What Counts as PHI When Using AI? A De-Identification Guide

Published July 19, 2026 45 min read

Executive Summary

Under the Health Insurance Portability and Accountability Act of 1996 (HIPAA), protected health information (PHI) is individually identifiable health information that is created, received, maintained, or transmitted by a covered entity or its business associate, and that relates to an individual's past, present, or future physical or mental health, the provision of care, or payment for care (Source: www.hhs.gov) (Source: www.ecfr.gov). The Safe Harbor method of the HIPAA Privacy Rule lists 18 categories of identifiers, from names and Social Security numbers to IP addresses and full-face photographs, that must be stripped from a data set before it can be treated as de-identified and used freely, including in artificial intelligence (AI) systems (Source: www.hhs.gov). Whether an AI tool is handling PHI turns not on the technology itself but on who is using it, what data goes in, and whether a signed Business Associate Agreement (BAA) is in place: HIPAA "does not create exceptions" for AI, and the same privacy and security standards that apply to human staff apply to algorithms (Source: www.paubox.com).

As of July 2026, major AI vendors offer HIPAA-eligible tiers gated behind enterprise contracts and executed BAAs. OpenAI covers **ChatGPT for Healthcare**, ChatGPT Enterprise with a Regulated Workspace, and specific API endpoints under a Business Associate and Healthcare Addendum (Source: help.openai.com). Anthropic covers Claude Enterprise and the first-party API only after an organization's Primary Owner activates HIPAA compliance and signs the BAA, and several beta features (Cowork, Claude in Chrome, MCP connectors) remain explicitly excluded (Source: support.claude.com). By contrast, consumer-tier chatbots accessed through a browser with no BAA in force are not HIPAA-eligible surfaces, a distinction that institutions like Mass General Brigham enforce operationally by prohibiting the entry of PHI, even de-identified data, into any large language model outside an approved internal platform (Source: libguides.massgeneral.org).

Enforcement activity underscores the stakes of getting this wrong. The Federal Trade Commission fined **GoodRx** $1.5 million in 2023 for disclosing users' health information to advertisers despite marketing itself as "HIPAA Secure" while not actually being a covered entity (Source: www.justice.gov), and telehealth provider **Cerebral** agreed to pay more than $7 million and accept a first-of-its-kind ban on using health data for advertising after leaking mental health details through tracking pixels (Source: themarkup.org). Separately, HHS's Office for Civil Rights (OCR) reported that healthcare data breaches affecting 500 or more individuals have exposed the records of more than one billion Americans since 2009, with the single largest incident, the 2024 Change Healthcare ransomware attack, compromising an estimated 192.7 million individuals (Source: www.hipaajournal.com).

This report defines what counts as PHI when AI enters the picture, walks through all 18 Safe Harbor identifiers, explains the two HIPAA de-identification methods (Safe Harbor and Expert Determination), clarifies when AI output itself becomes PHI, distinguishes PHI from personally identifiable information (PII), and lays out a practical implementation framework for de-identifying data before it reaches a model. It draws on federal regulatory text, HHS guidance, peer-reviewed legal analysis, vendor BAA documentation, enforcement records, and named case studies including Mayo Clinic's enterprise deployment of the Abridge ambient-AI scribe under a signed BAA (Source: www.healthcaredive.com) and the American Medical Association's finding that 81% of physicians now use AI tools professionally, more than double the rate reported in 2023 (Source: www.ama-assn.org). The throughline for any organization evaluating AI tools against PHI risk is straightforward: identify whether a covered entity or business associate relationship exists, confirm a BAA covers the specific product surface being used, and either de-identify data under a defensible method or restrict PHI to contractually covered environments.

Introduction and Background

Generative AI tools have moved from novelty to daily workflow inside healthcare and life sciences organizations faster than most privacy frameworks have adapted. The AMA's 2026 Physician Survey on Augmented Intelligence, fielded from January 15 to February 2, 2026, among nearly 1,700 physicians, found that health AI is now used to summarize medical research, draft discharge instructions, document visit notes, and generate responses to patient messages, with reported usage climbing from 39% to well beyond that for research summarization alone since 2023 (Source: www.ama-assn.org). That same survey found 86% of physicians rate data privacy as important to broader AI adoption, and 88% want robust safety and efficacy validation before trusting these tools further (Source: www.ama-assn.org). About 40% of physicians describe themselves as equally excited and concerned about AI, citing patient privacy and the integrity of the patient-physician relationship as their top concerns (Source: www.ama-assn.org). The tension is obvious: the same conversational interface that makes a large language model useful for drafting a discharge summary is also the interface through which a clinician, researcher, or commercial-operations analyst can inadvertently transmit protected health information to a vendor with no contractual obligation to protect it.

Protected health information (PHI) is a term of art defined under the HIPAA Privacy Rule, codified at 45 CFR Part 160 and Part 164. It is not simply "health data" in a colloquial sense; it is individually identifiable health information held by a **covered entity** (a health plan, healthcare clearinghouse, or healthcare provider that transmits health information electronically in connection with a covered transaction) or a **business associate** of a covered entity (Source: www.ecfr.gov).This distinction matters enormously for AI governance: the same patient record entered into the same chatbot can be PHI in one organizational context and simply "sensitive personal data" outside HIPAA's jurisdiction in another, such as when a direct-to-consumer wellness app that is not a covered entity collects health information directly from users (Source: www.foreseemed.com). That was precisely the fact pattern the FTC used against GoodRx, which marketed itself with a "HIPAA Secure" seal while operating outside HIPAA's scope entirely, and was instead pursued under the FTC's Health Breach Notification Rule (Source: www.justice.gov).

For life sciences and pharmaceutical organizations, the stakes extend beyond clinical documentation. Commercial teams working in customer relationship management platforms, clinical operations teams processing trial data, and medical affairs teams summarizing adverse event reports all routinely touch data that can qualify as PHI once it is linked to an identifiable patient, and increasingly these teams are experimenting with AI copilots layered on top of existing systems. IntuitionLabs, a life sciences and AI consultancy and Veeva Vault CRM X-Pages Partner, lists HIPAA, GDPR, FDA 21 CFR Part 11, and EU Annex 11 among the compliance frameworks its implementations are built to satisfy, reflecting how tightly AI enablement and regulatory data handling are now coupled in pharma commercial and clinical operations (Source: intuitionlabs.ai). This report unpacks what counts as PHI when AI is involved, how the two HIPAA de-identification methods work, when AI-generated output itself becomes PHI, how PHI differs from the broader category of personally identifiable information (PII), and what a defensible AI intake process looks like in practice.

Defining PHI: The Legal Taxonomy Under HIPAA

The starting point for any AI governance conversation is the statutory definition. Under 45 CFR 160.103, **health information** means information, whether oral or recorded in any form or medium, that is created or received by a healthcare provider, health plan, public health authority, employer, life insurer, school, or healthcare clearinghouse, and that relates to an individual's past, present, or future physical or mental health or condition, the provision of care, or payment for care (Source: www.ecfr.gov). **Individually identifiable health information** narrows that further to health information that either identifies the individual or offers a reasonable basis to believe it could be used to identify the individual (Source: www.ecfr.gov). **PHI** is that individually identifiable health information when it is transmitted or maintained by electronic media, or transmitted or maintained in any other form or medium by a covered entity or business associate (Source: www.ecfr.gov).

Three elements must all be present for information to qualify as PHI, and an AI governance program should test for each explicitly:

- **Health relationship:** The information relates to physical or mental health, the provision of care, or payment for care. A patient's zip code alone, listed in a phone directory, is not PHI; the same zip code attached to a note that the patient was treated at a specific clinic is (Source:

www.hhs.gov).

- **Identifiability:** The data either directly identifies the individual or provides a reasonable basis to do so, whether alone or combined with other reasonably available information (Source: www.hhs.gov).
- **Covered custody:** The information is held or transmitted by a HIPAA covered entity or its business associate. Information collected directly by a consumer app, employer, or life insurer outside a covered relationship generally falls outside HIPAA, even though it may be regulated by other laws such as the FTC's Health Breach Notification Rule or state consumer health data statutes (Source: www.justice.gov).

A **business associate** is any person or entity that, on behalf of a covered entity, creates, receives, maintains, or transmits PHI for a regulated function, including claims processing, data analysis, utilization review, billing, or management services involving PHI disclosure (Source: www.ecfr.gov). Peer-reviewed legal analysis of AI compliance frames this the same way, describing a business associate as a person or organization that conducts certain activities on PHI on behalf of, or provides services such as financial, administrative, management, legal, or data aggregation for, a covered entity (Source: pmc.ncbi.nlm.nih.gov). This definition is the hinge on which AI vendor status turns. When a hospital or physician inputs patient data into an AI chat tool to generate discharge notes, draft patient letters, or summarize a chart, the AI vendor processing that data becomes a business associate under HIPAA and must execute a BAA (Source: www.paubox.com). Critically, automated processing counts: an AI system that never has a human review PHI is still considered to be accessing that data under HIPAA, so the absence of human eyes on the record does not by itself exempt a vendor from business associate status. The "conduit exception," which excuses entities that merely transmit data without accessing its content, generally does not apply to AI systems that must process and "see" data to function, such as an AI model analyzing a medical image (Source: www.paubox.com).

The 18 HIPAA Safe Harbor Identifiers and Where AI Tools Encounter Them

The most concrete, checklist-style answer to "what counts as PHI" comes from the Safe Harbor method's enumerated list of 18 identifier categories. Under 45 CFR 164.514(b)(2), a covered entity may treat a data set as de-identified if all 18 categories of identifiers relating to the individual, and to their relatives, employers, or household members, are removed, and the covered entity has no actual knowledge that the remaining information could be used to re-identify the individual (Source: www.hhs.gov).

Table 1 below lists all 18 categories exactly as codified, alongside plain-language notes on how each surfaces in everyday AI workflows such as chatbot prompts, ambient scribes, and document-summarization tools.

#	HIPAA SAFE HARBOR IDENTIFIER	HOW IT COMMONLY ENTERS AN AI PROMPT OR UPLOAD
1	Names	Patient names typed into a chatbot prompt or contained in an uploaded chart, note, or transcript.
2	Geographic subdivisions smaller than a state (street address, city, county, precinct, ZIP code)	Home addresses in intake forms; ZIP codes in population-health prompts (first three digits may remain if the area has more than 20,000 residents) (Source: www.hhs.gov).
3	All elements of dates (except year) directly related to an individual, including birth, admission, discharge, and death dates, and ages over 89	Visit dates or lab-result timestamps pasted into a summarization prompt; exact ages of elderly patients.
4	Telephone numbers	Contact fields copied from an EHR into a drafting tool.
5	Fax numbers	Referral cover sheets uploaded for AI-assisted routing.
6	Email addresses	Patient portal message threads fed into a reply-drafting assistant.
7	Social Security numbers	Billing or insurance data pasted for claims-related AI processing.
8	Medical record numbers	Chart identifiers retained in copy-pasted clinical text.
9	Health plan beneficiary numbers	Insurance eligibility text used in prior-authorization automation.
10	Account numbers	Patient billing account numbers in revenue-cycle AI tools.
11	Certificate or license numbers	Professional license numbers appearing in credentialing workflows.
12	Vehicle identifiers and serial numbers, including license plates	Rare, but present in EMS or trauma documentation.
13	Device identifiers and serial numbers	Implant or wearable device serial numbers in device-monitoring AI.
14	Web URLs	Patient-specific portal links embedded in message threads.
15	IP addresses	Metadata captured by telehealth or remote-monitoring AI platforms.
16	Biometric identifiers, including finger and voice prints	Voice recordings ingested by ambient AI scribes before transcription.
17	Full-face photographs and comparable images	Dermatology or wound-care images uploaded to diagnostic AI tools.
18	Any other unique identifying number, characteristic, or code	Clinical trial record numbers, rare-diagnosis descriptions, or unique job titles that could single out an individual (Source: www.hhs.gov).

Table 1 makes clear that PHI risk in AI workflows rarely arrives as a single obvious field like a Social Security number; it more often arrives as a combination of quasi-identifiers, a birth date plus a rare diagnosis plus a small-town ZIP code, that together create a reasonable basis for re-identification even when no single field is directly identifying. HHS guidance is explicit that Safe Harbor does not permit partial identifiers either: a data set containing only patient initials or the last four digits of a Social Security number still fails the standard, because these are treated as derivatives of listed identifiers (Source: www.hhs.gov). HHS guidance also flags "any other unique identifying number, characteristic, or code" as a catch-all: an occupation description like "current President of State University" can single out an individual just as effectively as a name, even after the 17 enumerated categories are removed (Source: www.hhs.gov). Commercial AI vendors building healthcare-specific tooling have engineered products

around exactly this list: Amazon Comprehend Medical's DetectPHI operation maps entity types such as NAME, ADDRESS, ID, and DATE directly onto the 18 Safe Harbor categories so developers can programmatically flag and redact them before data reaches a downstream model (Source: docs.aws.amazon.com). Amazon Comprehend Medical assigns a confidence score to each entity it flags, and AWS explicitly advises developers to identify the right confidence threshold for their use case and filter out entities that fall below it, rather than treating any single detection run as a definitive compliance determination (Source: docs.aws.amazon.com).

De-Identification Methods: Safe Harbor and Expert Determination

HIPAA provides exactly two recognized paths to transform PHI into data that is no longer subject to the Privacy Rule at all. Once information is properly de-identified under either method, HIPAA "does not restrict the use or disclosure" of that information, because it no longer meets the definition of PHI (Source: www.hhs.gov). This is the legal foundation for the common practice of de-identifying clinical or claims data before feeding it into an AI model for research, analytics, or product development, and it is confirmed independently in the peer-reviewed literature: de-identified health information that can no longer be used to identify the data subject falls outside the definition of PHI, and there is no restriction on its use or disclosure under HIPAA (Source: pmc.ncbi.nlm.nih.gov).

The Safe Harbor method requires removing all 18 categories of identifiers listed in Table 1 and confirming the covered entity has no actual knowledge that the remaining information could still identify someone (Source: www.hhs.gov). "Actual knowledge" is a high bar defined by HHS as clear and direct knowledge, not mere awareness of academic re-identification research; a covered entity is not expected to presume that every recipient of de-identified data possesses sophisticated re-identification capabilities (Source: www.hhs.gov). But actual knowledge is triggered in concrete scenarios HHS has spelled out: if a record reveals an occupation as unusual as "former president of State University," if the recipient is known to have a family member represented in the data, or if a rare clinical event (such as an unusually large multiple birth) was widely publicized, Safe Harbor is not satisfied merely by stripping the 18 fields (Source: www.hhs.gov).

The Expert Determination method takes a statistical rather than checklist approach. A person with appropriate knowledge of generally accepted statistical and scientific methods applies those principles to determine that the risk is "very small" that the information could be used, alone or combined with other reasonably available information, to identify an individual, and documents the methods and results that justify the determination (Source: www.hhs.gov). There is no fixed credential that qualifies someone as an "expert"; OCR reviews the person's professional experience and training on a case-by-case basis during enforcement (Source: www.hhs.gov), nor is there a universal numeric risk threshold; the acceptable level of "very small" risk depends on the anticipated recipient and data-sharing context (Source: www.hhs.gov). Statistical disclosure limitation techniques commonly used under this method include **k-anonymization** (generalizing or suppressing quasi-identifiers such as age and ZIP code until every combination of attributes matches at least *k* individuals in the data set) (Source: www.hhs.gov), differential privacy, pseudonymization, and synthetic data generation, all of which are catalogued in NIST Special Publication 800-188, "De-Identifying Government Datasets: Techniques and Governance," published in September 2023 (Source: csrc.nist.gov). Notably, an expert may assign a re-identification code to a de-identified data set, including one derived using a one-way cryptographic hash function, provided the key is never disclosed to recipients, allowing a covered entity to preserve its own ability to re-link records without compromising the de-identified status of the shared copy (Source: www.hhs.gov).

Neither method produces a zero-risk outcome. HHS is explicit that both methods, even when properly applied, yield de-identified data that retains some residual risk of re-identification, small but not zero (Source: www.hhs.gov). This residual risk is the reason compliance commentators flag "data triangulation" as a growing concern: with massive access by dominant technology companies such as Meta, Google, and Microsoft to patients' personal information, there is a significant risk of privacy violation through re-identification of health data sets that were de-identified through the Safe Harbor mechanism (Source: pmc.ncbi.nlm.nih.gov). Organizations that need to de-identify data before it reaches an AI tool typically use automated detection software such as Amazon Comprehend Medical's DetectPHI, Microsoft's open-source Presidio framework for PII/PHI identification and anonymization across text, images, and structured data (Source: microsoft.github.io), or purpose-built medical natural-language-processing vendors, combined with human review before data is released for AI training, analytics, or prompt-based use, since automated tools alone are explicitly not guaranteed to catch every identifier (Source: microsoft.github.io).

Between full de-identification and raw PHI, the Privacy Rule also recognizes an intermediate category: the **limited data set**, from which most direct identifiers are removed but certain fields such as dates and geographic subdivisions larger than street address may be retained if the recipient signs a data use agreement restricting further use and disclosure. HHS guidance notes that a covered entity may require the recipient of de-identified information to enter into a data use agreement to access files with a known disclosure risk, such as is required for release of a limited data set under the Privacy Rule (Source: www.hhs.gov). This distinction is operationally relevant to AI use cases in research cohorts or population-health analytics, where service dates or general geographic detail may be necessary for the analysis to be useful; a limited data set paired with a data use agreement can sometimes thread that needle, though the data remains PHI in the eyes of the Privacy Rule and BAA obligations continue to apply to any AI vendor processing it.

Is AI Output Considered PHI? Business Associates, BAAs, and Consumer Tools

A frequent question is whether AI-generated output, a summary, a draft note, a chatbot's answer, is itself PHI, separate from whatever PHI was in the input. The answer follows directly from the statutory definition: if the output is derived from individually identifiable health information and remains capable of identifying the individual, it is PHI regardless of whether a human or a model produced it. HIPAA's use and disclosure restrictions apply based on whether PHI is created, received, maintained, or transmitted, not on the type of AI technology used to process it. A discharge summary drafted by an ambient AI scribe from a recorded patient encounter is PHI in exactly the same way a discharge summary typed by a resident would be; the authorship mechanism is irrelevant to the legal classification.

What does change with AI is the vendor relationship. Once an AI system is processing PHI on behalf of a covered entity, for functions such as claims processing, data analysis, transcription, diagnostic support, or draft correspondence, its provider becomes a business associate and a BAA is required before any PHI should be transmitted (Source: www.paubox.com). Encrypting PHI in transit or at rest is a valuable security safeguard, but it does not substitute for the underlying contractual requirement to have a BAA in place wherever a vendor handles PHI. Downstream subcontracting also matters: if an AI vendor routes model hosting or inference through another company, that subcontractor becomes a business associate in its own right and must be covered by appropriate agreements as well (Source: www.ecfr.gov).

A separate and often overlooked gap involves PHI that a patient discloses directly to a consumer AI chatbot rather than data a covered entity inputs on the patient's behalf. When a patient shares personal health information with a chatbot seeking medical advice, the AI developer or vendor is neither a covered entity nor a business associate, so HIPAA does not apply to that exchange at all (Source: pmc.ncbi.nlm.nih.gov). Peer-reviewed analysis of this gap notes it matters because a considerable number of AI developers and vendors are technology companies that operate outside the traditional scope of HIPAA's covered-entity and business-associate framework (Source: pmc.ncbi.nlm.nih.gov), which is precisely the jurisdictional gap that made the FTC's Health Breach Notification Rule, rather than HIPAA, the applicable enforcement tool against GoodRx and Cerebral, discussed further in the case studies below. This gap extends even to informational conversations initiated voluntarily: there can be instances when patients engage in a customized conversation and share their own PHI with an AI chat tool for potential medical questions and recommendations, all outside HIPAA's reach (Source: pmc.ncbi.nlm.nih.gov).

Legal practitioners advising healthcare AI deployments increasingly frame BAA negotiation as inseparable from a broader data-mapping exercise rather than a one-time contract review. Morgan Lewis's healthcare AI practice advises that organizations must map how data enters, moves through, and exits AI systems, because a covered entity cannot satisfy HIPAA's Security Rule risk-analysis requirement without a clear, operational understanding of how AI systems actually interact with electronic PHI (Source: www.morganlewis.com). The same guidance is unambiguous that protected health information should not be used in public AI tools or to train general-purpose models, and that organizations should avoid or limit training models on PHI unless the model operates in a closed, proprietary environment developed for the covered entity (Source: www.morganlewis.com). Because AI systems change over time, software updates, configuration changes, and new integrations can alter how data flows and who has access to it, so data mapping is not a one-time compliance exercise but an ongoing obligation tied to HIPAA's requirement to regularly review and update security measures (Source: www.morganlewis.com).

Table 2 summarizes BAA coverage across the AI platforms most commonly evaluated by healthcare and life sciences organizations as of July 2026. Coverage is uneven and surface-specific: a platform being "HIPAA-eligible" in general does not mean every feature, tier, or beta product is covered, and organizations must confirm coverage feature by feature rather than assuming a blanket guarantee. It is also worth noting that no HHS-approved certification program for HIPAA compliance exists for cloud service providers; Microsoft, for instance, points customers to NIST SP 800-66, an introductory resource guide for implementing the HIPAA Security Rule, as the crosswalk it relies on to map its FedRAMP and ISO 27001 attestations to HIPAA Security Rule safeguards (Source: learn.microsoft.com). Google Cloud similarly emphasizes that while it provides a secure, BAA-covered infrastructure, the customer remains responsible for ensuring that the environment and applications built on top of that infrastructure are properly configured and secured according to HIPAA requirements (Source: cloud.google.com).

PLATFORM	HIPAA-ELIGIBLE SURFACES (REQUIRE SIGNED BAA)	NOTABLE EXCLUSIONS OR CONDITIONS
OpenAI / ChatGPT	ChatGPT for Healthcare, ChatGPT Enterprise with Regulated Workspace, ChatGPT FedRAMP, ChatGPT for Clinicians, API with Modified Retention, API FedRAMP with Modified Retention (Source: help.openai.com)	Free/Plus consumer ChatGPT is not HIPAA-eligible; additional workplace features can be enabled outside the BAA but are "intended only for uses that do not involve transmission, storage, or processing of PHI" (Source: help.openai.com)
Anthropic / Claude	Claude Enterprise (chat, Projects, Artifacts, file execution, voice, web search, Research, Skills) and 1P API, only after the Primary Owner activates HIPAA compliance and signs the BAA (Source: support.claude.com)	Consumer Free/Pro/Max plans excluded; MCP/Connectors, Enterprise Search, Claude in Chrome, Cowork, and several Claude Code surfaces are excluded or covered only with zero data retention enabled (Source: support.claude.com)
Microsoft Azure (incl. Azure OpenAI)	HIPAA BAA available by default via Microsoft's Product Terms and Data Protection Addendum to all customers who are covered entities or business associates (Source: learn.microsoft.com)	Only "in-scope" Azure services are covered; a signed BAA does not itself guarantee HIPAA compliance of the customer's deployed solution (Source: learn.microsoft.com)
Google Cloud (incl. Vertex AI, Gemini Enterprise)	BAA covers Google Cloud's entire infrastructure and a defined product list including Vertex AI Workbench, AutoML, Document AI, and Gemini Enterprise Agent Platform (Source: cloud.google.com)	Customers must avoid Pre-GA offerings for PHI and must proactively execute the BAA; Google Workspace HIPAA coverage is handled under a separate agreement (Source: cloud.google.com)
AWS (incl. Amazon Comprehend Medical)	Comprehend Medical is purpose-built to detect the 18 Safe Harbor identifier categories in clinical text for redaction (Source: docs.aws.amazon.com)	Confidence scores are probabilistic; AWS recommends "additional human review" for compliance-critical redaction rather than relying on automated detection alone (Source: docs.aws.amazon.com)

The practical takeaway from *Table 2* is that "can you use ChatGPT with patient data" or "is Claude HIPAA compliant" are not yes/no questions; they depend entirely on which product tier, which feature, and whether a BAA has actually been executed for that organization's account. This is exactly the operational distinction Mass General Brigham built into its institutional AI policy, permitting only approved internal tools (an internal "AI Zone" platform and Microsoft Copilot) for any work involving PHI, while restricting public tools like ChatGPT, Claude, and Perplexity to "limited, non-sensitive tasks" performed through a browser with no patient, research subject, or employee information involved, and banning DeepSeek outright (Source: libguides.massgeneral.org). Notably, MGB's policy prohibits feeding even de-identified data into a public LLM, a more conservative stance than HIPAA strictly requires, reflecting an institutional judgment that re-identification risk and reputational risk both warrant a wider berth than the regulatory floor (Source: libguides.massgeneral.org).

Implementation Guidance: Building a HIPAA-Compliant AI Workflow

Translating the legal framework above into an operational program requires a repeatable intake and governance process rather than one-off vendor evaluations. The following steps reflect the consensus approach recommended across HHS guidance, cloud-provider compliance documentation, peer-reviewed legal analysis, and institutional AI policies reviewed for this report:

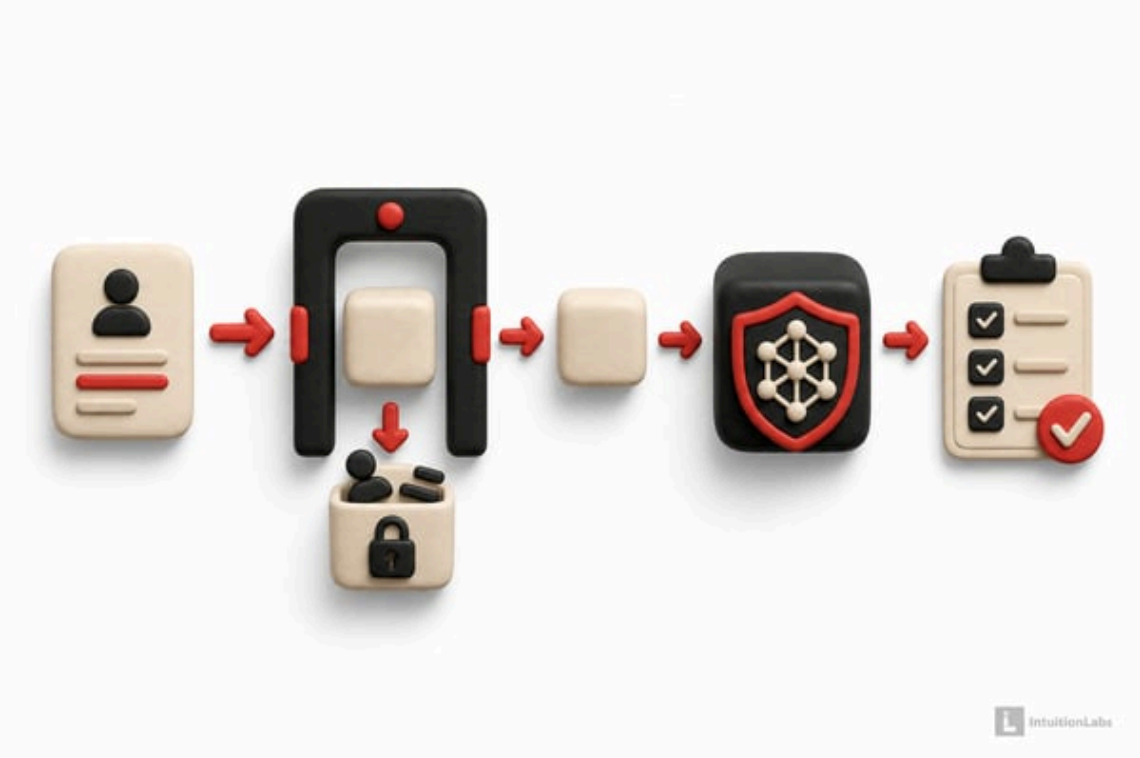

- **Map data flows before evaluating any tool.** Organizations must map how data enters, moves through, and exits AI systems, including what happens to outputs, logs, or embeddings retained by the vendor, because that mapping is the foundation on which a defensible risk analysis is built (Source: www.morganlewis.com).
- **Classify the relationship explicitly.** Confirm whether the organization is a covered entity, whether the AI vendor meets the business associate definition under 45 CFR 160.103, and whether the vendor's "HIPAA compliant" marketing claim is backed by an actual signed BAA, since a compliance claim never substitutes for the contract (Source: www.paubox.com).
- **Confirm feature-level BAA coverage, not just platform-level.** As Table 2 shows, beta features, connectors, and certain API endpoints are frequently excluded from an otherwise HIPAA-ready platform; every new feature enabled for a workspace should be checked against the vendor's current coverage documentation before PHI touches it (Source: support.claude.com).
- **De-identify wherever the use case allows it.** If the AI task (population analytics, research cohort building, benchmarking) does not require identifiable data, apply Safe Harbor or Expert Determination before data leaves the secure environment, rather than relying on the AI vendor's downstream safeguards (Source: www.hhs.gov).
- **Layer human review onto automated de-identification.** Automated PHI detectors such as Comprehend Medical or the Presidio analyzer, anonymizer, and image-redactor modules materially reduce manual effort but are explicitly not guaranteed to catch every identifier (Source: microsoft.github.io), particularly quasi-identifiers in free text and the catch-all "any other unique identifying" category, so spot-check review remains necessary for compliance-critical releases (Source: docs.aws.amazon.com).
- **Restrict public consumer AI tools to genuinely non-PHI tasks.** Follow the model MGB and similar institutions have adopted: permit public chatbots for generic, non-patient tasks accessed via browser, and route anything touching PHI, even de-identified, through an approved, contractually covered environment (Source: libguides.massgeneral.org).
- **Document the risk assessment for every de-identification decision.** Both HIPAA methods require documentation, either of the expert's statistical methodology and results, or of the covered entity's basis for concluding it lacks "actual knowledge" of residual re-identification risk (Source: www.hhs.gov).
- **Treat subcontractors and model hosts as downstream business associates.** A subcontractor that creates, receives, maintains, or transmits PHI on behalf of a business associate is itself a business associate and must be covered by appropriate agreements, so organizations should ask AI vendors directly how inference is hosted and whether any subcontractor accesses PHI (Source: www.ecfr.gov).
- **Define data rights contractually, not just data access.** AI vendor agreements should specify who owns and may use inputs, intermediate artifacts, and outputs, and who is responsible for securing consents or authorizations and for retaining or deleting data (Source:

www.morganlewis.com).

- **Restrict third-party access even within vendor models.** When proprietary or vendor AI models are used, PHI should not be accessible to third parties absent appropriate HIPAA-compliant agreements and safeguards, a standard that applies whether the model is hosted by the covered entity or licensed from an outside AI vendor (Source: www.morganlewis.com).
- **Evaluate re-identification risk separately from BAA status.** Organizations should assess heightened privacy risks associated with AI, including re-identification risk, which increases with large, commingled data sets, and account for state-law protections that may impose obligations beyond HIPAA even where a BAA is in place (Source: www.morganlewis.com).
- **Disclose AI use in plain language.** Clear, plain-language notices should describe how AI is used, the level of human oversight involved, and the system's limitations, so patients and staff understand when a tool, not a person, is processing their information (Source: www.morganlewis.com).
- **Refresh data maps as AI systems evolve.** Software updates, configuration changes, and new integrations can silently alter how data flows through an AI system and who can access it, so risk analyses and data-flow documentation need to be revisited on an ongoing basis rather than signed off once at go-live (Source: www.morganlewis.com).

For consultancies and system integrators advising healthcare and life sciences clients, this intake discipline is increasingly a prerequisite to any AI enablement engagement rather than an afterthought bolted on at the end. Firms operating in regulated pharma environments generally build compliance requirements, HIPAA, GDPR, FDA 21 CFR Part 11 electronic-records rules, and EU Annex 11, into the architecture of AI-enabled tooling from the outset rather than retrofitting it, because retrofitting a CRM extension or analytics layer to satisfy a BAA after PHI has already flowed through an uncovered surface is materially harder than designing the data flow correctly the first time (Source: intuitionlabs.ai).

Data Analysis and Evidence

The quantitative picture surrounding PHI, AI adoption, and enforcement risk has shifted markedly in the past three years. On the adoption side, the AMA's 2026 survey of nearly 1,700 physicians found 81% now use AI professionally, more than double the rate the AMA recorded when it first polled physicians on health AI in 2023 (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 (Source: www.ama-assn.org), and the JAMA Network signed a strategic content agreement with the OpenEvidence platform, which the AMA reports is used daily by an estimated 40% of U.S. doctors (Source: www.ama-assn.org). The same survey found seven in 10 physicians view AI as a tool to automate tasks that contribute to work-related burnout, and 76% say the technology can help with patient care, with the single most common reported clinical use, drafting discharge instructions, care plans, or progress notes, cited by 30% of surveyed physicians (Source: www.ama-assn.org). Yet the same survey found 88% of physicians harbor at least some concern about AI-related skill loss, and 85% want to be directly consulted on AI adoption decisions (Source: www.ama-assn.org), signaling that adoption is outpacing formal governance in many organizations. Physicians identified clear liability frameworks as the regulatory action most likely to build trust and increase adoption of AI tools, ranking it above other proposed safeguards (Source: www.ama-assn.org).

On the breach and enforcement side, OCR's public breach portal, which has tracked incidents affecting 500 or more individuals since October 2009, shows a clear escalation: 7,418 large healthcare data breaches were reported between 2009 and 2025, exposing the protected health information of more than 1,013,066,481 records, over 2.9 times the current U.S. population (Source: www.hipaajournal.com). Hacking and other IT incidents now account for more than 80% of large healthcare data breaches, a category that has grown sharply: OCR reported a 239% increase in hacking-related breaches and a 278% increase in ransomware attacks between January 2018 and September 2023 (Source: www.hipaajournal.com), even as other categories moved in the opposite direction: there has also been a downward trend in improper disposal incidents and unauthorized access or disclosure incidents, although the latter increased again in 2025 (Source: www.hipaajournal.com). The single largest breach on record, the 2024 Change Healthcare ransomware attack, affected an estimated 192.7 million individuals (Source: www.hipaajournal.com), and 2025 alone saw mega-breaches at Conduent Business Services (62.2 million individuals), Aflac (13.9 million), and Episource (6.7 million) (Source: www.hipaajournal.com). Early 2026 data suggests a modest deceleration in reported volume: from January 1 to April 30, 2026, 252 large healthcare data breaches were reported to OCR, 9.5% fewer than the corresponding period in 2025, though OCR itself cautions that reporting has lagged following a lengthy government shutdown in late 2025 (Source: www.hipaajournal.com).

On the financial-penalty side, OCR's 2026 civil monetary penalty structure, adjusted for inflation and published in the Federal Register on January 28, 2026, sets a Tier 1 ("reasonable efforts") minimum of $145 per violation and a Tier 4 ("willful neglect, uncorrected") ceiling of $2,190,294 per violation, with an identical annual cap (Source: www.hipaajournal.com). In practice, most enforcement in 2024 through 2026 has centered on risk-analysis failures, with settlements such as Montefiore Medical Center's $4.75 million (2024) for failing to conduct a comprehensive risk analysis and monitor information-system activity (Source: www.hipaajournal.com), BayCare Health System's $800,000 (2025) for information access management and minimum-necessary-standard failures (Source: www.hipaajournal.com), and Spencer Gifts LLC's $450,000 (2026) for risk-analysis failure and a lack

of HIPAA Privacy, Security, and Breach Notification Rule policies (Source: www.hipaajournal.com). These enforcement patterns are directly relevant to AI governance because an AI system that lacks a documented risk analysis, or whose data flows have never been mapped, is exactly the fact pattern OCR has repeatedly penalized when the underlying technology was far simpler than a generative AI pipeline.

Finally, on the AI-vendor commercialization side, the ambient AI scribe segment illustrates how quickly HIPAA-covered AI deployment has scaled inside major health systems: a JAMA-published multisite study across five academic medical centers using Ambience, Nuance DAX Copilot, and/or Abridge found ambient scribes reduced total EHR time by 13.4 minutes and documentation time by 16.0 minutes per clinician, while enabling 0.49 more patient visits per week (Source: www.aha.org). Separate JAMA-published research found Emory Healthcare saw a 30.7% increase in documentation-related well-being after adopting ambient documentation technology, and Mass General Brigham observed a 21.2% reduction in burnout prevalence after 84 days of use (Source: www.aha.org). Cooper University Healthcare in Camden, New Jersey found that Dragon Copilot saved clinicians 4.15 minutes in documentation time per patient, adding up to roughly an hour or more saved daily, while Intermountain Health in Salt Lake City saw a 27% reduction in time spent on notes per appointment among clinicians who used the tool for 10 or more encounters between April 2024 and December 2025 (Source: www.aha.org) (Source: www.aha.org). At Mercy, a health system based in Chesterfield, Missouri, one nurse using Dragon Copilot reported that the technology saved her about two hours of charting in a 12-hour shift (Source: www.aha.org). None of this adoption would be legally viable at scale without the underlying BAA infrastructure described in Table 2, underscoring how tightly PHI governance and AI value realization are now linked in practice.

Case Studies and Real-World Examples

GoodRx: The Cost of Claiming HIPAA Protection Without Covered-Entity Status

GoodRx, operating as GoodRx Gold, GoodRx Care, and Hey Doctor, disclosed millions of users' personal health information, including medication and health-condition details, to third-party advertisers between 2017 and 2020 without user authorization (Source: www.justice.gov). Critically, GoodRx marketed itself with a "HIPAA Secure: Patient Data Protected" seal while it was not, in fact, a HIPAA covered entity and never complied with HIPAA requirements (Source: www.justice.gov). The DOJ and FTC jointly resolved the case in February 2023 via a consent order requiring GoodRx to pay a $1.5 million civil penalty, notify affected users, and cease disclosing health information for advertising purposes going forward (Source: www.justice.gov). This case illustrates why "is this AI tool HIPAA compliant" is the wrong first question for many consumer-facing health platforms; the prior question is whether HIPAA applies at all, and if it does not, the FTC's Health Breach Notification Rule and state consumer-protection statutes fill the gap.

Cerebral: Mental Health Data Leaked Through Advertising Trackers

Telehealth provider Cerebral disclosed contact details, medical histories, insurance information, and prescription data for millions of users to third parties including Meta, Google, and TikTok between 2019 and 2023 through advertising pixel trackers embedded in its website and onboarding survey (Source: themarkup.org). In April 2024, the FTC announced a settlement requiring Cerebral to pay more than $7 million and accept a "first-of-its-kind" prohibition on using any health information for most advertising purposes (Source: themarkup.org). FTC Chair Lina Khan stated Cerebral "violated its customers' privacy by revealing their most sensitive mental health conditions across the Internet and in the mail" (Source: themarkup.org). Separately, Cerebral self-disclosed a breach to OCR affecting more than 3.1 million individuals (Source: www.hipaajournal.com). The lesson for AI governance is direct: analytics and tracking integrations embedded in patient-facing digital tools are exactly the kind of "AI-adjacent" data flow, feeding health-adjacent survey answers into third-party algorithmic ad-targeting systems, that can silently create PHI or FTC-regulated health-data exposure well before anyone considers whether a chatbot is involved.

MMG Fusion: Business Associate Breach Notification Failures at Scale

MMG Fusion, LLC, a Maryland software company that receives PHI from covered entities and communicates directly with patients on their behalf, experienced an intrusion in December 2020 in which an unauthorized actor accessed and later posted on the dark web the PHI of approximately 15 million individuals (Source: www.hhs.gov). OCR's March 2026 settlement, its 12th enforcement action under the agency's Risk Analysis Initiative, found MMG had impermissibly disclosed PHI, failed to conduct an accurate risk analysis, and failed to timely notify affected covered entities of the breach (Source: www.hhs.gov). MMG paid $10,000, a figure OCR explicitly reduced based on the company's financial condition, alongside a three-year corrective action plan (Source: www.hhs.gov). OCR Director Paula M. Stannard emphasized that timely breach notification by business

associates is "crucial for a covered entity to meet its own breach notification obligations" (Source: www.hhs.gov). Any AI vendor operating as a business associate inherits this exact obligation chain, and a health system evaluating an AI vendor's BAA should confirm the vendor's breach-notification timeline and risk-analysis practices, not just its data-encryption claims.

Mayo Clinic and Abridge: A Model for BAA-Governed Ambient AI at Scale

Mayo Clinic expanded its partnership with ambient AI documentation company Abridge in January 2025, initially rolling the tool out to roughly 2,000 clinicians serving about one million patients, with a phased expansion to additional eligible providers over the following year (Source: www.healthcaredive.com). Mayo, Abridge, and EHR vendor Epic had already collaborated on a related AI documentation product for nursing staff (Source: www.healthcaredive.com), and Mayo performed a "rigorous clinical note quality assessment" before deployment (Source: www.healthcaredive.com). This case demonstrates the intended end state of the framework described earlier in this report: a defined PHI data flow (recorded patient-clinician conversations), a named business associate under a signed BAA, a documented quality and safety review process, and a phased rollout tied to demonstrated clinical-note accuracy rather than a blanket organization-wide deployment on day one.

Advanced Care Hospitalists: The BAA Failure Pattern That Predates Generative AI

Not every cautionary tale in this space involves AI directly, and that is precisely the point. Advanced Care Hospitalists PL (ACH), a Florida physicians' group, had entered into a $500,000 settlement and resolution agreement with OCR announced on December 5, 2018 (Source: www.jdsupra.com). OCR's investigation found that ACH had impermissibly disclosed the protected health information of 9,255 patients to a third party for billing processing services without the protections of a business associate agreement (Source: www.jdsupra.com), and that ACH had been in business since 2005, but had not conducted any risk analysis or implemented any HIPAA privacy, security, and breach notification policies before 2014 (Source: www.jdsupra.com). The case predates generative AI by nearly a decade, but it is the exact fact pattern, PHI shared with a vendor absent a signed BAA, that recurs whenever an organization treats an AI product as "just another SaaS integration" rather than a business associate relationship requiring the same contractual diligence applied to any other vendor touching PHI.

Implications and Future Directions

Several trends will shape how "what counts as PHI when using AI" is answered over the next several years. First, enforcement is shifting from breach response toward proactive risk-analysis scrutiny: OCR's own year-end update confirmed 22 investigations resulted in civil monetary penalties or settlements in 2024 alone, one of its busiest enforcement years, and the agency has stated it will expand its risk-analysis initiative to include risk management in 2026 (Source: www.hipaajournal.com). AI vendors and their healthcare customers should expect that undocumented AI data flows will increasingly be treated the same way undocumented server infrastructure has been: as a risk-analysis failure in its own right, independent of whether a breach ever occurs.

Second, state law is filling gaps HIPAA leaves open. Washington's My Health My Data Act, signed in April 2023, was explicitly designed to "close the gap" between HIPAA's coverage and the broader universe of consumer health data collected outside traditional covered entities, requiring consent or necessity as a legal basis for collecting or sharing consumer health data (Source: iapp.org). The breadth of covered data under the Act is significant for AI vendors and their customers because the law recognizes a private right of action allowing consumers to sue companies directly for violating any of its provisions (Source: iapp.org). Comparable consumer-health-data provisions have since been introduced or enacted in other states, meaning AI tools that process health-adjacent data entirely outside a HIPAA covered relationship, wellness apps, symptom checkers, direct-to-consumer genetic services, still face a growing patchwork of state obligations that mirror HIPAA's spirit without HIPAA's narrower jurisdictional trigger. Even where HIPAA does not apply at all, entities inputting health data into AI tools still need to account for state law requirements and common-law privacy expectations, since individuals retain a reasonable expectation of privacy in their personal information regardless of whether a HIPAA covered relationship exists (Source: www.morganlewis.com). Peer-reviewed analysis of these gaps adds a further wrinkle: even PHI that originates inside a covered entity's own AI-enabled platform loses HIPAA protection the moment a patient exports or transfers it into an uncovered space, such as a personal health device outside the covered entity's control (Source: pmc.ncbi.nlm.nih.gov).

Third, the AI vendor landscape is bifurcating between platforms building dedicated, contractually gated healthcare tiers (OpenAI's ChatGPT for Healthcare, Anthropic's HIPAA-ready Claude Enterprise, the hyperscalers' infrastructure-wide BAAs) and a long tail of point-solution vendors, ambient scribes, clinical NLP tools, and specialty diagnostic AI, that must each independently demonstrate BAA readiness. Legal and compliance advisors increasingly recommend that organizations confirm BAAs explicitly permit access to PHI by every entity involved in the AI operating environment, including subcontractors and hosting providers, rather than assuming a single signature covers the entire technical stack (Source: www.morganlewis.com). For life sciences organizations layering AI onto commercial, clinical, and regulatory systems built on platforms like Veeva

Vault CRM, the practical implication is that AI enablement projects increasingly need to be scoped jointly with data-governance and compliance functions from the earliest design phase, rather than treated as a pure technology integration exercise, since the underlying question of what counts as PHI, and therefore what a BAA must cover, has to be resolved before, not after, a workflow goes live.

Fourth, expect continued divergence between what HIPAA technically requires and what leading institutions actually permit. Mass General Brigham's decision to bar even de-identified data from public LLMs (Source: libguides.massgeneral.org) reflects a broader institutional trend toward risk postures more conservative than the regulatory floor, driven by reputational risk, residual re-identification concerns, and the difficulty of auditing exactly what a third-party model does with submitted text after the fact.

Frequently Asked Questions (FAQs)

What is the difference between PHI and PII? PII (personally identifiable information) is any data that can identify a specific individual across any industry context; PHI is the narrower subset of identifiable information that also relates to health, treatment, or payment and is held by a HIPAA covered entity or business associate (Source: www.foreseemed.com). All PHI is PII, but not all PII is PHI: a name and phone number in a public phone book is PII but not PHI, while the same name and phone number recorded alongside a diagnosis in a medical record is PHI (Source: www.hipaajournal.com).

Is AI output considered PHI under HIPAA? Yes, if it is derived from individually identifiable health information and remains capable of identifying the individual. The Privacy Rule's definition of PHI turns on whether information is transmitted or maintained by a covered entity or business associate, not on what technology created it (Source: www.ecfr.gov).

Can you use ChatGPT with patient data? Only within OpenAI's HIPAA-eligible products (ChatGPT for Healthcare, ChatGPT Enterprise with Regulated Workspace, ChatGPT for Clinicians, or specific API endpoints with Modified Retention) after a Business Associate and Healthcare Addendum has been executed (Source: help.openai.com). Free and Plus consumer ChatGPT accounts are not HIPAA-eligible surfaces and should not receive PHI.

What are the 18 HIPAA identifiers? Names; geographic subdivisions smaller than a state; all elements of dates (except year) tied to an individual, plus ages over 89; telephone numbers; fax numbers; email addresses; Social Security numbers; medical record numbers; health plan beneficiary numbers; account numbers; certificate or license numbers; vehicle identifiers; device identifiers; web URLs; IP addresses; biometric identifiers; full-face photographs; and any other unique identifying number, characteristic, or code (Source: www.hhs.gov) (see Table 1 above for the full list with AI-workflow context).

What is the Expert Determination method? A statistical de-identification path in which a qualified expert applies generally accepted scientific and statistical methods to determine that re-identification risk is very small, and documents that determination, as an alternative to removing all 18 Safe Harbor identifiers outright (Source: www.hhs.gov).

Does every AI tool that touches health data require a Business Associate Agreement? Only when the organization providing or using the AI tool is a HIPAA covered entity or business associate and the tool creates, receives, maintains, or transmits PHI on that entity's behalf (Source: www.ecfr.gov). If the data has been properly de-identified under Safe Harbor or Expert Determination, or if the tool's provider is not acting on behalf of a covered entity, HIPAA's BAA requirement does not apply, though other laws such as state consumer-health-data statutes may still govern.

Conclusion

What counts as PHI when using AI is governed by the same three-part test that has applied to paper records for two decades: does the information relate to health, treatment, or payment; is it individually identifiable; and is it held by a covered entity or its business associate. AI does not change that test, but it does change how quickly and easily identifiable health information can move outside a protected environment, whether through a well-intentioned clinician pasting a chart excerpt into a public chatbot, an ambient scribe recording a patient encounter, or a commercial-operations analyst uploading a data extract into an unsanctioned AI copilot. The 18 Safe Harbor identifiers give organizations a concrete checklist for de-identification, the Expert Determination method offers a more flexible statistical path where the checklist is too blunt, and a signed, feature-specific Business Associate Agreement remains the non-negotiable gate before PHI reaches any third-party AI system.

The enforcement record, from GoodRx's false HIPAA marketing claim to Cerebral's pixel-driven mental-health data leak to MMG Fusion's breach-notification failures to Advanced Care Hospitalists' decade-long BAA gap, shows that most penalties trace back not to sophisticated AI misuse but to the same fundamentals that have always driven HIPAA enforcement: missing risk analyses, undocumented data flows, and contractual gaps with vendors. Organizations that treat AI evaluation as an extension of existing vendor risk management, mapping data flows, confirming feature-level BAA coverage, de-identifying wherever the use case allows, and documenting every judgment call, are positioned to capture the documented productivity

gains of tools like ambient AI scribes without repeating the compliance failures that have defined two decades of HIPAA enforcement. As AI adoption continues to outpace formal governance across the healthcare and life sciences sector, the organizations that will avoid becoming the next enforcement case study are those that answer the PHI question before deployment, not after an incident forces the answer on them.

Tags: protected health information, phi and ai, hipaa de-identification, safe harbor rule, hipaa compliant ai tools, phi vs pii, business associate agreement, expert determination method, 18 hipaa identifiers, hipaa 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.