Is Claude HIPAA Compliant? BAA, Enterprise & API Rules (2026)

Published July 4, 2026 40 min read

Is Claude HIPAA Compliant? BAA, Enterprise & API Rules (2026)

Executive Summary

Claude, the AI model family built by Anthropic, is not HIPAA compliant by default on any of its consumer plans, and compliance for its business plans depends entirely on which product surface an organization uses and whether a Business Associate Agreement (BAA) has actually been signed and activated. As of July 2026, Anthropic states plainly that it "provides a BAA covering our HIPAA-ready services, such as use of our first-party API or Enterprise plans" (Source: privacy.claude.com), but "standard Claude Enterprise plans do not include BAA coverage without action from a Primary Owner" (Source: privacy.claude.com). Free, Pro, and Max, the individual subscription tiers priced from $0 to $100 or more per month, carry no BAA option at all, and "Team plans and individual plans (Free, Pro, and Max) can't enable HIPAA" (Source: support.claude.com). Only Claude Enterprise, which starts at $20 per seat per month billed annually with a 20-seat minimum (Source: claude.com), and the Claude Developer Platform (API) configured as a HIPAA-ready organization, can be brought under a BAA.

Even within Enterprise and the API, coverage is granular rather than blanket. Anthropic's own feature table shows core Chat, Projects, Artifacts, and the Messages API as eligible, while Claude Cowork, Claude Console, Workbench, and several beta surfaces such as Claude in Office and Claude Design are explicitly excluded, and the BAA "excludes features such as Workbench, Claude Console, Claude Cowork, or features currently in beta" (Source: privacy.claude.com). Claude Code, Anthropic's coding agent, is only covered "with zero data retention (ZDR) enabled, and only on qualified accounts" (Source: support.claude.com), and Anthropic confirms "there's currently no configuration that allows BAA-covered access to Covered Models in Claude Code or Cowork" (Source: support.claude.com). Anthropic holds HIPAA-ready configuration alongside SOC 2 Type I and Type II, ISO 27001:2022, and ISO/IEC 42001:2023 credentials (Source: privacy.claude.com), documented on its public Trust Center (Source: trust.anthropic.com).

This report compares Claude's posture against OpenAI's ChatGPT, where "only ChatGPT Enterprise or Edu customers that have a sales-managed account are eligible for a BAA" and OpenAI "don't offer a BAA for ChatGPT Business" (Source: help.openai.com) (Source: help.openai.com), against Microsoft Copilot, which is covered under Microsoft's standing HIPAA/HITECH BAA program for in-scope services (Source: learn.microsoft.com), and against Google Cloud's Gemini and Vertex AI offerings, where Google states it "will enter into Business Associate Agreements with customers as necessary under HIPAA" for specifically enumerated covered products (Source: cloud.google.com). It documents Anthropic's HIPAA enforcement backdrop, including the fact that the U.S. Department of Health and Human Services Office for Civil Rights (OCR) had, as of late 2024, settled or imposed civil penalties in 152 cases totaling $144,878,972 (Source: www.hhs.gov), and that 2025 closed with 21 OCR settlements and civil monetary penalties, the second-highest annual total on record (Source: www.hipaajournal.com). It walks through named deployments, including Novo Nordisk's use of Claude Code to cut clinical study report drafting from "10+ weeks to 10 minutes" (Source: claude.com) and Banner Health's reported BannerWise deployment using Claude to more than 55,000 employees (Source: www.beckershospitalreview.com). The bottom line: Claude can be used in a way that supports HIPAA compliance, but only on the Enterprise plan or a HIPAA-ready API organization, only after a signed BAA is activated, and only within a carefully bounded set of features; HIPAA-regulated covered entities and business associates should not submit protected health information (PHI) through Claude.ai Free, Pro, Max, Team, or any feature outside an applicable BAA and approved configuration.

Introduction and Background

The question "is Claude HIPAA compliant" is asked so often because it conflates two separate ideas: whether an AI vendor is capable of supporting HIPAA-covered workflows, and whether a specific product configuration is currently covered by a signed agreement. The Health Insurance Portability and Accountability Act of 1996 (HIPAA) does not certify software products; there is no such thing as a HIPAA "compliant" stamp that a vendor earns once. Instead, HIPAA imposes obligations on covered entities (health plans, healthcare clearinghouses, and most healthcare providers) and their business associates, defined by the U.S. Department of Health and Human Services as any "person or entity that performs certain functions or activities that involve the use or disclosure of protected health information on behalf of, or provides services to, a covered entity" (Source: www.hhs.gov). HHS's own list of illustrative business associates, ranging from a claims-processing third-party administrator to "an independent medical transcriptionist that provides transcription services to a physician" (Source: www.hhs.gov), shows that the underlying legal category long predates AI vendors; Anthropic, OpenAI, Microsoft, and Google simply fall into a category HHS has regulated for two decades. Before any PHI can lawfully flow to a vendor such as Anthropic, the covered entity must obtain "satisfactory assurances" from that vendor, and those assurances "must be in writing, whether in the form of a contract or other agreement" (Source: www.hhs.gov). That written contract is the BAA.

Anthropic launched Claude in 2023 and has since built a substantial healthcare and life sciences business, expanding from Claude for Life Sciences in October 2025 into Claude for Healthcare in January 2026, timed to the J.P. Morgan Healthcare Conference (Source: www.fiercehealthcare.com). As IntuitionLabs' own technical review of the space frames it, Anthropic "announced 'Claude for Healthcare' and an expansion of 'Claude for Life Sciences,' delivering HIPAA-compliant AI tools and domain-specific integrations" (Source: intuitionlabs.ai). That expansion sits alongside a growing library of compliance credentials: a HIPAA-ready configuration with a BAA available, ISO 27001:2022 for information security management, ISO/IEC 42001:2023 for AI management systems, and SOC 2 Type I and Type II reports (Source: privacy.claude.com) (Source: privacy.claude.com). These credentials are published on the Anthropic Trust Center, where the company states, "Here you can find our compliance artifacts, request documentation, and view high-level details on controls we adhere to" (Source: trust.anthropic.com).

At the same time, healthcare and life sciences organizations remain the single most expensive sector for data breaches, with the average breach in the industry costing 10.93 million USD according to the IBM and Ponemon Institute Cost of a Data Breach study, and healthcare breaches taking 213 days to discover on average, longer than the 194-day cross-industry mean (Source: www.ibm.com) (Source: www.ibm.com). Against that backdrop, getting the answer to "is Claude HIPAA compliant" precisely right, plan by plan and feature by feature, is not an academic exercise; it is the difference between a defensible compliance posture and a reportable breach. This report works through every Claude product surface, the mechanics of Anthropic's BAA, how Claude compares to ChatGPT, Microsoft Copilot, and Google Gemini/Vertex AI, and what a healthcare organization or health-tech developer actually needs to do to use Claude lawfully with PHI as of July 2026.

What "HIPAA Compliant" Actually Means for an AI Vendor

HIPAA regulates the handling of protected health information, defined broadly as individually identifiable health information created, received, maintained, or transmitted by a covered entity or its business associates. No federal certification exists that a software vendor can obtain to be officially declared "HIPAA compliant"; Google's own compliance documentation is explicit that "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), and the same principle applies to every AI vendor discussed in this report. Microsoft echoes this: "there's currently no certification standard that the Department of Health and Human Services approves to demonstrate compliance with HIPAA or the HITECH Act by a business associate" (Source: learn.microsoft.com).

What exists instead is a **Business Associate Agreement**, a legally binding contract under which the vendor accepts specific obligations to safeguard PHI. A BAA under HIPAA must, at minimum, describe the permitted and required uses of PHI, prohibit the business associate from using or disclosing PHI beyond what the contract or law allows, and require appropriate safeguards against misuse, per HHS's own contract-provision guidance (Source: www.hhs.gov). Signing a BAA with Anthropic, OpenAI, Microsoft, or Google does not, by itself, make a healthcare organization's overall system HIPAA compliant. As Aptible's own HIPAA guidance for Claude puts it, "'HIPAA eligible' means Anthropic will sign a BAA. 'HIPAA compliant' means your full implementation meets HIPAA requirements. Getting a BAA is necessary" but not sufficient (Source: www.aptible.com). HIPAA Vault reinforces the same boundary from the technical-safeguards side, noting that "HIPAA's Security Rule requires safeguards such as access controls, audit mechanisms, and transmission security. Claude is designed as a general-purpose AI assistant, not a regulated healthcare platform by default" (Source: www.hipaavault.com).

Applied to Claude specifically, three conditions must all be true before an organization can lawfully send PHI into any Claude surface:

- **A BAA must be signed.** For Anthropic, that means "your organization's Primary Owner will need to sign a BAA and then reach out to your Anthropic contact or our Sales team" for API access (Source: privacy.claude.com), or activating the click-to-accept flow for Enterprise.
- **The specific surface used must be an "Eligible Service" under that BAA.** Anthropic's BAA "only covers the single organization that accepted it" and excludes named products entirely (Source: privacy.claude.com).
- **The organization must independently implement HIPAA's technical safeguards**, including audit logging, access controls, and encryption, none of which Anthropic's BAA supplies on the customer's behalf. These obligations apply to any system that handles PHI, including the calls an application makes to an LLM.

Understanding this three-part test is the key to answering every variant of "is Claude HIPAA compliant," because the honest answer changes completely depending on which Claude product, which feature within that product, and which organizational configuration a person means.

Claude Free, Pro, and Max: Consumer Plans Cannot Process PHI

Claude's individual subscription plans, **Free** ($0) (Source: claude.com), **Pro** ($17 per month with an annual subscription discount, $20 if billed monthly), and **Max** (from $100 per month) (Source: claude.com), have no BAA path whatsoever. Anthropic's privacy documentation is explicit that its BAA article "is about our commercial products such as Claude for Work and the Anthropic API. For our consumer products such as Claude Free, Pro, Max and when accounts from those plans use Claude Code," a separate consumer privacy collection applies instead (Source: privacy.claude.com), and that consumer collection contains no BAA offering at all (Source: privacy.claude.com).

The HIPAA-ready Enterprise support article states this limitation even more directly for the enablement flow itself: "You can enable the HIPAA configuration from organization settings if your organization is on an Enterprise plan. Team plans and individual plans (Free, Pro, and Max) can't enable HIPAA" (Source: support.claude.com). Practically, this means a clinician, therapist, or administrator using their personal Claude Pro or Max account to summarize a chart, draft a discharge note, or paste in a patient message is operating entirely outside any contractual PHI protection, regardless of how careful Anthropic's underlying infrastructure may be. Third-party HIPAA compliance vendor HIPAA Vault states the general principle bluntly: "Claude AI is not HIPAA compliant by default" and "should not be used to store, process, or transmit PHI unless it is deployed under Anthropic's HIPAA-ready Enterprise plan with a signed Business Associate Agreement (BAA)" (Source: www.hipaavault.com) (Source: www.hipaavault.com).

It matters that this restriction extends to features bundled into consumer plans that sound enterprise-grade, such as Claude Code, Claude Cowork, and Claude Design, all included in Pro and Max. None of these change the underlying answer: as long as the account is a Free, Pro, or Max subscription, there is no BAA, and therefore no lawful path for PHI, no matter how the product is marketed for "everyday productivity." Anthropic also offers a separate consumer health-data integration, allowing individual Pro and Max subscribers in the US to connect personal lab results and health records through HealthEx, Function, Apple Health, and Android Health Connect (Source: www.anthropic.com), but this is a consumer feature governed by consumer terms of service, not a HIPAA business associate relationship, and it is designed for individuals managing their own health information rather than covered entities processing patients' PHI on their behalf.

Claude Team: Enterprise Features Without HIPAA Coverage

The **Team** plan sits between individual subscriptions and Enterprise, priced at $20 per seat per month billed annually ($25 monthly) for a Standard seat, and $100 per seat per month billed annually ($125 monthly) for a Premium seat with five times the usage (Source: claude.com) (Source: claude.com), designed for organizations of "2 to 150" people. Team includes single sign-on, central billing, connectors, and no model training by default, which makes it look, on the surface, like a plausible home for regulated data. It is not. Anthropic's Trust Center compliance matrix lists Claude Team as covered by SOC 2, ISO 27001, ISO 42001, and CSA STAR, but its HIPAA column is marked "N/A," in contrast to Claude Enterprise and the API, both of which show HIPAA as a supported scope (Source: trust.anthropic.com).

This is confirmed directly in the HIPAA-ready Enterprise plans documentation, which groups Team with the individual plans as unable to enable HIPAA at all (Source: support.claude.com). Aptible's compliance guide reaches the same conclusion in its feature-by-feature comparison table, listing "Claude.ai Team" as having no BAA available. This is a recurring point of confusion in practice: a real Reddit thread from a healthcare-adjacent buyer describes exactly this trap, recounting that after finally hearing back from Anthropic about a HIPAA inquiry, "they said go teams," without addressing the HIPAA question at all (Source: www.reddit.com). The same commenter, evaluating options for a team of fewer than twenty people, wrote: "Would love to have an off the shelf HIPAA compliant option for their teams plan that is self enrolled... We have a team of less than 20 and wanted the healthcare option" (Source: www.reddit.com). Organizations that need HIPAA coverage but have fewer than the 20-seat Enterprise minimum face a genuine structural gap in Anthropic's current plan lineup: there is no self-serve, HIPAA-ready option below Enterprise scale, and Team, despite its enterprise-style admin controls, is not an appropriate option for a HIPAA-regulated workflow that sends PHI to Anthropic because the plan cannot enable HIPAA or be covered by Anthropic's BAA.

Claude Enterprise: The HIPAA-Ready Offering and How Activation Works

Claude Enterprise is the plan where HIPAA coverage becomes possible, and it is available in two forms: self-serve and sales-assisted, both priced at $20 per seat per month billed annually, with usage billed separately at API rates and a 20-seat minimum (Source: claude.com). Enterprise on its own, however, does not include a BAA. Anthropic's HIPAA-ready documentation calls this out directly: "For Claude Enterprise features to be covered under a Business Associate Agreement (BAA), the Primary Owner of the organization must activate HIPAA compliance in the HIPAA-ready Claude Enterprise organization settings under 'Data and privacy' and accept Anthropic's BAA. Standard Claude Enterprise plans do not include BAA coverage without action from a Primary Owner" (Source: privacy.claude.com).

The offering itself is framed around a specific population: "This offering is designed for HIPAA-covered entities and their business associates, including: healthcare providers (e.g., hospitals, clinics, physicians); health plans and insurers; healthcare data processors; business associates that handle PHI on behalf of covered entities; other HIPAA regulated entities" (Source: support.claude.com). Activation is a self-serve flow, not a negotiated contract: "Eligible Enterprise organizations can enable HIPAA-ready configuration directly from organization settings, no sales or legal cycle required. The Business Associate Agreement (BAA) is included in the flow as click-to-accept, so there's no separate document to sign and return" (Source: support.claude.com). Only one person can do this, however: "Only the Primary Owner of the organization can accept the BAA and enable HIPAA. Other Owners or Admins can't complete this flow on the org's behalf" (Source: support.claude.com), and once accepted, it cannot be undone: "This is a one-way decision. Once HIPAA is enabled and the BAA is accepted, the change can't be reversed from organization settings" (Source: support.claude.com).

Once active, HIPAA-ready Claude Enterprise covers a meaningful set of daily-use features: users "can chat with Claude, create projects and artifacts, and use voice mode," and depending on admin configuration can also use "connectors, enterprise search, file creation and code execution, web search, research, and skills" (Source: support.claude.com). But two carve-outs are worth emphasizing because they surprise buyers. First, "enabling HIPAA readiness alone doesn't bring Claude Code under your BAA. Claude Code is covered under your BAA only with zero data retention (ZDR) enabled, and only on qualified accounts" (Source: support.claude.com). Second, "Cowork is not yet covered under Anthropic's BAA" in any configuration (Source: support.claude.com). Anthropic's own enterprise sales page frames the question as a headline FAQ item, "Do you support HIPAA? Can we get a BAA?" (Source: claude.com), underscoring how central this question is to enterprise healthcare buyers.

Timing also matters for organizations with legacy contracts. "If your organization signed a BAA for Claude API usage before December 2, 2025, that agreement only covers API usage, it does not extend to the HIPAA-ready Enterprise plan. To add this Enterprise plan access, you'll need to sign a new BAA with your account team. BAAs signed after December 2, 2025 can cover both API usage and the Enterprise plan under a single agreement" (Source: support.claude.com). Anthropic supports these enablement conversations through a dedicated healthcare sales channel and documents the flow in an Implementation Guide for HIPAA Entities available on the Trust Center, downloadable during setup (Source: support.claude.com).

Claude API and Developer Platform: HIPAA-Ready Organizations, ZDR, and Covered Models

For developers building products on top of Claude rather than using the chat interface, the relevant surface is the Claude Developer Platform, formerly called the Claude API. Anthropic offers two distinct data-handling arrangements here, and understanding the difference is essential: "Anthropic offers two data handling arrangements for the Claude API: zero data retention (ZDR) and HIPAA readiness" (Source: platform.claude.com). These are not interchangeable, and Anthropic is explicit that healthcare workloads should use the HIPAA arrangement, not ZDR: "HIPAA readiness applies a broader set of privacy and security safeguards than ZDR (encryption, access controls, and audit logging that protect PHI throughout its lifecycle) rather than requiring immediate deletion. If your organization handles PHI, HIPAA readiness is the arrangement to use; you do not also need ZDR" (Source: platform.claude.com).

Setting up HIPAA-ready API access is, like Enterprise, largely self-serve: "Most organizations can enable it directly in the Claude Console with Anthropic's standard BAA; organizations that require a negotiated BAA should work with their account team" (Source: platform.claude.com). Once enabled, it is permanent: "Enablement takes effect immediately... Once HIPAA readiness is enabled for your organization, the configuration is permanent and cannot be disabled by an administrator" (Source: platform.claude.com). The API enforces this at the request level; sending a call that uses a non-eligible feature returns an HTTP 400 error such as, "The requested features are not available for HIPAA-regulated organizations without Zero Data Retention: code_execution" (Source: platform.claude.com).

A long list of API capabilities carries different eligibility depending on the arrangement. Prompt caching, structured outputs, extended thinking, the 1M-token context window, the memory tool, PDF support sent inline through the Messages API, and citations are all HIPAA-eligible; by contrast, Batch processing, the Files API, Agent Skills via the Skills API, Code execution, Computer use, the MCP connector, and Claude Managed Agents are excluded from both HIPAA readiness and ZDR because they are inherently stateful services that retain jobs, files, or container state (Source: platform.claude.com). Crucially, "Claude Code is not covered under HIPAA readiness" at the API level at all (Source: platform.claude.com); on Enterprise it can only be brought under the BAA through the separate ZDR path described above, and even then, doing so has a side effect: it blocks access to Anthropic's newest, most capable models.

That is because of Anthropic's **Covered Model** designation. As of July 2026, the two models so designated are Claude Mythos 5 (limited access to approved partners) and Claude Fable 5 (generally available across Claude applications, the Claude Platform, Amazon Bedrock, Google Cloud Agent Platform, and Microsoft Foundry) (Source: support.claude.com) (Source: support.claude.com). Covered Models "require 30-day data retention on every platform where they're offered, as part of our safety work, and they can't be accessed from organizations or workspaces with zero data retention (ZDR) enabled" (Source: support.claude.com). This creates a genuine tradeoff for regulated customers: a HIPAA-ready API organization using standard retention can access Covered Models under the BAA, but a Claude Code deployment covered via ZDR cannot, so "requests to a Covered Model from a ZDR-enabled organization or workspace return an error" (Source: support.claude.com). Organizations that need both BAA-covered Claude Code and BAA-covered access to Covered Models must run separate Anthropic organization IDs with different configurations, since "HIPAA readiness and ZDR cannot coexist on a single 1P API organization" (Source: support.claude.com).

What Anthropic's BAA Covers and Excludes: A Feature-by-Feature View

Because Anthropic's BAA operates at the level of individual product surfaces rather than blanket account coverage, the most reliable way to assess "is Claude HIPAA compliant" for a specific use case is to check the exact feature against Anthropic's published eligibility table. *Table 1* below summarizes the coverage status of the most commonly used Claude Enterprise and Claude Platform features as of July 2026, condensed from Anthropic's official BAA documentation.

FEATURE / SURFACE	BAA COVERAGE STATUS	SOURCE DETAIL
Chat, Projects, Artifacts (Claude Enterprise)	Covered	"Covered as Eligible Services under Anthropic BAA" (Source: privacy.claude.com)
File creation & code execution (Enterprise)	Covered, with exclusion	Covered "excluding network access and use of external websites" (Source: privacy.claude.com)
Voice, Web Search, Research, Skills (Enterprise)	Covered	Listed as Eligible Services under the BAA (Source: privacy.claude.com)
MCPs / Connectors, Enterprise Search, Claude in Chrome	Partial	"Available to use but sending data to 3rd parties via this feature isn't covered under Anthropic's BAA" (Source: privacy.claude.com)
Cowork	Not covered	"Cowork is not an Eligible Service under Anthropic's BAA in any configuration" (Source: support.claude.com)
Claude for Office, Claude Design (beta)	Not covered	"Available to use but feature is in beta and not covered under Anthropic's BAA" (Source: privacy.claude.com)
Claude Code CLI (via 1P API console or Enterprise OAuth)	Covered only with ZDR	"Only covered under the BAA with ZDR enabled" (Source: privacy.claude.com)
Claude Code (Desktop remote, Web, Review, Security, Computer Use)	Not covered	"This feature is incompatible with ZDR" (Source: privacy.claude.com)
Messages API core (prompt caching, structured outputs, memory, web search, bash tool, text editor)	Covered	"Covered as Eligible Services under Anthropic BAA" (versions accepted after April 1, 2026) (Source: privacy.claude.com)
Batch API, Files API, Skills API, Code Execution, Computer Use, Web Fetch (1P API)	Not covered	"Not covered under Anthropic BAA and not accessible for HIPAA-Ready API users" (Source: privacy.claude.com)

This table demonstrates a consistent pattern: features that are stateless, single-request, and fully processed within Anthropic's own infrastructure tend to be BAA-eligible, while features that persist state, route data to third parties, or remain in beta tend to be excluded. Buyers should treat the table as a checklist to run against their specific implementation rather than assume that "Claude Enterprise has a BAA" answers the question for every feature their team happens to use. A common and costly mistake, according to compliance vendor Aptible, is enabling connectors or Claude Code without separately verifying ZDR status, since "if your teams work across multiple organizations with different coverage, you're responsible for ensuring PHI is only submitted through Eligible Services" (Source: support.claude.com). Anthropic's Trust Center offering a dedicated Implementation Guide for HIPAA Entities exists precisely to walk administrators through these interactions in more depth than a summary table can capture (Source: support.claude.com).

Claude vs. ChatGPT vs. Microsoft Copilot vs. Google Gemini: HIPAA Across the Major AI Vendors

For teams evaluating Claude against alternatives, the question "is ChatGPT HIPAA compliant vs Claude" comes up constantly, and the honest answer is that all four major frontier AI vendors follow the same basic shape: no BAA on free or self-serve consumer tiers, a BAA available on managed enterprise tiers, and specific feature exclusions layered on top. *Table 2* below compares the four vendors most healthcare buyers are actually choosing between as of July 2026.

VENDOR	CONSUMER / SELF-SERVE TIER	ENTERPRISE / MANAGED TIER BAA PATH	NOTABLE GAPS
Anthropic Claude	No BAA on Free, Pro, Max, or Team (Source: support.claude.com)	Self-serve click-to-accept BAA on Enterprise, or sales-assisted API HIPAA readiness (Source: support.claude.com)	Cowork, Console, Workbench, Claude Code without ZDR, most beta features not covered (Source: privacy.claude.com)
OpenAI ChatGPT	No BAA for Free, Plus, Team, or ChatGPT Business (Source: help.openai.com)	BAA available for ChatGPT Enterprise and Edu with a sales-managed account, and for the API via case-by-case review, since "an enterprise agreement is not required to sign a BAA" for API services (Source: help.openai.com); a dedicated ChatGPT for Healthcare product launched January 2026 (Source: openai.com)	ChatGPT Health, the consumer wellness product, is never BAA-eligible, "as the purpose of the product is to support personal health literacy rather than regulated healthcare operations" (Source: www.hipaajournal.com)
Microsoft Copilot	Not applicable to consumer Copilot	Microsoft "offers its covered entity and business associate customers a Business Associate Agreement that covers in-scope Microsoft services" (Source: learn.microsoft.com), including "Microsoft Copilot for Security" and Microsoft 365 Copilot by name among in-scope services (Source: learn.microsoft.com)	Microsoft states plainly, "No, Microsoft can't use a customer's Business Associate Agreement" (Source: learn.microsoft.com), since it operates one standardized BAA across its hyperscale services
Google Gemini / Vertex AI	Not applicable to consumer Gemini app	Google "will enter into Business Associate Agreements with customers as necessary under HIPAA" for enumerated Covered Products on Google Cloud, including "Vertex AI Workbench instances" (Source: cloud.google.com)	Customers must "disable or otherwise ensure that you do not use Google Cloud Products that are not explicitly covered by the BAA... when working with PHI" and must not use "Pre-GA offerings... in connection with PHI, unless expressly noted otherwise" (Source: cloud.google.com)

The pattern across all four vendors is that the free or lightly managed consumer product is never HIPAA-eligible, and the differences that matter to buyers show up in the details: how self-serve the BAA activation is, how many product features are excluded from the agreement, and whether the vendor has built healthcare-specific tooling on top of the base offering. Anthropic's self-serve, click-to-accept BAA flow for Enterprise (Source: support.claude.com) is comparatively fast next to OpenAI's model, where API BAA requests are reviewed case by case: OpenAI states that "we review each BAA request on a case-by-case basis and may need additional information" (Source: help.openai.com), and reconsideration for declined use cases "is only available for customers that are working with a member of our sales team" (Source: help.openai.com). In practice, however, buyers report friction with Anthropic's own process too; one Reddit account from a prospective Claude Enterprise HIPAA customer described submitting "the BAA form for enrolling our company in the HIPAA-ready Claude Enterprise plan about a month ago, and we haven't heard anything since from Anthropic" (Source: www.reddit.com). Aptible's own guidance for developers reflects the same experience at scale: "if you're on the API or self-managed path, plan for it to take weeks, not days. Anthropic's compliance process moves more slowly than cloud providers like AWS, where a BAA is part of the standard enterprise agreement" (Source: www.aptible.com).

On the question specifically of "is ChatGPT HIPAA compliant" relative to Claude, the HIPAA Journal's independent analysis draws a distinction worth preserving carefully: "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. However, OpenAI now offers ChatGPT for Healthcare that can support HIPAA compliance under specific conditions" (Source: www.hipaajournal.com). This mirrors Claude's own structure almost exactly: general chat products are excluded; a purpose-built enterprise healthcare product is not. OpenAI's ChatGPT for Healthcare rolled out in January 2026 to institutions including "AdventHealth, Baylor Scott & White Health, Boston Children's Hospital, Cedars-Sinai Medical Center, HCA Healthcare, Memorial Sloan Kettering Cancer Center, Stanford Medicine Children's Health, and University of California, San Francisco (UCSF)" (Source: openai.com), and that launch came one day after Anthropic's own Claude for Healthcare announcement, both companies racing into the same market at the same J.P. Morgan Healthcare Conference window (Source: www.fiercehealthcare.com).

For an adjacent life-sciences and AI consultancy such as IntuitionLabs advising clients on this landscape, the practical conclusion is that vendor selection should never be decided on marketing claims of "HIPAA compliant" AI; it should be decided on which specific product tier, feature set, and BAA scope actually matches the workflow in question. IntuitionLabs' own technical guidance stresses that Claude for Healthcare and Claude for Life Sciences "are not new standalone LLMs, but tailored deployments of the Claude models with domain-specific tuning, connectors, and compliance layers" (Source: intuitionlabs.ai), not a guarantee of compliance in itself.

Implementation Guidance: Getting a BAA and Building a Compliant Deployment

Organizations that determine Claude Enterprise or a HIPAA-ready API organization fits their needs should plan implementation as a multi-step process, not a single toggle. The following sequence reflects how Anthropic structures the workflow and how third-party compliance advisors describe the surrounding work required:

- **Confirm the use case actually involves PHI.** Anthropic's guidance frames this as the threshold question: "will you be processing protected health information through Claude? If yes, you need the HIPAA-ready offering with a BAA" (Source: support.claude.com). Non-PHI uses, such as medical education content or de-identified research summaries, can generally proceed on standard plans, while HIPAA Vault cautions that therapy or counseling notes, clinical documentation, patient communications, and intake summaries all belong on the list of "Not Acceptable Uses" for any non-BAA-covered surface (Source: www.hipaavault.com).
- **Select the right Anthropic product surface.** Enterprise for chat-first workflows, a HIPAA-ready Developer Platform organization for custom applications, and a decision up front about whether Claude Code needs BAA coverage (which requires the separate ZDR arrangement on a "qualified account").
- **Have the Primary Owner review and accept the BAA and Implementation Guide.** The enablement steps are: "1. Sign in to Claude as the Primary Owner and go to Organization settings > Data and privacy. 2. Go to HIPAA Compliance. 3. Click 'Enable' to open the consent flow. 4. Download the Business Associate Agreement, review it, then click 'Next.' 5. Download the Implementation Guide, review it, then click 'Next.' 6. Click 'Accept and enable HIPAA'" (Source: support.claude.com).
- **Build the compliance layer Anthropic's BAA does not provide.** This includes application-level audit logging, encrypted storage of any logs that capture PHI, key management and rotation, and access controls, none of which the BAA supplies automatically. HIPAA requires specified compliance documentation to be retained for six years, but it does not impose a universal six-year retention period on every technical audit log. Organizations should set audit-log retention through their risk analysis, applicable documentation and accounting obligations, contracts, state law, and counsel (HHS OCR audit protocol.
- **Account for multi-vendor exposure.** Organizations that route some workloads to Claude and others to OpenAI or Amazon Bedrock generally need a signed BAA with each provider individually, since compliance obligations attach separately to whichever vendor is actually processing PHI at a given moment.
- **Govern developer and internal tooling, not just production.** Engineers running Claude Code against a codebase that contains PHI in test fixtures, database dumps, or log files bring that usage into HIPAA's scope even if the customer-facing product is fully compliant, a risk Aptible calls "the developer workflow risk" (Source: www.aptible.com).

This is also where shadow AI becomes a live operational risk independent of any vendor's BAA terms. Aptible cites reporting that "17% of healthcare workers admit using unapproved AI tools at work, most commonly for documentation" (Source: www.aptible.com), a reminder that a well-negotiated Enterprise BAA is worthless if clinicians and staff continue to paste PHI into their personal Claude.ai or ChatGPT accounts out of habit or convenience. Governance, training, and technical controls that block unsanctioned consumer AI access are as important to real-world HIPAA compliance as the BAA itself.

Data Analysis and Evidence

Several quantitative threads help frame why this level of precision matters. First, Anthropic's compliance credential stack is broader than HIPAA alone: Claude via the API and Claude Enterprise both carry SOC 2 Type 2, ISO 27001, ISO 42001, CSA STAR, HIPAA, and NIST 800-171 attestations according to the Trust Center's own scope matrix, while Claude Team carries every credential except HIPAA, and Claude for Government carries FedRAMP High and DoD Impact Level authorizations instead (Source: trust.anthropic.com). Anthropic's HIPAA-specific audit artifacts include a "2025 Type 1 HIPAA - Report (C4E).pdf" for Claude Enterprise and a separate "2025 Type 1 HIPAA Report (1P API).pdf" for the API, both available to customers who request Trust Center access (Source: trust.anthropic.com).

Second, the cost of getting this wrong is not abstract. Beyond the average 10.93 million USD cost of a healthcare data breach cited above (Source: www.ibm.com), the same IBM and Ponemon Institute research found that "only one-third of breaches were detected by internal security staff" across industries (Source: www.ibm.com), and that organizations "making serious use of automation and AI cybersecurity enjoyed an average cost reduction of 1.76 million USD compared to those without such technologies" (Source: www.ibm.com), an argument healthcare organizations weigh against the risk of ungoverned AI deployment.

HIPAA's own civil penalty structure is steep and tiered. Violation fines range "from $145 to $2,190,294 per violation, depending on the level of culpability," with the top tier reserved for "willful neglect" not corrected within 30 days (Source: www.hipaajournal.com); a 2019 enforcement discretion notice clarified that the maximum annual penalty "should only apply to the most serious Tier 4 violation category" (Source: www.hipaajournal.com). Since the compliance date of the Privacy Rule in 2003, OCR "has received over 374,321 HIPAA complaints" and has resolved 99 percent of them (Source: www.hhs.gov), including 152 cases resulting in settlements or civil money penalties totaling $144,878,972 (Source: www.hhs.gov), and, separately, 67,873 cases resolved through direct technical assistance rather than formal enforcement, since "in 67,873 cases, OCR intervened early and provided technical assistance" (Source: www.hhs.gov). Enforcement has not slowed; 2025 ended with "21 settlements and civil monetary penalties, the second highest annual total to date" (Source: www.hipaajournal.com), and OCR's 2025 penalties included a $1,500,000 civil monetary penalty against Warby Parker for "multiple HIPAA Security Rule failures" (Source: www.hipaajournal.com), a $3,000,000 settlement against Solara Medical Supplies tied to an impermissible ePHI disclosure (Source: www.hipaajournal.com), and an $800,000 settlement against BayCare Health System for failures in "information access management (minimum necessary standard), risk management, information system activity review" (Source: www.hipaajournal.com).

Third, adoption of AI in the sectors most exposed to this risk is accelerating faster than governance. Physicians' use of AI "nearly doubled in a year" according to the American Medical Association, cited in OpenAI's own healthcare materials (Source: openai.com). That combination, rapidly rising clinical AI use alongside a data class (PHI) that carries the highest breach costs in any industry and a regulator that is actively increasing enforcement, is exactly why the plan-by-plan and feature-by-feature precision detailed in this report matters more for Claude than for a typical SaaS purchasing decision.

Finally, on the vendor side, IntuitionLabs' own technical review of Claude's healthcare and life sciences suite notes that Anthropic's latest models have shown material gains on medical and scientific benchmarks, citing a Protocol QA benchmark score of 0.83 for Sonnet 4.5, above both the human expert baseline of 0.79 and the prior Sonnet 4 model's 0.74 (Source: intuitionlabs.ai), reinforcing that the capability question and the compliance question are separate axes: a model can be highly capable at clinical reasoning while still being deployed through a product surface that carries no BAA coverage whatsoever.

Case Studies and Real-World Examples

Novo Nordisk: Claude Code for Regulated Pharmaceutical Documentation

Novo Nordisk, the Danish pharmaceutical company and maker of Ozempic, used Claude Code to build **NovoScribe**, an AI-powered documentation platform combining retrieval-augmented generation with domain-expert-approved templates. The results are striking: "time spent producing clinical study documentation reduced from 10+ weeks to 10 minutes" and a "95% reduction in resources needed to create device verification protocols" (Source: claude.com) (Source: claude.com). Before automation, the manual process remained prone to errors and subjective interpretation, and "each day of delay in bringing a new medicine to market can cost pharmaceutical companies up to $15 million in potential revenue" (Source: claude.com). Digitalization Strategy Director Waheed Jowiya stressed the compliance discipline required to get there: "in a highly regulated industry, we can't just throw our data and information into a large language model and hope for the best" (Source: claude.com). This case illustrates regulated-industry Claude Code usage centered on clinical trial documentation rather than PHI directly, underscoring that pharmaceutical regulatory workflows and HIPAA-covered clinical care workflows require different compliance frameworks even when the same underlying model is used.

Banner Health: Reported BannerWise Deployment at Scale

Phoenix-based Banner Health deployed a Claude-powered internal AI platform called **BannerWise**, offered to "the health system's 55,000-plus employees across six states in late 2025," supporting document analysis and summarization, content creation, and code optimization support (Source: www.beckershospitalreview.com). Banner's Chief Technology Officer Mike Reagin explained the selection criteria: "We were drawn to Anthropic's focus on AI safety and Claude's Constitutional AI approach to creating more helpful, harmless, and honest AI systems" (Source: www.fiercehealthcare.com). Anthropic's own announcement framed the broader Claude for Healthcare launch, which "rolled out the platform Jan. 11" (Source: www.beckershospitalreview.com), as capable of tools that "speed up prior authorization requests so that patients can get life-saving care more quickly" (Source: www.beckershospitalreview.com). The public reporting establishes that BannerWise is a large Claude deployment at a HIPAA-covered health system, but it does not disclose Banner's Anthropic contract, BAA execution, covered-feature configuration, or whether PHI is processed. It should therefore be treated as a deployment example, not proof of a particular HIPAA configuration.

OpenAI's Parallel Rollout to Boston Children's Hospital and Peer Institutions

As a useful point of comparison, OpenAI's ChatGPT for Healthcare rolled out in January 2026 to a similar tier of health systems, including "AdventHealth, Baylor Scott & White Health, Boston Children's Hospital, Cedars-Sinai Medical Center, HCA Healthcare, Memorial Sloan Kettering Cancer Center, Stanford Medicine Children's Health, and University of California, San Francisco (UCSF)" (Source: openai.com). Boston Children's Hospital's SVP and Chief Innovation Officer John Brownstein described the arrangement as building on prior custom work: "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" (Source: openai.com). This case demonstrates that Claude is competing directly against a nearly identical BAA-gated enterprise healthcare product from OpenAI, and that large health systems are actively running the same due-diligence process, verifying BAA scope, feature eligibility, and governance controls, against both vendors in parallel.

(Hypothetical Example) A Digital Health Startup Migrating from Claude Team to HIPAA-Ready Enterprise

Consider a hypothetical 15-person telehealth startup that initially adopted Claude Team for its developers' coding workflows and internal documentation, unaware that Team carries no HIPAA eligibility. As the product matures toward handling real patient intake data, the startup's compliance lead discovers, via Anthropic's own documentation, that "Team plans and individual plans (Free, Pro, and Max) can't enable HIPAA" (Source: support.claude.com), and that Enterprise requires a 20-seat minimum the 15-person team does not yet meet (Source: claude.com). This scenario, a composite drawn from the real friction reported in the Reddit thread and Aptible's guidance cited earlier in this report rather than a single named company, illustrates a structural gap smaller regulated-data teams frequently hit: they must either grow to Enterprise minimums, pursue a negotiated sales-assisted API HIPAA arrangement instead, or delay handling any real PHI until the organizational and contractual prerequisites are in place.

Genmab and Sanofi: Life Sciences Deployments Adjacent to HIPAA Scope

Beyond direct patient care, Anthropic counts pharmaceutical and biotech partners including Genmab and Sanofi among its Claude for Life Sciences customers, with Genmab stating that "by reducing manual burden, our partnership with Anthropic will empower our teams to focus more time on high-value scientific and strategic work, accelerating our path to patient impact" (Source: www.anthropic.com) and Sanofi noting that "Claude is integral to Sanofi's AI transformation and is used by most Sanofians daily" (Source: www.anthropic.com). Anthropic's head of biology and life sciences, Eric Kauderer-Abrams, framed this broader push in terms of institutional fit rather than raw capability: "Anthropic is a very natural fit for the healthcare and life sciences world because our identity as an AI company is built around safety and responsibility and rigor and reproducibility, and these are all the central tenants of the healthcare and life sciences industries" (Source: www.fiercehealthcare.com). These deployments, together with Fierce Healthcare's broader reporting that "Banner Health, Stanford Healthcare, Novo Nordisk, Sanofi, AbbVie and Genmab" are already using Claude's suite of LLMs for administrative workflows, clinical documentation, regulatory submissions, and clinical trial analysis (Source: www.fiercehealthcare.com), demonstrate that most life sciences R&D use cases (protocol drafting, regulatory document generation, trial data analysis) involve regulated pharmaceutical data governed by FDA and Good Clinical Practice requirements rather than HIPAA-defined PHI, a distinction healthcare and life sciences buyers frequently need clarified when scoping their own Claude deployments.

Implications and Future Directions

Anthropic's HIPAA posture has changed materially over a short period, and the trend line points toward broader, faster-activating coverage rather than a static policy. The shift to allowing BAAs signed after December 2, 2025 to cover both Enterprise and API usage under a single agreement (Source: support.claude.com), and the move to make Chat, Projects, Artifacts, Voice, Web Search, Research, and Skills all covered "under versions of the BAA accepted after 12/2/25" (Source: privacy.claude.com), both suggest Anthropic is actively widening BAA scope release by release rather than treating it as a fixed legal boilerplate. The launch of Claude for Healthcare in January 2026, with native connectors to the CMS Coverage Database, ICD-10 codes, the NPI Registry, and PubMed built directly into the HIPAA-ready Enterprise offering (Source: www.anthropic.com), signals that Anthropic sees healthcare not as an edge case to accommodate but as a core vertical worth purpose-building for, alongside its parallel Claude for Life Sciences track and its Claude for Government FedRAMP High offering for federal health agencies (Source: trust.anthropic.com).

At the same time, the persistent exclusions, Cowork, Console, Workbench, most beta features, and Claude Code outside a ZDR configuration, indicate that Anthropic is deliberately keeping its newest and most experimental surfaces outside the BAA until they mature, a pattern likely to repeat with future product launches such as Claude Design and Claude for Office as they move from beta into general availability. Buyers should expect this list of covered versus excluded features to keep shifting and should treat the published feature-eligibility tables, not organizational memory of a prior sales conversation, as the source of truth at the time of any new deployment decision.

The competitive dynamic with OpenAI, Microsoft, and Google will likely keep pushing all four vendors toward faster BAA activation and broader default coverage, since healthcare buyers are now comparing the mechanics of these agreements as a genuine differentiator rather than treating HIPAA support as table stakes assumed to be identical across vendors. For consultancies and advisors such as IntuitionLabs guiding pharmaceutical, biotech, and healthcare clients through AI vendor selection, and for Veeva's own ecosystem of life sciences customers evaluating how Claude's HIPAA-ready and Veeva-integrated tooling might complement their existing regulated content and clinical systems, the durable lesson is that "HIPAA compliant" is never a one-time label; it is a configuration that must be re-verified against the vendor's current, published terms every time a new feature, model, or workflow enters scope.

Frequently Asked Questions (FAQs)

Is Claude HIPAA compliant? Claude is not HIPAA compliant on any plan by default. It can support HIPAA compliance only on Claude Enterprise after a Primary Owner activates HIPAA readiness and accepts Anthropic's BAA, or on a Claude Developer Platform organization configured as HIPAA-ready, and only for the specific features Anthropic lists as Eligible Services (Source: privacy.claude.com).

Does Anthropic offer a Business Associate Agreement (BAA)? Yes. Anthropic "provides a BAA covering our HIPAA-ready services, such as use of our first-party API or Enterprise plans" (Source: privacy.claude.com), activated through a self-serve click-to-accept flow on Enterprise or a similar flow in the Claude Console for API organizations (Source: support.claude.com).

Is Claude Enterprise HIPAA compliant? Claude Enterprise supports HIPAA compliance only once its Primary Owner activates the HIPAA-ready configuration; "standard Claude Enterprise plans do not include BAA coverage without action from a Primary Owner" (Source: privacy.claude.com). Even then, Cowork and several beta features remain outside the BAA (Source: support.claude.com).

Is the Claude API HIPAA compliant? The Claude API can be HIPAA-ready if an organization signs Anthropic's BAA and enables a HIPAA-ready API configuration, which applies "a broader set of privacy and security safeguards than ZDR" for eligible endpoints such as the core Messages API (Source: platform.claude.com), but excludes Batch, Files, Skills, Code Execution, Computer Use, and Web Fetch entirely (Source: privacy.claude.com).

Is ChatGPT HIPAA compliant, and how does it compare to Claude? Like Claude, generic ChatGPT is not HIPAA compliant; OpenAI "won't enter into a Business Associate Agreement for the Free, Plus, Team, or Enterprise versions" without a sales-managed account or the dedicated ChatGPT for Healthcare product (Source: www.hipaajournal.com) (Source: help.openai.com). Both vendors follow the same underlying structure: consumer tiers excluded, managed enterprise tiers eligible.

Can Claude be used for healthcare data at all without a BAA? Only if the data is not PHI, for example fully de-identified research summaries, general medical education content, or an individual consumer voluntarily connecting their own health records through the personal HealthEx, Apple Health, or Android Health Connect integrations on Pro or Max, which operate under consumer terms rather than a business associate relationship (Source: www.anthropic.com).

Does the Claude Team plan support HIPAA? No. Anthropic explicitly groups Team with Free, Pro, and Max as unable to enable HIPAA (Source: support.claude.com), and the Trust Center's compliance matrix marks Team's HIPAA scope as "N/A" while Enterprise and the API show HIPAA as supported (Source: trust.anthropic.com).

Is Claude Code covered under a BAA? Only conditionally. Claude Code CLI and Desktop local mode are covered "only with zero data retention (ZDR) enabled, and only on qualified accounts" (Source: support.claude.com), and Desktop remote mode, the web version, Code Review, Security, Computer Use, and Remote Control are never covered (Source: privacy.claude.com).

What does Anthropic's HIPAA compliance credential set actually include? A HIPAA-ready configuration with BAA availability, ISO 27001:2022, ISO/IEC 42001:2023, and SOC 2 Type I and Type II reports, published through the Anthropic Trust Center (Source: privacy.claude.com) (Source: privacy.claude.com).

Conclusion

The precise, honest answer to "is Claude HIPAA compliant" is that it depends entirely on which plan, which feature, and which configuration a person means, and that answer changes if any of those three variables changes. Covered entities and business associates should not use Free, Pro, Max, or Team for workflows that disclose PHI to Anthropic because those plans cannot enable HIPAA or be covered by Anthropic's BAA. Claude Enterprise and a HIPAA-ready Claude Developer Platform organization can support HIPAA compliance, but only after a Primary Owner deliberately activates the BAA, and even then a defined set of surfaces, Cowork, Console, Workbench, most beta features, and Claude Code outside a zero-data-retention configuration, remain outside the agreement's protection. Anthropic's newest and most capable models carry their own retention requirements as Covered Models, adding yet another layer buyers must reconcile against their BAA configuration before assuming coverage.

This structure is not unique to Anthropic. OpenAI, Microsoft, and Google Cloud each draw a similarly sharp line between BAA-ineligible consumer products and BAA-eligible managed enterprise or purpose-built healthcare offerings, and each publishes its own list of excluded features that buyers must check independently. What differs is the pace of activation, the granularity of the published documentation, and, as real user accounts illustrate, how consistently the sales and support process delivers on the self-serve promise. For a healthcare provider, health plan, or health-tech company, the operative task is not to ask a single yes-or-no question about Claude, but to map the exact workflow, the exact Claude surface it touches, and the exact BAA status of that surface, then build the audit logging, access controls, and encryption that HIPAA requires around it, since Anthropic's agreement covers its own infrastructure, not the customer's implementation. Approached this way, Claude can absolutely be part of a HIPAA-compliant deployment as of July 2026; approached casually, through a personal Pro account or an unconfigured Team workspace, it cannot, and no marketing claim from any vendor changes that fact.

Tags: is claude hipaa compliant, claude hipaa baa, anthropic business associate agreement, claude enterprise hipaa, claude api hipaa, chatgpt vs claude hipaa, claude for healthcare, hipaa compliant ai, claude team plan, anthropic compliance

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.