Is ChatGPT HIPAA Compliant? What Practices Must Know in 2026

Published July 19, 2026 40 min read

Executive Summary

Standard, consumer facing **ChatGPT** (the Free, Plus, and Team tiers, along with self serve **ChatGPT Business**) is not **HIPAA** compliant and cannot lawfully be used to process **Protected Health Information (PHI)**, because OpenAI will not sign a Business Associate Agreement (BAA) for those products (Source: help.openai.com) (Source: www.hipaajournal.com). That baseline has not changed even as OpenAI's healthcare ambitions have accelerated. On January 8, 2026, OpenAI launched **OpenAI for Healthcare**, a suite anchored by **ChatGPT for Healthcare**, which is already rolling out at institutions including Boston Children's Hospital, Cedars-Sinai Medical Center, HCA Healthcare, and UCSF (Source: openai.com) (Source: www.hipaajournal.com). OpenAI now recognizes six HIPAA eligible products: ChatGPT for Healthcare, ChatGPT for Enterprise with Regulated Workspace, ChatGPT FedRAMP, ChatGPT for Clinicians, and the API with Modified Retention, plus its FedRAMP variant, each covered by OpenAI's Business Associate and Healthcare Addendum once a BAA is executed (Source: help.openai.com).

Crucially, none of this makes ChatGPT compliant "out of the box." HIPAA Journal states plainly that "ChatGPT for Healthcare is not HIPAA compliant \'out of the box'" and that compliance depends on how the product is deployed, configured, governed, and paired with workforce training (Source: www.hipaajournal.com). A signed BAA is the legal prerequisite; the technical and administrative safeguards required by the HIPAA Security Rule are a separate, ongoing obligation that sits with the covered entity (Source: www.hhs.gov). Individual clinicians have a narrower, self-serve path: ChatGPT for Clinicians, free at launch for verified U.S. physicians, nurse practitioners, physician assistants, and pharmacists, includes an in-product BAA flow, but users are explicitly warned not to enter PHI "unless a BAA is in place and you are authorized to sign a BAA for your account" (Source: help.openai.com).

The stakes for getting this wrong are substantial. HIPAA civil penalties range from $145 to $2,190,294 per violation category depending on culpability tier (Source: www.hipaajournal.com), and a USC Schaeffer Center analysis warns that "once you enter something into ChatGPT, it is on OpenAI servers and they are not HIPAA compliant. That's the real issue, and that is, technically, a data breach" (Source: schaeffer.usc.edu). Physician adoption of AI has nearly doubled, reaching 81% in the American Medical Association's 2026 survey of nearly 1,700 physicians, yet 86% of doctors still rate data privacy as a critical concern for broader adoption (Source: www.ama-assn.org) (Source: www.ama-assn.org).

This report walks through what HIPAA actually requires, how ChatGPT's product tiers differ in BAA eligibility, the mechanics of obtaining and operationalizing an OpenAI BAA, and how competing HIPAA eligible options from Microsoft Azure, Google Cloud, Amazon Web Services (AWS), and specialized clinical AI vendors such as Nabla and Abridge compare. It also surveys the fast moving landscape of state laws in Texas, Illinois, California, and Maine that now require disclosure or restrict autonomous AI decision making in care settings, with dozens more bills advancing in 2026 (Source: www.beckershospitalreview.com) (Source: www.beckershospitalreview.com), and a pending federal proposal that would tighten HIPAA Security Rule obligations regardless of vendor (Source: intuitionlabs.ai). The bottom line for practices in 2026: generic ChatGPT remains off limits for PHI, a properly contracted and configured HIPAA eligible OpenAI product can support compliant use, and no software purchase, from OpenAI or any competitor, substitutes for the risk analysis, workforce training, and governance that HIPAA itself requires.

Introduction and Background

The question "is ChatGPT HIPAA compliant" has a deceptively simple answer that has grown more nuanced as OpenAI's product line has expanded. Through most of 2023, 2024, and 2025, the honest answer for anyone asking about the consumer chatbot was an unambiguous no. OpenAI updated its data usage policy in March 2023 to state that it "will not use data submitted by customers to train or improve its models unless customers expressly opt-in" and that it "also will enter into business associate agreements in support of applicable customers' compliance" with HIPAA, but that commitment applied only to specific enterprise and API arrangements, not to the ordinary consumer product most clinicians had already begun experimenting with (Source: www.sheppard.com). Several large employers outside healthcare, including JPMorgan Chase and Verizon Communications, restricted employee use of ChatGPT altogether over fears that proprietary information typed into the chatbot could be exposed or retained (Source: www.sheppard.com).

That picture changed materially in January 2026. OpenAI introduced OpenAI for Healthcare, a set of products it describes as designed "to help healthcare organizations deliver more consistent, high-quality care for patients, while supporting their HIPAA compliance requirements," arriving alongside evidence that physician demand for AI tools had outpaced what compliant, enterprise-sanctioned deployments could offer (Source: openai.com). This was not a cosmetic rebrand: many clinicians, an American Medical Association survey found, still have to rely on their own tools because their organizations are not adopting AI fast enough, often due to the constraints of regulated environments. That gap between individual clinician appetite and institutional compliance infrastructure is exactly why this question matters, and why the answer now requires distinguishing between several materially different ChatGPT products rather than treating "ChatGPT" as one monolithic tool.

This report is written for practice owners, health system compliance leaders, and clinicians evaluating whether, and how, ChatGPT and its OpenAI branded sibling products can fit into a HIPAA-governed workflow as of July 2026. It draws on OpenAI's own policy and product documentation, U.S. Department of Health and Human Services (HHS) guidance, peer-reviewed literature on generative AI risk in medicine, state legislative trackers, and named case studies from health systems that have already deployed generative AI at scale. The analysis proceeds in stages: first, a grounding in what HIPAA itself actually requires of any technology vendor; second, a walk through ChatGPT's distinct product tiers and their differing BAA eligibility; third, the practical mechanics of obtaining and operationalizing an OpenAI BAA; fourth, a survey of HIPAA-eligible alternatives from major cloud providers and specialized clinical AI vendors; and finally, quantitative evidence, real-world case studies, and a discussion of where the regulatory and competitive landscape is heading. Where sources disagree, for example on which ChatGPT tiers have historically qualified for a BAA, this report states the discrepancy openly rather than papering over it. As a life sciences and AI consultancy that advises healthcare and pharmaceutical organizations on regulated data infrastructure without itself selling a competing chatbot product, IntuitionLabs approaches this question from the vantage point of an implementation advisor: the technology question of whether OpenAI can sign a BAA is usually easier to answer than the governance question of whether an organization has actually configured, trained for, and audited compliant use.

HIPAA Fundamentals: What Compliance Actually Requires

Before evaluating any AI vendor, it is worth being precise about what HIPAA (the Health Insurance Portability and Accountability Act of 1996) actually obligates, because "HIPAA compliant" is frequently used loosely in marketing copy despite having a specific legal meaning. HIPAA's Privacy Rule applies directly only to **covered entities**: health plans, health care clearinghouses, and health care providers who transmit health information electronically in connection with certain transactions (Source: www.law.cornell.edu). The regulation defines a health care provider broadly as "any other person or organization who furnishes, bills, or is paid for health care in the normal course of business," a definition wide enough to sweep in solo practices and small clinics alongside hospital systems (Source: www.law.cornell.edu). Because most providers rely on outside vendors to perform functions involving Protected Health Information (PHI), HIPAA also regulates **business associates**: any person or entity that, on behalf of a covered entity, "creates, receives, maintains, or transmits protected health information" for a regulated function or activity (Source: www.law.cornell.edu) (Source: www.law.cornell.edu). HHS's own guidance lists examples ranging from third-party claims administrators to "an independent medical transcriptionist that provides transcription services to a physician," an analogy directly relevant to AI dictation and documentation tools, and notes that business associate services can also include "legal; actuarial; accounting; consulting; data aggregation; management; administrative; accreditation;

and financial" functions, categories broad enough to plausibly cover many generative AI use cases (Source: www.hhs.gov) (Source: www.hhs.gov). HHS also confirms that "a covered health care provider, health plan, or health care clearinghouse can be a business associate of another covered entity," a nuance that matters as health systems increasingly act as data intermediaries for one another (Source: www.hhs.gov).

Whenever a covered entity discloses PHI to a business associate, the Privacy Rule requires that the covered entity first obtain "satisfactory assurances" that the business associate will safeguard the information, documented through a **Business Associate Agreement (BAA)** (Source: www.hhs.gov). Without an executed BAA, disclosing PHI to a vendor, including pasting a clinical note into a chatbot, is itself an impermissible disclosure, regardless of whether the vendor's infrastructure happens to be secure. If a covered entity later discovers a business associate is not honoring that agreement and cannot resolve the breach, HHS guidance requires escalation: covered entities must "take reasonable steps to cure the breach or end the violation, and if such steps are unsuccessful, to terminate the contract or arrangement," reporting unresolved problems "to the Department of Health and Human Services (HHS) Office for Civil Rights (OCR)" (Source: www.hhs.gov). Separately, the HIPAA Security Rule requires "administrative, physical, and technical safeguards to ensure the confidentiality, integrity, and availability of electronic protected health information," a standard that applies to covered entities and business associates alike and exists independently of whatever contractual assurances a BAA provides (Source: www.hhs.gov). HHS also maintains a public Security Risk Assessment Tool whose "features make it useful in assisting small and medium-sized health care practices and business associates as they perform a risk assessment," underscoring that the risk-analysis obligation applies just as much to a small practice piloting an AI scribe as it does to a large hospital system (Source: www.hhs.gov).

This distinction, between having a signed BAA and actually being configured for compliant use, is the single most important concept for evaluating any AI vendor, including OpenAI. As one industry guide focused on healthcare startups summarizes, "HIPAA does not offer a government-issued 'certification'" and compliance is instead demonstrated through passing audits, responding to OCR inquiries, or negotiating settlements after the fact (Source: intuitionlabs.ai). Several practical elements matter most for AI tools specifically:

- **PHI identifiers**: HIPAA recognizes 18 specific identifiers, patient names (including nicknames), dates of birth, admission or discharge dates, and geographic detail smaller than a state among them, any of which converts otherwise generic clinical text into regulated PHI (Source: schaeffer.usc.edu).
- **Minimum necessary standard**: covered entities must limit any use or disclosure of PHI, including to an AI system, "to the minimum necessary to accomplish the intended purpose" (Source: www.sheppard.com).
- **De-identification as an escape valve**: PHI that has been de-identified using a HIPAA Privacy Rule-permitted method is no longer PHI and can be processed through non-BAA tools, provided the de-identification is genuine rather than superficial (Source: www.hipaajournal.com).
- **Business associate liability chains**: a covered entity can be held liable for a business associate's HIPAA violations if it "knew, or by exercising reasonable diligence, should have known" of a pattern of noncompliance and failed to act (Source: www.hipaajournal.com).

These fundamentals apply identically whether the technology in question is a claims clearinghouse, a transcription service, or a large language model. What differs with generative AI is the ease with which PHI can be pasted into an unauthorized tool by a single clinician acting in good faith, which is precisely the failure mode explored in the case studies below.

ChatGPT's Compliance Tiers: From Consumer Chatbot to HIPAA-Eligible Enterprise Product

"ChatGPT" is not a single product from a compliance standpoint; it is a family of offerings with materially different contractual terms. HIPAA Journal's summary captures the starting point: "Generic ChatGPT services are not HIPAA compliant and cannot be used in a HIPAA-compliant manner because they do not offer the safeguards and Business Associate Agreements required under the HIPAA Security and Privacy Rules to protect PHI" (Source: www.hipaajournal.com). A separate healthcare-focused legal risk analysis reaches the same conclusion from a different angle: "ChatGPT is not HIPAA compliant, which poses a risk for healthcare use," since the way the consumer product handles data "lacks transparency" around PHI security (Source: www.paubox.com). However, OpenAI has steadily carved out a set of "HIPAA eligible" products, meaning products for which OpenAI will execute a BAA once the customer meets its eligibility and configuration requirements. As of 2026, OpenAI's own support documentation lists six such products: 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 (Source: help.openai.com).

Table 1 below summarizes how the major ChatGPT product tiers differ in BAA eligibility, cost structure, and permitted use of PHI as of July 2026.

PRODUCT TIER	BAA AVAILABLE?	WHO CAN USE IT	PHI PERMITTED?	KEY DETAIL
ChatGPT Free / Plus / Team	No	Any individual consumer or small team subscriber	No	OpenAI "won't offer a BAA for ChatGPT Business," and generic tiers lack HIPAA-required safeguards and audit trails (Source: help.openai.com)
ChatGPT Enterprise (self-serve)	Only via sales-managed accounts, and disputed across sources	Organizations with a dedicated OpenAI account team	No, unless a BAA is separately negotiated	HIPAA Journal states OpenAI "will not enter into a Business Associate Agreement for the Free, Plus, Team, or Enterprise versions of its product" (Source: www.hipaajournal.com)
ChatGPT for Healthcare	Yes, case-by-case enterprise sales process	Healthcare organizations (hospitals, health systems); no self-serve signup	Yes, once BAA executed and workspace configured	Launched January 8, 2026; already deployed at Boston Children's Hospital, Cedars-Sinai, HCA Healthcare, and UCSF (Source: www.hipaajournal.com)
ChatGPT for Clinicians	Yes, in-product BAA flow for eligible individuals	Free for verified U.S. physicians (MD/DO), nurse practitioners, physician assistants, and pharmacists (Source: help.openai.com)	Only after individual clinician signs BAA under Settings > Agreements	Image generation is unsupported, and PHI sharing with third-party GPTs or connected apps is separately restricted
OpenAI API with Modified Retention	Yes, via baa@openai.com request process	Developers and healthcare software vendors building custom applications	Yes, across a defined list of HIPAA-eligible endpoints once BAA and Modified Retention are provisioned	HIPAA eligibility is "contingent on Customer's account being provisioned with Modified Retention" (Source: help.openai.com)
ChatGPT Health (consumer)	No	Any consumer, outside the EEA, Switzerland, and the UK at initial rollout	No	A consumer wellness product used by "over 230 million people globally" every week to ask health questions, but "not intended for diagnosis or treatment" and governed by consumer terms, not HIPAA (Source: openai.com)

Two things stand out in this comparison. First, there is a genuine discrepancy across sources on how far self-serve ChatGPT Enterprise eligibility for a BAA actually extends: OpenAI's own help center article on API BAAs still describes sales-managed Enterprise and Edu accounts as eligible, stating that "only ChatGPT Enterprise or Edu customers that have a sales-managed account are eligible for a BAA for ChatGPT at this time" (Source: help.openai.com), while HIPAA Journal's 2026 analysis implies ChatGPT for Healthcare has effectively superseded Enterprise as the path for regulated healthcare use. Practices should treat this as an area requiring direct confirmation from OpenAI's sales team rather than relying on either source alone. Second, ChatGPT for Enterprise, ChatGPT Edu, and ChatGPT for Healthcare all "successfully completed a SOC 2 Type 2 audit," which is a security attestation distinct from, and not a substitute for, HIPAA compliance itself (Source: openai.com).

ChatGPT for Healthcare in Practice

ChatGPT for Healthcare packages several capabilities specifically for regulated clinical environments: "evidence retrieval with transparent citations," drawing from "millions of peer-reviewed research studies, public health guidance, and clinical guidelines," alongside role-based access controls, single sign-on, audit logs, and customer-managed encryption keys (Source: openai.com). On the retrieval side, OpenAI "configures HIPAA eligible ChatGPT workspaces to use OpenAI's search index" rather than routing queries to a third-party search provider, which keeps external query traffic inside the BAA-covered environment (Source: help.openai.com). Critically, content shared with ChatGPT for Healthcare is not used to train models, addressing one of the most common fears clinicians have about pasting notes into a chatbot (Source: openai.com). Yet HIPAA Journal notes there is "no self-serve option for healthcare organizations to subscribe to the ChatGPT for Healthcare product," meaning every deployment goes through OpenAI's enterprise sales channel, and BAAs are prepared "on a case-by-case basis" with organizations "only allowed to preview Agreements under an NDA" (Source: www.hipaajournal.com) (Source: www.hipaajournal.com).

ChatGPT for Clinicians: A Narrower, Individual Path

For clinicians whose institutions have not yet adopted ChatGPT for Healthcare, OpenAI created a parallel, individual-level path. ChatGPT for Clinicians is "designed for individual use, including the ability to sign a Business Associate Agreement," and eligible clinicians can complete verification through a third-party provider using their National Provider Identifier (NPI) (Source: help.openai.com). The product includes "trusted clinical search with citations," "deep research across medical literature," and documentation support for drafting referral letters, prior authorization letters, and patient instructions. But its scope is deliberately limited: it covers only the individual clinician's own workspace, does not extend to organization-wide deployment, and OpenAI's guidance is explicit that clinicians "should independently verify information as appropriate" since the tool is "intended to support, not replace, professional judgment" (Source: help.openai.com).

The API with Modified Retention

For software vendors and health systems building custom applications, the OpenAI API remains the most flexible route. HIPAA eligibility hinges on a specific technical configuration: the customer's account must be provisioned with "Modified Retention," after which a defined list of endpoints, including `/v1/chat/completions`, `/v1/responses`, `/v1/embeddings`, and `/v1/audio/transcriptions`, can process PHI once a BAA is executed. Unlike ChatGPT Enterprise, a full enterprise agreement is not a prerequisite: "an enterprise agreement is not required to sign a BAA" for API services, making this the most accessible HIPAA-eligible path for smaller developers and digital health startups (Source: help.openai.com). This is precisely the mechanism that companies like Abridge, Ambience, and EliseAI use to build ambient documentation and appointment-scheduling products on top of OpenAI's models without those companies being "ChatGPT" in any consumer-facing sense.

Obtaining and Operationalizing a Business Associate Agreement with OpenAI

Understanding which products are eligible is only the first step; the process of actually obtaining and operationalizing a BAA is where most compliance failures originate. For API customers, the process begins by emailing baa@openai.com with details about the company and its intended use case, after which OpenAI reviews the request and, if approved, executes the agreement before PHI can be processed (Source: help.openai.com). For ChatGPT Enterprise, Edu, or ChatGPT for Healthcare, the equivalent path runs through OpenAI's enterprise sales channel rather than a self-serve form.

Executing the BAA is a legal precondition, not a technical one, and OpenAI is explicit that a signed agreement alone does not configure anything. Organizations must separately "deploy ChatGPT for Healthcare in a HIPAA-eligible environment, configure the administrative controls for identity management, audit-logging, and data retention, and apply role-based permissions" (Source: www.hipaajournal.com). Once configured, workspace admins gain access to an audit log of conversations through OpenAI's Enterprise Compliance API, though "authorized OpenAI employees will only ever access your conversations for the purposes of resolving incidents, recovering end user conversations with your explicit permission, or where required by applicable law" (Source: openai.com). By default, deleted conversations are removed from OpenAI's systems within 30 days for ChatGPT Enterprise, Edu, and Healthcare workspaces, "unless we are legally required to retain them" (Source: openai.com), and API inputs and outputs are similarly retained for up to 30 days unless the customer has separately requested zero data retention for eligible endpoints (Source: openai.com).

Beyond configuration, HIPAA Journal emphasizes that "no technology is HIPAA compliant by itself" and that compliance "depends on how technology is deployed, configured, and used," including a requirement to train workforce members on the compliant use of the product (Source: www.hipaajournal.com). This training now needs to cover a genuinely new distinction that did not exist a few years ago: telling the difference "between consumer AI tools, AI-assisted services, and HIPAA-eligible AI environments," since a clinician who has used ChatGPT personally for years may not

intuitively know that their hospital's ChatGPT for Healthcare workspace operates under an entirely different data-handling regime than the app on their personal phone (Source: www.hipaajournal.com). Not every feature within a HIPAA-eligible workspace is automatically covered, either: workplace administrators can selectively enable "additional ChatGPT functionality that is not covered under BAA," but such features "are disabled by default" and are intended only for uses that do not involve PHI (Source: help.openai.com). Practices implementing any HIPAA-eligible OpenAI product should therefore build a compliance checklist that includes:

- **Confirming BAA scope** directly with OpenAI's sales or compliance team rather than assuming coverage extends to every ChatGPT feature or every workspace an employee might use.
- **Disabling non-BAA-covered features** by default and enabling them selectively through role-based access control groups only for use cases confirmed not to involve PHI.
- **Executing a business associate agreement before go-live**, since legal guidance for healthcare organizations advises that, "prior to implementing any AI technology that processes, secures or accesses PHI, covered entities should enter into a business associate agreement with the vendor" of that technology (Source: www.sheppard.com).
- **Auditing third-party GPTs and connected apps**, since OpenAI's own guidance warns not to share PHI with them "unless you have the appropriate permissions and agreements in place" (Source: help.openai.com).
- **Delivering role-specific HIPAA training** that explicitly distinguishes personal ChatGPT accounts from institutionally licensed, BAA-covered workspaces, and that reminds staff clinicians "should avoid entering any protected health information into a chatbot, but that can be harder than it sounds" once transcripts include casual identifying chit-chat (Source: schaeffer.usc.edu).
- **Tracking state-specific consent and disclosure rules**, since some states now require patient consent before PHI is disclosed to generative AI tools or human verification of AI-generated outputs (Source: www.hipaajournal.com).

HIPAA-Compliant Alternatives and the Broader Healthcare AI Landscape

ChatGPT is not the only path to generative AI in a HIPAA-governed setting, and for many organizations it is not even the most common one, since most large language model access in healthcare flows through cloud infrastructure providers or purpose-built clinical vendors rather than through the ChatGPT interface itself. All three major hyperscale clouds now offer HIPAA-eligible generative AI services under a BAA. Microsoft states that it "will enter into BAAs with its covered entity and business associate customers" and that the Microsoft HIPAA BAA is incorporated automatically into the Microsoft Product Terms for licensing customers, without a separate contract to sign, having "enabled the physical, technical, and administrative safeguards required by HIPAA and the HITECH Act inside the in-scope Azure services" (Source: learn.microsoft.com) (Source: learn.microsoft.com). Microsoft is careful to note the limits of that assurance: "there is currently no certification program approved by the US Department of Health and Human Services (HHS)" through which any cloud provider can formally certify HIPAA compliance, and having a Microsoft BAA "doesn't automatically impart compliance onto your cloud solutions" (Source: learn.microsoft.com) (Source: learn.microsoft.com). Microsoft also clarifies that customers cannot simply reuse an existing agreement, since its hyperscale, multi-tenant cloud services must be operated in a standardized manner across all customers (Source: learn.microsoft.com).

Google Cloud takes a similar contractual approach, stating that it "will enter into Business Associate Agreements with customers as necessary under HIPAA" and that its BAA "covers Google Cloud's entire infrastructure (all regions, all zones, all network paths, all points of presence)" rather than a carved-out subset (Source: cloud.google.com) (Source: cloud.google.com). Google echoes the same caution as Microsoft: "there is no certification recognized by the US HHS for HIPAA compliance and that complying with HIPAA is a shared responsibility between the customer and Google" (Source: cloud.google.com), while also noting a genuine pricing differentiator: "other public clouds charge more money for their HIPAA cloud, we do not" (Source: cloud.google.com). AWS goes further still on breadth of eligible services, offering "over 166 HIPAA eligible services" through a standardized Business Associate Addendum, while cautioning that "customers do not automatically inherit HIPAA compliance by using HIPAA eligible services" and must separately configure controls within them (Source: aws.amazon.com) (Source: aws.amazon.com). AWS also flags a model-specific wrinkle unique to generative AI: legal teams must separately verify that a given foundation model's acceptable use policy permits high-risk healthcare use, since "there may be product restrictions or requirements (such as requiring human-in-the-loop) for certain HCLS use cases" (Source: aws.amazon.com), and stresses that key management discipline matters regardless of vendor, warning that "loose key management practices represent a significant risk to your data and business" (Source: aws.amazon.com).

Alongside the hyperscalers, a distinct category of purpose-built clinical AI vendors has emerged, most visibly the ambient AI documentation ("AI scribe") tools that record and transcribe clinical encounters. *Table 2* below compares the major categories of HIPAA-eligible alternatives to consumer ChatGPT.

CATEGORY	REPRESENTATIVE VENDORS	BAA / HIPAA MODEL	PRIMARY USE CASE	NOTABLE DETAIL
Cloud hyperscaler LLM platforms	Microsoft Azure OpenAI Service, Google Cloud Vertex AI, AWS Bedrock	BAA incorporated into master licensing or cloud services agreement	Custom application development by health systems and vendors	AWS Bedrock restricts use to models whose acceptable use policies permit high-risk healthcare use cases, reviewed on a per-model basis
Ambient AI clinical documentation ("AI scribes")	Nabla, Abridge, Nuance DAX Copilot, Suki, Ambience Healthcare	Vendor signs BAA directly with covered entity; underlying model access separately BAA-covered upstream	Real-time transcription and note generation during patient visits	Nabla's ambient assistant at University of Iowa Health Care "does not retain audio recordings (HIPAA-compliant)" (Source: www.nabla.com)
OpenAI direct products	ChatGPT for Healthcare, ChatGPT for Clinicians, API with Modified Retention	Direct BAA with OpenAI	General-purpose clinical reasoning, documentation, research, and custom application backends	Powers Abridge, Ambience, and EliseAI as an upstream model provider in addition to its own front-end products

For organizations without in-house engineering resources to configure a hyperscaler's HIPAA-eligible tier directly, a specialized vendor with a pre-negotiated BAA and an EHR-integrated workflow is often the faster and less error-prone route than building directly on the OpenAI API. Many large healthcare purchasers now also expect independent assurance beyond a bare BAA: "many large healthcare customers and payers require SOC 2 or HITRUST certification from their technology vendors as preconditions for doing business," a bar that applies equally to OpenAI, the hyperscalers, and specialty clinical AI vendors (Source: intuitionlabs.ai). That tradeoff, speed and configuration simplicity versus the flexibility of building custom workflows on raw model access, is precisely where organizations most often need outside implementation guidance: distinguishing between "the vendor has signed a BAA" and "our specific configuration and workforce practices meet HIPAA's Security Rule" is a governance exercise, not a shopping decision, and it is where consultancies advising on regulated health data infrastructure add the most value regardless of which underlying model or cloud a health system ultimately selects.

Data Analysis and Evidence

The quantitative backdrop to this compliance question shows two trends moving in tension with each other: physician adoption of AI is accelerating rapidly, while regulators, state legislatures, and health systems are simultaneously tightening the rules governing how that AI can touch patient data. The American Medical Association's 2026 Physician Survey on Augmented Intelligence, based on responses from "nearly 1,700 doctors from across a diverse range of physician specialties, practice settings and career stages," found that the share of physicians reporting AI use has reached 81%, "more than double what it was when the AMA first polled doctors on health AI in 2023" (Source: www.ama-assn.org) (Source: www.ama-assn.org). More than three-quarters of physicians now believe AI improves their ability to care for patients, up from 65% in 2023, and "the greatest expected advantages are in diagnostic accuracy and work efficiency," with "seven in 10 physicians" viewing AI as a tool to automate tasks that contribute to burnout and 76% saying it can help with patient care (Source: www.ama-assn.org) (Source: www.ama-assn.org) (Source: www.ama-assn.org). Yet the same survey shows privacy anxiety has not receded alongside adoption: about 40% of physicians said they are "equally excited and concerned about AI, citing patient privacy and the integrity of the patient-physician relationship as top concerns," 86% rated data privacy as critically important to broader AI adoption, 85% said they want to be "consulted or directly involved in decisions about AI adoption" within their own organizations, and "clear liability frameworks rank highest among regulatory actions essential to build physician trust and increase adoption of AI tools" (Source: www.ama-assn.org) (Source: www.ama-assn.org) (Source: www.ama-assn.org).

On model quality specifically for clinical use, OpenAI's HealthBench evaluation, built "in partnership with 262 physicians who have practiced in 60 countries" and containing "5,000 realistic health conversations" graded against "48,562 unique rubric criteria," provides a rare independently-structured benchmark (Source: openai.com) (Source: openai.com). OpenAI reports that its frontier models "improved by 28% on HealthBench" in the months following the benchmark's introduction, though it is worth noting this is a vendor-run and vendor-designed evaluation rather than an independently administered one (Source: openai.com). Independent, peer-reviewed literature offers a more measured counterweight: a 2024 analysis published in Frontiers in Radiology notes that "large language models (LLMs) like ChatGPT offer impressive capabilities" but "come with significant

privacy implications that need to be carefully addressed," cautioning that "patient data may be compromised, leading to the unauthorized access, harvesting and sharing of sensitive personal information" and that "established AGI models may inadvertently perpetuate existing biases present in the data used for training," concluding that "proper oversight and regulation are necessary to ensure the integration of AGI technologies into the healthcare system is both safe and effective" (Source: pmc.ncbi.nlm.nih.gov) (Source: pmc.ncbi.nlm.nih.gov) (Source: pmc.ncbi.nlm.nih.gov). A separate healthcare compliance analysis adds that "autonomous AI working without oversight presents real risks in the healthcare space," and that "sensitive health data, even when de-identified, can be re-identified or leaked through model vulnerabilities like inference or inversion attacks" (Source: www.paubox.com) (Source: www.paubox.com).

The regulatory downside risk, meanwhile, remains substantial and is quantified in specific dollar terms. Under the current civil penalty structure, HIPAA violations are penalized across four culpability tiers, with fines ranging from a Tier 1 minimum of $145 per violation up to a Tier 4 maximum of $2,190,294 per violation category per year (Source: www.hipaajournal.com). *Table 3* below sets out the full inflation-adjusted penalty structure as of the 2025 inflation multiplier applied January 28, 2026.

PENALTY TIER	CULPABILITY	MINIMUM PENALTY PER VIOLATION	MAXIMUM PENALTY PER VIOLATION	MAXIMUM ANNUAL CAP
Tier 1	Lack of Knowledge	$145	$73,011	$2,190,294
Tier 2	Reasonable Cause	$1,461	$73,011	$2,190,294
Tier 3	Willful Neglect (corrected)	$14,602	$73,011	$2,190,294
Tier 4	Willful Neglect (uncorrected after 30 days)	$73,011	$2,190,294	$2,190,294

(Figures as published by HIPAA Journal, reflecting the 2025 inflation multiplier applied by HHS on January 28, 2026 (Source: www.hipaajournal.com).)

These are not abstract figures. In one frequently cited enforcement precedent, the remote cardiac monitoring service CardioNet "was fined $2.5 million for failing to fully understand the HIPAA requirements and subsequently failing to conduct a complete risk assessment," after PHI belonging to 1,391 individuals was exposed when an employee's laptop was stolen (Source: www.hipaajournal.com). More broadly, OCR "closed" 22 enforcement actions in 2024 alone through settlements or civil monetary penalties (Source: www.hipaajournal.com), and separate industry analysis puts the scale of the underlying breach problem in context: 734 large healthcare data breaches were reported to OCR in 2024, exposing over 276.8 million records, more than 80% of the U.S. population, and the average healthcare data breach cost $7.42 million in 2025, the costliest of any industry for the 14th consecutive year (Source: intuitionlabs.ai) (Source: intuitionlabs.ai). The compliance stakes are not merely regulatory: one industry analysis warns that "60% of small businesses" fail within six months of a major data breach, a reminder that the operational risk of a careless AI deployment can compound the strictly legal risk (Source: intuitionlabs.ai). Against that backdrop, the risk of a single clinician pasting an unredacted clinical note into a non-BAA-covered chatbot is not a hypothetical compliance nicety; it is one additional, largely invisible vector into an already strained breach environment.

Case Studies and Real-World Examples

Boston Children's Hospital and the OpenAI for Healthcare Rollout

Boston Children's Hospital is among the earliest named adopters of ChatGPT for Healthcare, and its account of the rollout illustrates why enterprise deployment, rather than individual clinician workarounds, is increasingly the preferred path for large systems. John Brownstein, the hospital's Senior Vice President and Chief Innovation Officer, described the transition in OpenAI's own launch materials: "Our early work with a custom OpenAI-powered solution allowed us to move quickly, prove value in a secure environment, and establish strong governance foundations. ChatGPT for Healthcare offers a path toward operational scale, providing an enterprise-grade platform that can support broad, responsible adoption across clinical, research, and administrative teams" (Source: openai.com). Boston Children's Hospital is joined by AdventHealth, Baylor Scott & White Health, Cedars-Sinai Medical Center, HCA Healthcare, Memorial Sloan Kettering Cancer Center, Stanford Medicine Children's Health, and UCSF as early institutional adopters, and HIPAA Journal independently confirms the timing, describing ChatGPT for Healthcare as "an enterprise-grade AI product

designed specifically for hospitals, clinicians, and regulated healthcare environments" that "was launched in January 2026" (Source: www.hipaajournal.com). The pattern across these early adopters is consistent: each moved from an ad hoc, single-team pilot toward a centrally governed, BAA-covered workspace rather than allowing individual clinicians to rely on personal ChatGPT accounts.

Penda Health's AI Consult in Nairobi, Kenya

The most rigorously measured evidence of an OpenAI-powered clinical copilot's real-world impact comes not from the United States but from Penda Health, a primary care provider operating 16 clinics in Nairobi, Kenya since 2012. Penda's clinician copilot, AI Consult, runs on OpenAI's API and was studied across "39,849 patient visits across 15 clinics," comparing outcomes for clinicians with and without access to the tool (Source: openai.com). Clinicians using AI Consult had a "16% relative reduction in diagnostic errors and a 13% reduction in treatment errors compared to those without" the tool, with even larger effects in visits where the system would have issued a safety-critical "red alert" (Source: openai.com). Notably for a privacy-focused analysis, Penda's implementation was built with data minimization in mind from the outset: "documentation without patient identifiers is sent to the OpenAI API at key points," and reviewers found that "in no case did AI Consult recommendations lead to harm" across the safety reports raised during the study period (Source: openai.com) (Source: openai.com). This case illustrates a design principle relevant well beyond Kenya: an OpenAI-powered clinical tool can be engineered to strip identifiers before data ever reaches the model, reducing both privacy exposure and the compliance burden that comes with transmitting full PHI.

University of Iowa Health Care's Enterprise Deployment of Nabla

University of Iowa Health Care (UIHC), Iowa's only comprehensive academic health system with more than 1,200 physicians and dentists, illustrates the alternative path of adopting a specialized, HIPAA-covered ambient AI vendor rather than building on ChatGPT directly. UIHC's decision was driven by measurable documentation burden: a 2023 KLAS Arch Collaborative survey found that "70% of physicians cited documentation as a barrier to personal efficiency and patient care" (Source: www.nabla.com). After a five-week pilot of Nabla's ambient AI assistant, which "does not retain audio recordings (HIPAA-compliant)," UIHC expanded to enterprise-wide deployment, ultimately reaching "2,200 clinicians" generating "45,000 daily encounters" (Source: www.nabla.com) (Source: www.nabla.com). The measured impact included a "26% sustained reduction in burnout," sustained over a two-year follow-up period (Source: www.nabla.com). UIHC's path underscores that not every health system needs to build directly on the OpenAI API or negotiate a bespoke ChatGPT for Healthcare BAA; a pre-integrated, EHR-native vendor with its own BAA can achieve similar documentation-burden relief with a shorter implementation timeline.

A Turkish Hospital and the Limits of Consumer ChatGPT (Illustrative International Example)

Not every relevant case study involves a U.S. HIPAA-covered entity, but international examples illustrate the accuracy and reliability risks that make unsupervised, non-BAA consumer chatbot use dangerous even where HIPAA itself does not apply. A Reuters investigation published February 9, 2026, profiled Dr. Cem Aksoy, a medical resident at a hospital in Ankara, Turkey, whose 18-year-old patient turned to ChatGPT after a cancer diagnosis; "the bot said he might survive only five years," a prognosis that proved wrong after a plastic surgeon successfully removed the tumor (Source: www.reuters.com). Weeks later, the same patient, coughing after having recently started smoking, asked ChatGPT whether it might indicate cancer metastasis to his lungs, and became convinced he needed to write a will before his physician could reassure him his lungs were fine (Source: www.reuters.com). Dr. Aksoy summarized the underlying hazard: "When someone is distressed and unguided, \[an AI chatbot\] just drags them into this forest of knowledge without coherent context" (Source: www.reuters.com). An OpenAI spokesperson responded that its "newest models have significantly improved how they handle health questions" and reiterated that ChatGPT is not intended as a substitute for a medical professional's guidance (Source: www.reuters.com). This example is not a HIPAA violation in the technical legal sense, since it involves a patient's own use of a consumer product outside U.S. jurisdiction, but it directly illustrates the practice-level risk this report addresses: even where a BAA and PHI-handling rules are not in play, unsupervised consumer chatbot use in high-stakes medical contexts carries real accuracy risk that a properly governed, clinician-supervised deployment is specifically designed to mitigate. USC's Schaeffer Center makes a related observation about detection: "patients usually only hear about it through newspapers getting wind of it through the grapevine and reporting on an incident," meaning the gap between an unauthorized disclosure occurring and anyone outside the clinic learning about it can be long and largely invisible to the patient involved (Source: schaeffer.usc.edu).

Implications and Future Directions

Three forces are converging to reshape how practices should think about ChatGPT and HIPAA compliance over the next 12 to 24 months. First, OpenAI's own regulatory posture is in motion: the mismatch between its help center's older description of sales-managed ChatGPT Enterprise BAA eligibility and HIPAA Journal's 2026 assertion that only ChatGPT for Healthcare, not general Enterprise, now qualifies suggests OpenAI is actively consolidating its healthcare compliance offering around a single, more tightly governed product line. Practices currently relying on an older Enterprise BAA arrangement should proactively confirm its continued validity rather than assume grandfathering.

Second, state legislatures have filled the vacuum left by the absence of federal AI-specific healthcare legislation, and the pace is accelerating rather than slowing. As one legal analysis puts it, "the U.S. Congress has yet to pass significant legislation regulating the use of AI in healthcare," while states have "moved aggressively to enact laws governing the use of AI in healthcare" in 2024, 2025, and 2026 (Source: www.hklaw.com). Texas's Responsible Artificial Intelligence Governance Act requires that "healthcare providers must give patients or their personal representatives conspicuous written disclosure whenever an AI system is used in diagnosis or treatment," while Illinois separately requires notification and, in some cases, explicit consent, and bars AI from making independent therapeutic decisions (Source: www.beckershospitalreview.com). California has layered on its own AI-identity rules, with AB 489 prohibiting AI systems from using "terms, letters, phrases or design elements that imply the AI possesses a healthcare license," while SB 243 "imposes specific safety protocols on AI companion bots, requiring them to prevent harmful conversations, detect mental health crises and suicidal ideation, and establish guardrails for users younger than 18" (Source: www.beckershospitalreview.com). Maine's HB 2082 restricts licensed mental health professionals to using AI "for administrative functions and limited supplementary purposes," expressly barring therapeutic communications or independent patient interaction by AI (Source: www.beckershospitalreview.com). Idaho and Nebraska took a near-identical approach, each enacting a "Conversational AI Safety Act" that requires public-facing chatbots to disclose their non-human nature and follow crisis-response protocols (Source: www.hklaw.com). Utah, meanwhile, is testing a lighter-touch model: its insurer-facing SB 319 requires disclosure of AI use in prior authorization review while still requiring "healthcare professionals to render adverse determinations based on their independent medical judgment," and the state's separate AI Policy Act sandbox is piloting a program that would allow AI systems to "autonomously renew certain routine prescriptions for patients with chronic conditions" under state supervision (Source: www.hklaw.com). For a multi-state health system, this means a single national ChatGPT for Healthcare BAA is necessary but not sufficient; state-by-state disclosure, consent, and human-oversight requirements now sit on top of the federal HIPAA baseline and must be tracked separately.

Third, HHS itself is moving to tighten the Security Rule's baseline expectations regardless of which AI vendor a practice selects, and international comparisons show U.S. rules are, in some respects, already the looser regime: unlike HIPAA's 60-day breach notification window and implied consent for treatment uses, the European Union's GDPR "requires explicit opt-in consent for each use of health data and mandates data breach notification within 72 hours," a stricter standard that any multinational health system or research collaboration must separately architect for (Source: intuitionlabs.ai). Domestically, a January 2025 Notice of Proposed Rulemaking would eliminate the "addressable versus required" distinction for security safeguards, effectively making previously optional protections mandatory for all covered entities and business associates, with the rule expected to move toward finalization through 2026 (Source: intuitionlabs.ai). For practices already navigating ChatGPT for Healthcare or a competing HIPAA-eligible AI product, this means the security bar underneath any BAA is set to rise regardless of vendor choice, reinforcing that vendor selection and internal governance are two separate, equally important workstreams rather than one substituting for the other.

Frequently Asked Questions (FAQs)

Is ChatGPT HIPAA compliant in 2026? Standard ChatGPT (Free, Plus, Team) and self-serve ChatGPT Business are not HIPAA compliant, because OpenAI will not sign a BAA for them. Specific enterprise products, ChatGPT for Healthcare, ChatGPT for Clinicians, ChatGPT for Enterprise with Regulated Workspace, and the API with Modified Retention, can support HIPAA-compliant use once a BAA is executed and the workspace is properly configured.

Can doctors use ChatGPT with patient data? Only within a HIPAA-eligible, BAA-covered workspace, and even then only after configuring administrative controls and workforce training. Using a personal or free ChatGPT account with identifiable patient information is, according to legal experts interviewed by USC's Schaeffer Center, a violation regardless of whether the clinician later opts the account out of model training, since "you've just violated HIPAA because the data has left the health system" the moment it was entered (Source: schaeffer.usc.edu).

What is a ChatGPT business associate agreement, and who can get one? It is the contract required under HIPAA before PHI can be disclosed to OpenAI as a vendor. For ChatGPT products, only sales-managed Enterprise or Edu accounts, ChatGPT for Healthcare organizations, and individually verified clinicians using ChatGPT for Clinicians can currently obtain one; ChatGPT Business is explicitly excluded from any BAA arrangement (Source: www.hipaajournal.com).

Does OpenAI offer a BAA for healthcare use of its API? Yes. Organizations email baa@openai.com, and once approved and provisioned with Modified Retention, a defined set of API endpoints becomes HIPAA-eligible, without requiring a full enterprise agreement.

Is ChatGPT Enterprise HIPAA compliant? Only conditionally, and sources describe this differently. OpenAI's own API BAA support page still lists sales-managed ChatGPT Enterprise and Edu accounts as BAA-eligible, while HIPAA Journal's more recent 2026 analysis states OpenAI "will not enter into a Business Associate Agreement for the Free, Plus, Team, or Enterprise versions of its product," suggesting ChatGPT for Healthcare has become the primary path for regulated deployments. Organizations should confirm current eligibility directly with OpenAI sales rather than relying on either description alone.

What are the risks of using ChatGPT in medical practice? Peer-reviewed analysis identifies risks including inaccurate or outdated medical advice, patient privacy violations, fabricated content, and the perpetuation of training-data bias (Source: pmc.ncbi.nlm.nih.gov). Additional documented risks include automation bias, where users "over trust AI-generated medical responses and view them to be as valid as doctors" even when accuracy is low, and the underlying "black box" opacity of a system "whose internal workings are a mystery to its users," which makes it difficult for clinicians to explain or challenge an AI-generated recommendation (Source: www.paubox.com) (Source: www.paubox.com).

What are HIPAA-compliant alternatives to ChatGPT? Options include Microsoft Azure OpenAI Service, Google Cloud Vertex AI, and AWS Bedrock, each offering a BAA covering their respective generative AI services, alongside specialized clinical AI vendors such as Nabla, Abridge, Nuance DAX Copilot, Suki, and Ambience Healthcare, which build their own BAA-covered products on top of underlying model access.

What happens if a practice violates HIPAA by using ChatGPT improperly? Penalties scale by culpability tier, from $145 per violation at the least culpable Tier 1 up to $2,190,294 per violation category annually at Tier 4, and covered entities can face additional liability if they knew or should have known about a business associate's noncompliance.

Does OpenAI train its models on data entered into ChatGPT for Healthcare? No. Content shared in ChatGPT for Healthcare is not used to train models, and by default this same non-training commitment extends to ChatGPT Business, Enterprise, Edu, and the API Platform after March 1, 2023, unless a customer explicitly opts in, since OpenAI states that "data from ChatGPT Business, ChatGPT Enterprise, ChatGPT for Healthcare, ChatGPT Edu, ChatGPT for Teachers, and the API Platform (after March 1, 2023) isn't used for training our models" (Source: openai.com).

Conclusion

The honest, current answer to "is ChatGPT HIPAA compliant" is that it depends entirely on which product, which contract, and which configuration a practice is actually using, and that answer has become more favorable, not less, over the past year without becoming any simpler. Free, Plus, Team, and Business ChatGPT remain categorically off-limits for Protected Health Information because OpenAI will not sign a Business Associate Agreement for them. ChatGPT for Healthcare, launched in January 2026 and already deployed at major health systems, along with ChatGPT for Clinicians and the API with Modified Retention, can support compliant use, but only after a BAA is executed, the workspace is deliberately configured with role-based access controls, audit logging, and data retention limits, and the workforce is trained to distinguish a personal ChatGPT account from an institutionally licensed, BAA-covered one. Neither the BAA nor the SOC 2 attestation that several ChatGPT products carry substitutes for the underlying HIPAA Security Rule obligation to conduct a risk analysis and implement administrative, physical, and technical safeguards; that responsibility sits with the covered entity regardless of which vendor it selects.

For practices weighing ChatGPT against Microsoft Azure OpenAI Service, Google Cloud Vertex AI, AWS Bedrock, or specialized clinical vendors such as Nabla and Abridge, the comparison should center less on whether a BAA is theoretically available, since all of these paths now offer one, and more on which combination of deployment speed, EHR integration, data minimization design, and governance support best fits the organization's existing compliance infrastructure. The regulatory environment is also not static: state laws in Texas, Illinois, California, Maine, Utah, Idaho, and Nebraska are layering new disclosure, consent, and human-oversight requirements on top of the federal HIPAA baseline, and a pending HHS Notice of Proposed Rulemaking would tighten Security Rule expectations across the board. Practices that treat a signed BAA as the finish line rather than the starting point of compliance, and that build workforce training and governance processes to match, will be far better positioned than those that assume a vendor's contractual assurances alone resolve the question.

Tags: chatgpt hipaa compliant, hipaa compliance, chatgpt for healthcare, openai baa, business associate agreement, healthcare ai, hipaa compliant alternatives, chatgpt enterprise

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.