently verified adoption and outcomes evidence, while newer agentic categories remain largely unproven in production healthcare settings.

On raw model capability, OpenAI states that "GPT-5.2 has been evaluated by licensed physicians across realistic clinical scenarios and benchmarks like HealthBench and GDPval" (Source: help.openai.com) and that in real-world healthcare tasks it "surpasses earlier OpenAI models and performs better than human baselines across every role measured" in GDPval (Source: help.openai.com); this is a vendor-published claim rather than an independently replicated third-party benchmark result, and readers should weigh it accordingly. On accuracy for the ambient-scribe category specifically, vendor-reported figures should also be read cautiously since they are not independently audited: Nabla is intuitive enough that it "doesn't require any training to self-onboard, per the startup, and has 95% accuracy," and "doctors using Nabla Copilot report saving 1.5 hours a day with the tool" (Source: fiercehealthcare.com). Independent peer-reviewed evidence is stronger for outcomes than for raw accuracy: the JAMA Network Open study of Mass General Brigham and Emory Healthcare, surveying more than 1,400 physicians and advanced practice providers, found ambient documentation "was associated with a 21.2% absolute reduction in burnout prevalence at 84 days" at Mass General Brigham, and a 30.7% absolute increase in documentation-related wellbeing at Emory after 60 days; the study authors caveat that response rates were modest and that "these findings likely represent the experience of more enthusiastic users" rather than the full clinician population (Source: massgeneralbrigham.org). A 180-day study of 112 Nuance DAX users and 103 controls in family medicine, internal medicine, and general pediatrics found few statistically significant differences in EHR time, work outside work, note time, appointment completion, same-day closure, note length, or financial metrics, although high DAX users had about 7% lower documentation hours (NEJM AI; Healthcare IT News.

Data Analysis and Evidence

The quantitative backdrop against which HIPAA-compliant AI procurement decisions are made has shifted rapidly. Grand View Research valued the global AI-in-healthcare market at "USD 36.7 billion in 2025 and is projected to grow from USD 50.7 billion in 2026 to USD 505.6 billion by 2033, at a CAGR of 38.9%" (Source: grandviewresearch.com). Within that market, "the software solution segment dominated the market with the largest revenue share of over 46% in 2025" (Source: grandviewresearch.com), while by application, "the robot-assisted surgery segment dominated the market in 2025 with the largest revenue share of over 13%" (Source: grandviewresearch.com). The same report cites a March 2024 Microsoft-IDC study finding "79% of healthcare organizations are presently utilizing AI technology," with a return on investment "realized within 14 months, generating USD 3.20 for every USD 1 invested in AI" (Source: grandviewresearch.com).

On the clinician-adoption side, the AMA's 2026 Physician Survey on Augmented Intelligence, based on responses from nearly 1,700 physicians across specialties, practice settings, and career stages, found 81% now use AI professionally, more than double the 38% recorded when the AMA first polled physicians in 2023. Within that survey, physicians most frequently reported using health AI for "summaries of medical research and standards of care" (Source: ama-assn.org), at 39%, followed by discharge instructions, chart summaries, and billing documentation. "Seven in 10 physicians see AI as a tool to automate tasks that contribute to work-related burnout, and 76% say the technology can help with patient care" (Source: ama-assn.org), though about 40% remain equally excited and concerned about AI, citing patient privacy as a top concern. Physician confidence has grown alongside usage: "more than three-quarters of physicians believe AI improves their ability to care for patients, up from 65% in 2023" (Source: ama-assn.org).

On the compliance-cost side of the ledger, HIPAA civil money penalties are indexed to inflation annually; OCR's current penalty structure includes "civil monetary penalties ranging from $145 to $2,190,294 per violation, depending on the level of culpability" (Source: hipaajournal.com), and "each category of violation carries a separate HIPAA penalty," with OCR determining the specific amount based on factors such as duration, scope, and organizational culpability (Source: hipaajournal.com). OCR enforcement shows that missing or deficient BAAs can contribute to settlements, but the facts must be read case by case. CCDH paid $31,000 in a matter centered on the lack of a signed BAA, while North Memorial's $1.55 million settlement involved both a missing BAA and failure to perform an organization-wide risk analysis (HHS: CCDH; HHS: North Memorial. Across all HIPAA enforcement categories, not limited to AI or even to business-associate issues, OCR reports resolving 99% of cases since April 2003; the agency "has investigated and resolved over 31,191 cases by requiring changes in privacy practices and corrective actions" (Source: hhs.gov), "in another 15,561 cases, our investigations found no violation had occurred" (Source: hhs.gov), and in a further 67,873 cases "OCR intervened early and provided technical assistance to HIPAA covered entities, their business associates, and individuals exercising their rights" without a formal investigation (Source: hhs.gov), and 152 cases resulted in civil money penalties totaling $144,878,972.

Independent third-party validation is emerging as its own market signal alongside these regulatory figures. HITRUST, the healthcare-focused security assurance body, now offers a distinct AI Risk Management Assessment "aligned with ISO/IEC 23894:2023 and NIST AI RMF" (Source: hitrustalliance.net), using "51 curated controls mapped to the NIST AI RMF and ISO/IEC 23894:2023" that also incorporate OWASP's Top 10 for large language models (Source: hitrustalliance.net), and is designed for "organizations integrating AI into operations or developing AI tech that need to manage risks without pursuing certification" that are "well-suited for companies navigating complex regulatory environments or those building AI systems for high-risk industries" (Source: hitrustalliance.net). HITRUST separately reports "a 99.62% breach-free rate among HITRUST-certified environments" (Source: hitrustalliance.net). NIST's own AI RMF, first released January 26, 2023, "is intended for voluntary use and to improve the ability to incorporate trustworthiness considerations" into AI design and deployment (Source: nist.gov), and "the Framework was developed through a

consensus-driven, open, transparent, and collaborative process" involving multiple public workshops and comment periods (Source: nist.gov). In March 2023, "NIST launched the Trustworthy and Responsible AI Resource Center, which will facilitate implementation of, and international alignment with, the AI RMF" (Source: nist.gov), meaning adoption of NIST or HITRUST frameworks is a best-practice signal for vendor due diligence rather than a legal substitute for an executed BAA.

Case Studies and Real-World Examples

Mass General Brigham and Emory Healthcare: Measuring Burnout Reduction

Mass General Brigham launched its ambient documentation program as a proof-of-concept pilot involving 18 physicians in July 2023, and by April 2025 "the technologies have been made available to all Mass General Brigham physicians, with more than 3,000 providers routinely using the tools" (Source: massgeneralbrigham.org). A study led by Mass General Brigham researchers, published in *JAMA Network Open* and drawing on surveys of more than 1,400 physicians and advanced practice providers across Mass General Brigham and Atlanta's Emory Healthcare, found ambient documentation technologies were associated with a 21.2% absolute reduction in burnout prevalence at Mass General Brigham after 84 days, and a 30.7% absolute increase in documentation-related wellbeing at Emory after 60 days. Rebecca Mishuris, MD, Mass General Brigham's chief medical information officer and a co-senior study author, said "there is literally no other intervention in our field that impacts burnout to this extent" (Source: massgeneralbrigham.org). Physician burnout "affects more than 50% of U.S. doctors and has been linked to time spent in electronic health records, particularly after hours" (Source: massgeneralbrigham.org), giving the burnout-reduction figure direct operational significance beyond a pure satisfaction metric.

The Permanente Medical Group and Nabla: Scaling to 10,000 Physicians

In October 2023, Nabla, maker of an autonomous clinical AI scribe, partnered with the Permanente Medical Group to "roll out its tool to 10,000 doctors in Northern California" (Source: fiercehealthcare.com), a group that provides care to Kaiser Permanente members and is described as the largest physician-led medical group in the nation. The deployment followed a successful pilot and was designed to "reduce administrative burden to allow doctors to focus on patient care" (Source: fiercehealthcare.com). Nabla's tool captures physician-patient encounters, in person or virtual, across any specialty, using a combination of proprietary and industry large language models. Nabla CEO Alex Lebrun said the partnership would "help bring the organization to the forefront of AI-powered care delivery transformation" (Source: fiercehealthcare.com).

UCSF Health and Ambience Healthcare: Piloting a New AI Scribe Category

UCSF Health contracted with Ambience Healthcare in November 2023 to deploy an AI scribe that listens to clinical visits and drafts documentation in real time, "initially launched with a group of 100 physicians in ambulatory practice" and in the pediatric emergency departments in Oakland and Mission Bay (Source: it.ucsf.edu). Ambience Healthcare CEO and co-founder Mike Ng said the goal was to "deploy AI tools that relieve clinicians of their documentation burden and enable them to fully focus on what they do best" (Source: it.ucsf.edu). UCSF's Faculty Practice Office explicitly framed the rollout as governed through its AI Governance Committee and Medical Informatics team, working "to ensure we are rolling this technology out responsibly and successfully" (Source: it.ucsf.edu), noting that Ambience was "the first 'ambient scribe' technology UCSF Health is trialing," and that "the Faculty Practice and Health IT anticipate trialing other competing vendors in this space" as the category matures (Source: it.ucsf.edu). Dr. Tom Chi, an attending urologist involved in the rollout, said "scribes have had an incredibly positive impact on our doctors and patients at UCSF" (Source: it.ucsf.edu), illustrating a covered entity pairing a purpose-built HIPAA-eligible vendor with an internal governance layer rather than treating vendor BAA availability as sufficient on its own.

Epic Systems and Microsoft: Embedding Generative AI Directly in the EHR

In April 2023, Microsoft and Epic announced an expanded strategic collaboration "to develop and integrate generative AI into healthcare" by combining Azure OpenAI Service directly with Epic's EHR (Source: news.microsoft.com). This case matters for HIPAA-compliant AI procurement because it demonstrates a distinct architectural pattern from either a general-purpose chatbot or a stand-alone scribe: PHI flows through Azure's HIPAA-covered infrastructure but is invoked from inside a workflow the covered entity already governs (Epic's EHR), rather than requiring clinicians to copy PHI into a separate external tool, which meaningfully reduces the surface area for accidental non-compliant use compared to a clinician independently deciding whether a given consumer AI tool is safe to paste patient information into.

Center for Children's Digestive Health: A Documented Enforcement Case

The 2017 OCR settlement with the Center for Children's Digestive Health (CCDH), a small, for-profit pediatric subspecialty practice operating seven Illinois clinics, illustrates the enforcement risk that applies equally to any AI vendor relationship lacking a BAA. OCR's compliance review found that "neither party could produce a signed Business Associate Agreement" prior to October 12, 2015, a gap of roughly twelve years after CCDH began disclosing PHI to records-storage vendor FileFax, Inc. in 2003 (Source: hhs.gov). CCDH "has paid the U.S. Department of Health and Human Services" $31,000 to settle potential violations and agreed to a corrective action plan (Source: hhs.gov). The case predates generative AI entirely, but the underlying failure mode, treating a vendor relationship as informally trustworthy without executing the legally required paperwork, is precisely the mistake healthcare organizations risk repeating when a clinician or department adopts a consumer AI tool without confirming BAA coverage first.

Implications and Future Directions

Several structural trends will shape HIPAA-compliant AI procurement over the next several years. First, the KLAS finding that agentic AI adoption remains negligible suggests the next wave of HIPAA-compliance complexity will center on multi-step agentic workflows, tools that autonomously read PHI, take actions, and call other tools or connectors, rather than the single-turn documentation and chat use cases that dominate compliant deployment today. Anthropic's structural approach is instructive here: its BAA covers eligible Enterprise features, but not data sent to third parties through connectors or Enterprise Search. That distinction means organizations must separately assess downstream tools and data flows rather than infer that every integration is covered whenever the underlying chat product is covered.

Second, the regulatory environment beyond core HIPAA is in flux. In January 2025, HHS's Office for Civil Rights issued a "Dear Colleague" letter confirming "it will enforce Section 1557's nondiscrimination protections to the use of AI" in patient care decision-support tools, effective from May 1, 2025 for the duty to identify and mitigate discrimination risk (Source: reedsmith.com). OCR's letter illustrated the risk with a concrete example: "a hospital's advanced emergency triage system must not discriminate against individuals with disabilities by failing to consider how an individual's disability or age could affect the assessment factors" (Source: reedsmith.com), and among its recommended mitigations, regulated organizations should "establish written policies and procedures governing AI tool use, monitor AI tools' impacts, and develop ways to address discrimination complaints" (Source: reedsmith.com), "train staff on the proper use of AI tools and audit their performance in real-world scenarios to ensure compliance" (Source: reedsmith.com), and "disclose to patients the use of AI in patient care decision support tools that pose a risk of discrimination" (Source: reedsmith.com). However, that broader Section 1557 rule has itself faced significant legal challenges: prior to the rule taking effect, "a federal district court judge in Mississippi issued a nationwide preliminary injunction, temporarily preventing enforcement of the rule" as to its gender-identity provisions (Source: apexbg.com), and on October 22, 2025, that same court "permanently vacated the portions of the Section 1557 regulations that interpreted sex discrimination to include gender identity" (Source: apexbg.com). Legal analysis of the underlying statute adds nuance relevant to AI vendor relationships: "while most employer health plans will not be considered covered entities, to the extent a group health plan receives federal financial assistance, it would be subject to Section 1557" (Source: apexbg.com), "a third-party administrator for a self-funded group health plan may be considered a covered entity where it is affiliated with a health insurance issuer that receives funding" (Source: apexbg.com), and, more generally, "whether Section 1557 applies to a particular entity is a fact-specific analysis" (Source: apexbg.com), meaning even entities that do not think of themselves as traditional "covered entities" may still owe Section 1557 obligations when they adopt AI tools. This litigation targeted the gender-identity interpretation specifically rather than the AI clinical-decision-support provisions, but the overall instability of Section 1557 rulemaking means healthcare organizations should treat AI nondiscrimination compliance as a genuinely unsettled area requiring ongoing legal monitoring rather than a fixed rule to check off once, distinct from the comparatively stable core HIPAA Privacy, Security, and Breach Notification Rules that this report's BAA analysis rests on.

Third, third-party assurance frameworks are converging around AI-specific extensions of existing healthcare security standards rather than entirely new regimes: HITRUST's AI Risk Management Assessment explicitly maps to the pre-existing NIST AI RMF and ISO/IEC 23894:2023 rather than inventing new controls from scratch, suggesting that organizations already familiar with HITRUST CSF (Common Security Framework) certification for their broader HIPAA program face a smaller incremental learning curve for AI-specific assurance than organizations building AI governance from zero. Fourth, KLAS's observation that EHR vendors including Epic, Oracle Health, athenahealth, and eClinicalWorks are rolling out native ambient speech tools that "some organizations prefer for easier integration and less vendor complexity" (Source: techtarget.com) points toward consolidation: standalone AI vendors that do not integrate cleanly with a health system's existing EHR and its existing BAA relationship face a structural disadvantage against native EHR-vendor AI features, even when the standalone tool's underlying model quality is comparable or superior. Fifth, buyer guidance on evaluating any medical AI tool converges on the same caution regardless of vendor size: "some free tools invest heavily in compliance and data protection, while some paid tools fail to meet the security standards required in today's practices," meaning price and HIPAA-readiness are not reliably correlated and every product still requires individual BAA verification (Source: blog.doximity.com). For organizations that lack in-house expertise to evaluate this fast-moving landscape, consultancies working across life sciences and health IT compliance emphasize that developing

systems handling health data in the United States "means navigating the stringent requirements of the Health Insurance Portability and Accountability Act" regardless of whether the underlying technology is a traditional database or a generative AI system (Source: intuitionlabs.ai), making HIPAA-compliant AI adoption fundamentally an architecture and governance exercise rather than a one-time vendor selection decision.

Frequently Asked Questions (FAQs)

Is ChatGPT HIPAA compliant? Not by default. OpenAI lists ChatGPT for Healthcare, ChatGPT for Enterprise with Regulated Workspace, ChatGPT FedRAMP, ChatGPT for Clinicians, the API with Modified Retention, and API FedRAMP with Modified Retention as HIPAA-eligible products available with a BAA. Ordinary Enterprise or Edu workspaces should not be assumed covered unless provisioned as an eligible offering under the executed BAA.

Is Microsoft Copilot HIPAA compliant? Microsoft documents Copilot Studio as covered under its HIPAA BAA, while Microsoft 365 Copilot can support HIPAA compliance for prompts and responses in properly configured eligible implementations. Coverage is service- and configuration-specific, and generated public-web-search queries are not covered by Microsoft's DPA or BAA; buyers must verify their executed agreement and disable or control web search for PHI workflows (Microsoft.

Is Claude AI HIPAA compliant? Only for Claude Enterprise (after the Primary Owner explicitly activates HIPAA compliance) and the first-party API under a sales-assisted BAA; Claude Free, Pro, Max, Console, Workbench, and Cowork are not covered.

Is Google Gemini HIPAA compliant? Gemini Enterprise, Gemini Code Assist, and related Google Cloud Gemini products are covered under an executed Google Cloud BAA, but consumer Gemini and Gemini inside Google Workspace are governed by separate agreements and are not automatically covered.

What is a HIPAA-compliant chatbot for healthcare providers? In practice, it is either a general-purpose AI assistant used strictly within a BAA-covered tier and feature set, or a purpose-built clinical AI product (an ambient scribe, clinical decision-support tool, or patient-messaging assistant) engineered from inception around PHI handling, with BAA availability as a standard part of its enterprise contract rather than a special add-on.

What is an AI business associate agreement in healthcare (AI BAA)? It is the same legal instrument HIPAA has always required of any vendor handling PHI on a covered entity's behalf, applied to an AI vendor: a written contract establishing permitted uses of PHI, the business associate's security obligations, breach-notification duties, and termination and data-return provisions (Source: hollandhart.com). There is no "AI-specific" BAA template mandated by HIPAA; the same legal requirements that govern a billing vendor or cloud host apply to a large language model provider.

What are the best HIPAA-compliant AI tools for clinicians? Among general-purpose assistants, ChatGPT for Clinicians and Microsoft 365 Copilot (within a covered Microsoft 365 tenant) currently offer the lowest-friction paths to individual clinician use with BAA coverage. Among purpose-built tools, Abridge, ranked number one Best in KLAS for Ambient AI for the second straight year in 2026, and Nabla, deployed to 10,000 Permanente Medical Group physicians, have the strongest combination of independent ranking data and large-scale, named production deployments.

Are there HIPAA-compliant generative AI vendors built specifically for healthcare? Yes; beyond Abridge, Nabla, and Microsoft's Dragon Copilot covered in this report, the broader category includes Suki, Ambience Healthcare, Hippocratic AI, Glass Health, Notable Health, and Doximity, whose own clinician-facing marketing states its AI products are "easy to use, free, HIPAA-compliant, and secure, helping doctors, nurse practitioners, and PAs integrate AI" into their workflow (Source: blog.doximity.com), a claim consistent with the broader pattern of healthcare-specific vendors treating HIPAA-readiness and BAA availability as a core, rather than optional, feature.

Conclusion

There is no universal answer to "is this AI tool HIPAA compliant," only a defensible process for determining whether a specific product tier, under a specific signed BAA, configured a specific way, is safe for a specific PHI-handling use case. Every major general-purpose vendor examined in this report, OpenAI, Anthropic, Microsoft, and Google, now offers a path to HIPAA eligibility, but each gates that eligibility behind enterprise or API tiers, requires explicit contractual and configuration steps, and excludes at least some popular consumer-facing or agentic features even within an otherwise covered account. Purpose-built clinical AI vendors, led by ambient documentation tools that KLAS identified as the most widely adopted clinical AI use case, generally build BAA-readiness into their core product rather than treating it as an enterprise add-on, and independent evidence, from KLAS's Best in KLAS rankings to the JAMA Network Open burnout study, offers a stronger evidentiary base for procurement decisions than vendor marketing claims alone.

The financial and reputational stakes of getting this wrong are well established even outside the AI context: OCR's enforcement record includes a $31,000 matter centered on a missing BAA and a $1.55 million matter involving both a missing BAA and an inadequate organization-wide risk analysis, and civil penalties can now reach $2,190,294 annually per violation category. Healthcare organizations evaluating AI tools should therefore treat the question in three concrete steps: confirm a BAA is actually available for the specific product tier under consideration, obtain and review that BAA's explicit list of covered and excluded features rather than assuming blanket coverage, and configure the product, and the surrounding clinical workflow, to keep PHI strictly within covered features. Applied consistently, that process converts "is this AI tool HIPAA compliant" from an unanswerable marketing question into a documented, auditable, and defensible compliance decision.

Tags: hipaa compliant ai, healthcare ai tools, chatgpt hipaa, claude ai hipaa, microsoft copilot hipaa, gemini hipaa, business associate agreement, ai scribe healthcare, clinical ai compliance

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

Leading Claude & ChatGPT AI Enablement and Training: Help biotech and pharmaceutical teams adopt Claude and ChatGPT through practical training, governed workflows, use-case development, and implementation designed for regulated environments.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at intuitionlabs.ai/contact for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by Adrien Laurent, a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.