COMPLETE COMPLIANCE REFERENCE

Compliance and HIPAA

Every compliance chapter compiled into a single document for sharing with partners, counsel, or auditors.

Table of Contents

Section 03 · Operations

HIPAA in Plain English

HIPAA stands for the **Health Insurance Portability and Accountability Act**. It is a federal law from 1996. People talk about it as if it were one thing, but it is really four rules stacked together. Knowing the four rules makes everything else about compliance easier to understand, because every technical or legal decision you make eventually traces back to one of them.

The four rules

1. The Privacy Rule

The Privacy Rule answers the question **who is allowed to see Protected Health Information, and under what conditions.** It establishes the concept of Protected Health Information, defines what counts as identifying, and sets the baseline that you cannot use or share it without a permitted reason. It also gives patients rights: to see their own records, to ask for corrections, and to know who has accessed their information.

2. The Security Rule

The Security Rule answers the question **how must Protected Health Information be protected technically and administratively.** It lists the safeguards every system that handles regulated data must have in place: encryption, access controls, audit logs, training, risk assessments, and so on. This is the rule that drives almost every technical decision in the architecture of this software. The full safeguard list is on its own page in this section.

3. The Breach Notification Rule

The Breach Notification Rule answers the question **what must you do when something goes wrong.** If Protected Health Information is exposed, you have to tell affected patients within 60 days, you have to tell the U.S. Department of Health and Human Services, and if the breach affects more than 500 people you have to tell the media. Florida adds its own breach law on top, with a faster 30-day clock for Florida residents.

4. The Enforcement Rule

The Enforcement Rule answers the question **what happens if you break the other rules.** Penalties are tiered by how aware you were and how much you tried to prevent the violation. Fines start at about $137 per violation and reach roughly $2 million per category per calendar year. Willful neglect can be charged criminally. The financial scale alone is reason enough to take the other three rules seriously.

Why all four matter for our platform

We are a software platform that ABA clinics use to run their operations and that will eventually hold Protected Health Information about their clients. That puts us inside the scope of HIPAA, which means we have obligations under all four rules. We have to know who can see what (Privacy). We have to protect it correctly (Security). We have to respond properly if something leaks (Breach Notification). And we have to do all of this consistently, because if we get sloppy we get penalized (Enforcement).

Everything else in this Compliance section is essentially the practical answer to: **how do we satisfy these four rules in code, in vendor choices, in our day-to-day operations, and in our paperwork.**

A note on what HIPAA is not

HIPAA is not a checkbox certification. There is no "HIPAA certified" stamp you can earn and then forget about. It is an ongoing legal obligation. Vendors who say they are "HIPAA compliant" really mean they are willing to sign a Business Associate Agreement and have implemented the Security Rule safeguards in their service. The compliance lives in the contracts and the practices, not in a badge.

HIPAA is also a federal floor, not a ceiling. States can add their own stricter rules on top. Florida does, and those additions are covered on the Florida Specifics page.

The Three Roles

HIPAA splits the world into three kinds of organizations. Which role you play decides which rules apply to you and who you are responsible for. Understanding this is the single most important framing for everything else in this section, because it explains why we have to sign contracts with so many different vendors.

The three roles at a glance

Role	Who fits this	What it means for them
Covered Entity	Front-line healthcare provider: hospitals, doctors, dentists, mental health practices, ABA clinics.	Original holder of the patient relationship. Carries the primary HIPAA obligation.
Business Associate	Any organization a Covered Entity hires whose work involves Protected Health Information. **This is us.**	The same Security Rule safeguards that apply to the clinic apply to the Business Associate, because the same data is involved.
Subcontractor	Any organization a Business Associate hires whose work involves Protected Health Information. AWS, Bedrock, Cognito, etc.	Inherits the same HIPAA obligations the Business Associate carries. The chain can extend indefinitely.

Covered Entity (the clinic)

A Covered Entity is the front-line healthcare provider. They are the ones directly providing care to patients and billing for that care. Hospitals, doctors, dentists, mental health practices, and the ABA clinics that use this platform are all Covered Entities. They are the original holders of the patient relationship and they carry the primary HIPAA obligation.

Business Associate (us)

A Business Associate is any organization a Covered Entity hires whose work involves touching Protected Health Information. This is us. ABA clinics hire our platform to run parts of their operations, and we end up storing and processing their clients' information. By HIPAA's definition, that makes our company a Business Associate to every clinic we serve. The same HIPAA Security Rule safeguards that apply to the clinic apply to us, because we are handling the same data.

Subcontractor (AWS and others below us)

A Subcontractor is any organization a Business Associate hires whose work involves touching Protected Health Information. AWS hosts our database. AWS Bedrock processes documents through Claude. Cognito stores our user identities. Every single one of those vendors is a Subcontractor to us, and they inherit the same HIPAA obligations we carry. A Subcontractor that hires another Subcontractor passes the obligation down again, indefinitely.

Why this chain matters

HIPAA treats this as a chain of responsibility. The Covered Entity is responsible for picking compliant Business Associates. The Business Associate is responsible for picking compliant Subcontractors. And so on. **If any single link in that chain is broken, the entire chain is broken.** An exposed record at the deepest Subcontractor is still a HIPAA violation against the Covered Entity who originally held the patient relationship.

That is why we cannot simply pick a database based on price. The database vendor must be willing to sit inside this chain. The same goes for the AI provider, the file storage, the email service, and the error monitoring tool. Every vendor either accepts the obligation in writing or we cannot use them for regulated data.

What this means in practice

Reality	How we operate against it
The clinics are also the entities we are formally answerable to under HIPAA.	If we cause a breach, the clinics are the ones the law expects to act first, and they will look to us for accountability.
Every vendor we route Protected Health Information through has to sign a Business Associate Agreement.	The contract that locks each link of the chain in place is covered on the next page.
We have to keep documentation showing the chain exists.	If HHS audits us, they will ask to see signed agreements with each Subcontractor we use. We track that on the BAA Status page.

In short: we are the middle link in a legally-defined chain. We owe the clinics above us the safeguards HIPAA requires. We require those same safeguards from every vendor below us. The contract that proves all of that is the Business Associate Agreement.

What a BAA Is and Why It Matters

A **Business Associate Agreement**, almost always shortened to BAA, is the legal contract that turns the abstract HIPAA chain of responsibility into something enforceable. It is the single most important document in the compliance picture. Without it, no amount of technical security makes a vendor "HIPAA compliant" from our perspective. With it, we have a vendor we can legally route Protected Health Information through.

What the contract actually says

Every BAA, regardless of which vendor is offering it, contains roughly the same five promises. The wording varies. The structure does not.

1. **Purpose limitation.** The vendor agrees they will only use the Protected Health Information for the specific service we hired them for. They will not train their AI models on it, sell it, use it for marketing, or repurpose it in any way we did not authorize.
2. **Safeguards.** The vendor agrees to implement the HIPAA Security Rule safeguards inside their service. Encryption at rest and in transit, access controls, audit logging, all the technical protections covered on the Security Rule Safeguards page.
3. **Breach notification.** The vendor agrees to tell us promptly if Protected Health Information is exposed on their side, so that we can fulfill our own breach notification obligations on time. This is not optional and it is how the chain of responsibility actually functions during an incident.
4. **Subcontractor pass-through.** The vendor agrees that anyone they hire who will touch the same Protected Health Information must sign an equivalent BAA with them. This is how the chain keeps extending deeper without losing protection.
5. **Return or destroy.** When the relationship ends, the vendor agrees to either return the data to us or destroy it. They cannot keep our regulated data after we stop being their customer.

Why the paper itself matters

Some founders assume that if a vendor is technically secure, they are HIPAA compliant. They are not, until there is a signed BAA in place. The reason is procedural: when HHS investigates a breach, the first thing they ask for is the signed BAA. If there is no BAA, the violation is automatic, regardless of how the data was actually protected. The paper is what proves the chain existed.

This also means that a vendor who refuses to sign a BAA, or who only offers one at a higher enterprise pricing tier, cannot be used in the regulated path of our system. They might be a fine vendor in every other respect. They are simply outside the chain.

How vendors offer BAAs

The process varies by vendor. The most common patterns:

- **Self-service through a portal.** AWS is the gold standard here. You log into AWS Artifact, click through their templated BAA, and it is signed in five minutes. No sales calls.
- **Sales-mediated contract.** Some vendors require talking to their enterprise sales team. The conversation is short and the BAA is templated, but you have to ask.
- **Tier-gated availability.** Some vendors only offer BAAs on their higher pricing tiers. The vendor's cheaper plans are not BAA-eligible at all.
- **Inherited through a parent agreement.** Some platforms route through underlying providers whose BAAs cover them. AWS Bedrock works this way: the BAA you signed with AWS automatically covers Claude inference running through Bedrock.

The BAA chain for this platform

Because we chose to build on AWS for the regulated data path, our BAA chain is short and easy to manage. We sign one BAA with AWS, and that single agreement covers our database, file storage, AI inference, email sending, logging, secrets management, and content delivery. The other vendors outside the regulated path (the public-facing site, the development tooling) do not need BAAs because they never touch Protected Health Information.

The current status of every BAA we need is tracked on the BAA Status page, which becomes the canonical "is our chain intact" reference during audits.

What Counts as Protected Health Information

Protected Health Information is the type of data HIPAA actually protects. Knowing exactly what counts and what does not is essential, because everything else in compliance hinges on this question. The same field can be regulated or unregulated depending on what else sits next to it on the page.

The definition

Protected Health Information is any information that satisfies **both** of the following:

1. It **identifies** a specific person.
2. It **relates** to that person's past, present, or future physical or mental health, the provision of healthcare to them, or the payment for that care.

The identifying part is the trigger. A diagnosis on its own, with no name attached, is not Protected Health Information. A diagnosis on a page that also carries a name, a date of birth, or any of the other identifiers below becomes Protected Health Information the moment those two things appear together. Strip every identifier and what remains is no longer regulated.

The eighteen identifiers (HIPAA Safe Harbor)

HIPAA enumerates eighteen specific identifiers in its Safe Harbor de-identification rule. If a piece of data carries any of these and is tied to health information, it is Protected Health Information. If you remove all eighteen, the remainder is considered de-identified and can be used freely.

#	Identifier	What it covers
1	Names	First, last, nickname, family member names.
2	Geographic detail smaller than a state	Street, city, county, ZIP code under specific conditions.
3	Dates tied to an individual	Birth, admission, discharge, death, and all ages over 89.
4	Phone numbers	Mobile, landline, fax.
5	Email addresses	Parent, learner, caregiver.
6	Social Security numbers	Any portion.
7	Medical record numbers	MRN, chart number, intake number.

#	Identifier	What it covers
8	Health plan beneficiary numbers	Medicaid recipient ID, commercial member ID.
9	Account numbers	Billing account, payment plan.
10	Certificate or license numbers	Driver license, professional license tied to a patient.
11	Vehicle identifiers	Plate numbers, VIN if tied to a patient.
12	Device identifiers and serials	Communication device, AAC device, sensor.
13	Web URLs	Patient portal links containing identifiers.
14	IP addresses	Logged sessions tied to a patient.
15	Biometric identifiers	Fingerprints, voiceprints, facial geometry.
16	Full-face photos and comparable images	Therapy photos, intake photos.
17	Any other unique identifier	Internal codes that, with effort, point back to a person.
18	Genetic information	Family history, genetic test results.

A working test

Before you put anything into this software, ask two questions:

1. Does this **identify a specific person**?
2. Does this **relate to their care or its payment**?

If both answers are yes, treat the data as Protected Health Information. If you are uncertain about either, treat it as Protected Health Information until proven otherwise. The cost of being too cautious is small. The cost of being wrong in the other direction is a breach.

Examples that matter for ABA

These come up often in ABA consulting and they are not always obvious.

Example	Verdict	Why
Photo of a child on a clinic's social media	Protected Health Information	Identifying (face) plus context (the clinic provides ABA care). Allowed only with a specific HIPAA-compliant release on file.
Parent email mentioning the child by name and describing a behavior	Protected Health Information	Identifying (name) plus relation to care (the behavior is the subject of therapy).

Example	Verdict	Why
Spreadsheet of learners by initials and progress score	Treat as Protected Health Information	Initials can identify when paired with other context (clinic location, dates).
De-identified case study used for training	Not Protected Health Information	Every identifier on the list has been removed and the case cannot be re-identified through context.

2.1 · OUR APPROACH

Our Stack and Why

Every piece of this section exists because we made one architectural decision and want to be able to explain it cleanly. The decision: **we are going all-in on Amazon Web Services for the regulated path of this platform.** Every vendor that ever touches Protected Health Information is either AWS itself or a service we explicitly chose because it integrates cleanly with AWS under the same Business Associate Agreement.

The decision in one sentence

We picked AWS because we get one signed Business Associate Agreement that covers our entire regulated stack, the data never leaves AWS's network as it moves between services, the cost scales reasonably, and the platform is built precisely for the kind of long-running, sensitive workloads we are designing for.

What sits inside the AWS BAA

Every service below is covered by our single AWS Business Associate Agreement. The third column is the reason that service exists in our stack rather than a generic description.

Service	What it does	Why it touches Protected Health Information
App Runner	Runs the Next.js application as a managed container.	The HTTPS request carrying an uploaded document lands here first, and App Runner reaches the database over a private VPC link, so the hosting layer is part of the regulated path. App Runner is HIPAA-eligible under the same Business Associate Agreement as the rest of the stack.
RDS Aurora Postgres	Primary database for clinic profiles, user accounts, intake submissions, AI gap analyses, and audit log entries.	Almost every non-file record ends up here. Encryption at rest is on by default and row-level security keeps clinic data separated.

Service	What it does	Why it touches Protected Health Information
S3	Object storage for uploaded files.	Intake forms, insurance cards, treatment plans, and session notes live in encrypted S3 buckets with per-clinic isolation policies.
Bedrock with Anthropic Claude	The AI engine that analyzes uploaded documents.	Document content is sent to Bedrock, which runs Claude inside AWS's own network. The data never leaves AWS to reach the model.
Cognito	Authentication and user identity.	Holds the records that prove who is logged in, what clinic they belong to, and what they are allowed to access. Multi-factor authentication is configured here.
CloudWatch	Logs and monitoring.	A scrubbing layer runs before any log call so identifiers never reach CloudWatch even when something crashes.
SES	Transactional email (password resets, intake confirmations, staff notifications).	Configured to never include Protected Health Information in the body of an email.
CloudFront	Content delivery network for the application.	Serves static assets globally. Regulated calls always route through the application layer, not the edge cache.
Secrets Manager	Holds API keys, database credentials, and other sensitive secrets.	Encrypted, access-controlled, audited. The application reads secrets at runtime under a tightly scoped role.

Why not Vercel, Neon, or other simpler stacks

The alternative we considered was a hybrid path: Vercel for hosting, Neon for the database, and AWS just for AI and file storage. That path is real and works, but it would mean signing separate Business Associate Agreements with Vercel, Neon, and AWS, managing three vendor relationships during audits, and paying multiple vendor markups.

Hybrid (Vercel + Neon + AWS)	All-AWS (what we chose)
Three Business Associate Agreements to sign and renew.	One Business Associate Agreement, one set of vendor terms.

Hybrid (Vercel + Neon + AWS)	All-AWS (what we chose)
Data crosses the public internet between vendors on every request.	Data stays inside AWS's private network end to end.
Two layers of developer-experience markup over the underlying compute.	Pay AWS once, directly, with no middleman.
Three permission systems to reconcile during audits.	One AWS IAM model controls every part of the regulated stack.
Friendlier developer experience for small Next.js apps.	Steeper learning curve, offset by AI coding assistance and the AWS CLI.

The tradeoff we accepted

AWS has a steeper learning curve than Vercel. Vercel hides the underlying infrastructure on purpose to make Next.js deployment feel easy. AWS does not. We accepted the learning curve because the long-term simplicity (one bill, one Business Associate Agreement, one set of credentials, one network) is worth the upfront effort. We work through the AWS CLI with AI coding assistance, which makes the learning curve far less steep than it would have been a few years ago.

What stays outside AWS

Not everything is on AWS. The public-facing marketing site, the development tooling, and the internal team productivity stack (Linear, GitHub, Slack, etc.) live elsewhere. None of those touch Protected Health Information, so they do not need to be inside the Business Associate Agreement chain. The line is clean: **regulated data lives in AWS, everything else lives wherever serves it best.**

How to read this page later

Six months from now, you will likely look at the AWS bill and wonder why we did not pick a cheaper or simpler option. The answer is on this page, and it is worth reading before making a switch. The architecture is not optimized for being inexpensive in absolute terms. It is optimized for being legally defensible, technically clean, and capable of scaling to dozens or hundreds of clinics without revisiting these decisions.

The Security Rule Safeguards

The HIPAA Security Rule lists the specific protections every system that handles Protected Health Information must have in place. The list is the practical answer to the question "what does it actually mean to be HIPAA secure." The rule organizes safeguards into three categories: Technical, Administrative, and Physical. This page walks each category and notes how our platform satisfies it.

Technical safeguards

Protections built into the software and infrastructure. Most are configuration decisions made once and verified during audits.

Safeguard	How we satisfy it
Encryption in transit	Every browser-to-platform request uses HTTPS. Internal calls between AWS services use AWS's private network with encrypted transport. No path moves Protected Health Information unencrypted.
Encryption at rest	RDS Aurora and S3 encrypt stored data with AES-256 using AWS-managed keys. On by default for our resources, verified through the AWS console.
Strong authentication	Every login requires a password plus a second factor (authenticator app or hardware key). No shared accounts. Service-to-service authentication uses short-lived IAM credentials, never permanent API keys.
Access controls	Every record carries the clinic ID it belongs to. Row-level security in Postgres prevents clinic A from reading clinic B's data even if the application code has a bug. The same logic applies to S3 bucket policies and Cognito group membership.
Automatic session timeout	Logged-in sessions expire after a configurable period of inactivity, forcing fresh authentication before sensitive data is exposed again.
Audit logging	Every read or write of regulated data is recorded in an append-only audit log: who acted, what record they touched, what fields, when, and from where. Tamper-evident and retained for six years per HIPAA.
Integrity controls	Database transactions are atomic. File uploads are checksummed. Audit log entries are immutable. We can prove data has not been silently altered.
Scrubbing before logging	Application logs run through a scrubbing helper that strips known Protected Health Information fields before the log line reaches CloudWatch. A crash report does not become a HIPAA incident on its own.

Administrative safeguards

Organizational policies and human practices. The technical layer protects the data; the administrative layer protects the operation.

Safeguard	How we satisfy it
Designated security officer	One person is named in writing as responsible for HIPAA security. For now that is the founder. As the team grows, the role may be split or formally delegated.
Annual risk assessment	Once a year we document the threats, the controls in place, and the residual risk we accept. The Annual Risk Assessment Log page is where these are kept.
Workforce training	Every person with access to regulated systems has completed HIPAA awareness training and signed an acknowledgment. Training is repeated annually.
Incident response plan	A written runbook covers what to do when a breach is suspected: who is notified, in what order, what gets contained, what gets logged, and when the breach notification clocks start. Held on the Incident Response Runbook page.
Sanctions policy	If a workforce member violates HIPAA rules, the consequence is documented in advance, not improvised. A deterrent and an audit requirement.
Vendor management	The BAA Status page is the canonical list of every vendor in the regulated chain, what they do, and the date their agreement was signed. Adding a new vendor requires their agreement being signed first.

Physical safeguards

Protections for the physical hardware and workstations. Most are handled by our vendors; a few are on us.

Safeguard	How we satisfy it
Data center security	AWS handles physical access to servers: locked facilities, badge controls, surveillance, environmental protection. Covered by their Business Associate Agreement.
Workstation security	Any laptop or device used to access regulated systems must have full-disk encryption, a strong login password, and automatic screen lock after a short idle period. No passwords on sticky notes. No shared workstations.
Device disposal	Old laptops or hard drives that may have held regulated data are wiped using a documented procedure before disposal, transfer, or resale.

How this gets audited

If HHS audits the platform, they will not run our code. They will ask for evidence: the signed Business Associate Agreements, the workforce training records, the most recent risk assessment, sample audit log entries, the incident response runbook, and the sanctions policy. Most of the actual technical protection is verified by simply showing them how the system is configured and pointing to the vendor agreements that lock the infrastructure in place. The administrative paperwork is what they spend the most time on, which is why we keep all of it inside this Compliance section rather than scattered across drives and email threads.

Approximate Monthly Cost

Running a HIPAA-compliant platform on AWS is not free, but it is far cheaper than the equivalent stack built across multiple vendors with separate compliance contracts. This page lays out the monthly cost we expect at two stages: the early startup phase with a handful of pilot clinics, and a scaled phase with fifty or more active clinics handling real document volume. All numbers are rough estimates in U.S. dollars, intended to make budgeting realistic rather than to predict precise bills.

Startup phase

This is where we sit today and for the next six to twelve months. Light traffic, small database, occasional AI usage, a few clinics on board for testing and early production.

Service	Monthly range	Notes
App Runner (managed container)	$15 to $35	A small always-on container (0.25 to 1 vCPU) plus serving traffic. It does not scale to zero, so there is a modest idle floor in exchange for no cold starts and simpler operations.
RDS Postgres (db.t4g.small, single availability zone)	$30 to $50	Adequate for early traffic and small data volumes.
S3 storage	$1 to $5	Scales linearly with how many gigabytes of documents we hold.
Bedrock with Claude	$20 to $100	The swing factor. Light analysis runs cost little; heavy use scales the bill quickly.
Cognito	Free	Free for the first 50,000 monthly active users.
CloudWatch Logs	$5 to $15	Application logs at startup volumes.
SES (transactional email)	~$1	Password resets, intake confirmations, notifications.
CloudFront (content delivery)	$5 to $20	At light traffic.
Secrets Manager	$2 to $5	Scales with how many secrets we store.
Data transfer out	$5 to $20	Bandwidth out of AWS to clinic browsers.
Realistic startup total	**$80 to $250**	Biggest variable is Bedrock usage. Everything else is fairly predictable.

Scale phase

Fifty or more active clinics, regular AI usage on real document volume, hundreds of gigabytes of stored files, and full-time operational traffic.

Service	Monthly range	Notes
App Runner (managed container)	$60 to $180	Auto-scales container instances with traffic. Higher serving load plus more concurrent instances during busy hours.
RDS Aurora Serverless v2 Postgres	$150 to $400	Auto-scaling: cost rises with real traffic instead of provisioning for peak.
S3 storage (terabytes of documents)	$50 to $200	Lifecycle policies can move old files to cheaper tiers automatically.
Bedrock with Claude (heavy usage)	$500 to $2,000	The largest single line item. Roughly proportional to documents analyzed.
CloudFront (high traffic)	$50 to $200	Global content delivery at scale.
CloudWatch and monitoring	$30 to $80	Full observability across services.
Everything else combined	$50 to $150	Cognito, SES, Secrets Manager, data transfer, backups.
Realistic scale total	**$900 to $3,200**	Bedrock dominates the bill. The rest grows gradually.

How this compares to the alternatives

For context, the same platform built on Vercel Enterprise plus Neon Business plus Bedrock starts at roughly $1,500 to $3,000 per month before any usage, simply for the right to sign Business Associate Agreements with those vendors. The all-AWS path is meaningfully cheaper at every stage, but the gap is largest at scale.

Stack	Startup phase	Scale phase
All-AWS (our path)	$80 to $250	$900 to $3,200
Vercel Enterprise + Neon Business + Bedrock	$1,500 to $3,000	$3,000 to $6,500

What controls the bill

A few levers swing the cost most. These are the ones to revisit any time the bill jumps.

Lever	How it moves the bill
AI usage volume	Every document run through Bedrock costs money proportional to tokens sent and received. Analyzing every uploaded file is expensive; analyzing only on request is cheap. A product decision more than a technical one.
Database size and traffic	Dominated by hours of active serving and total stored data. Aurora Serverless v2 scales down during quiet hours, which helps significantly.
File volume	S3 is cheap per gigabyte, but adds up at hundreds of clinics over years. Lifecycle policies move cold files to cheaper tiers.
Backup retention	Longer retention means more storage. We default to the HIPAA-required six years and do not extend without a reason.

Billing safeguards

AWS will let a bill grow indefinitely without warning unless you configure alerts. We set the following on day one. These are required steps on the AWS Setup Checklist page.

Safeguard	How it works
Billing alerts	Email and phone alerts trigger at $50, $100, $250, and $500 per month so a runaway bill cannot grow silently.
Anomaly detection	AWS Cost Anomaly Detection flags unusual spikes, such as a runaway loop hitting Bedrock.
Per-service spending caps	Configured wherever AWS exposes them, so a single service cannot blow past its budget.

AI Provider Standards

The AI engine in this platform is the highest-exposure single component. When a clinic uploads a document for analysis, the full content of that document is sent to the model. If that path is not covered by a Business Associate Agreement, we have a HIPAA violation regardless of how well-secured every other part of the system is. This page walks the landscape of AI providers and explains the choice we made.

How AI providers handle Business Associate Agreements

Every AI provider falls into one of three categories. The first column is the category, the second is the rule, and the third is which providers sit there.

Category	Rule	Who sits here
Business Associate Agreement available standard	HIPAA-eligible by default. Route Protected Health Information once the agreement is signed.	AWS Bedrock, Azure OpenAI Service, Google Vertex AI
Business Associate Agreement available at higher tier	Offered only on enterprise pricing or through a sales-mediated contract. Standard developer accounts are not eligible.	Anthropic direct API, OpenAI direct API
No Business Associate Agreement available	Off-limits for any regulated workload, regardless of how the tool is used.	Consumer AI tools, unofficial integrations

The provider landscape side by side

The same five providers, compared on the things that actually matter for our decision.

Provider	Business Associate Agreement	Underlying model	Fits our stack	Why we did or did not choose it
AWS Bedrock	Yes, standard	Anthropic Claude (and others) hosted inside AWS	Yes	Our chosen path. Same agreement that covers the rest of our infrastructure. No additional vendor.

Provider	Business Associate Agreement	Underlying model	Fits our stack	Why we did or did not choose it
Anthropic direct API	Enterprise tier only	Claude	Adds a vendor	Bedrock gives us the same Claude models without a second vendor relationship.
Azure OpenAI Service	Yes, standard	OpenAI GPT family	No	Mature for Microsoft-shop healthcare. Wrong fit because we are not on Azure.
Google Vertex AI	Yes, on listed services	Gemini and partner models	No	Comparable to Azure OpenAI. Wrong fit because we are not on Google Cloud.
OpenAI direct API	Enterprise tier only	GPT family	No	Standard developer tier has no agreement, so Protected Health Information is off-limits.
Vercel AI Gateway and similar routers	Depends on underlying provider	Pass-through	Convenience layer	A router can be covered, but the underlying model still has to be covered. Adding a router is convenience, not compliance.

Settings that matter once an agreement is in place

A signed Business Associate Agreement is necessary but not sufficient. The AI integration also has to be configured correctly. Every AI call that may touch Protected Health Information has to clear the five settings below.

Setting	What it requires
Zero data retention	The provider does not store our inputs after the request completes and does not use them for training. This is part of the agreement terms with Bedrock and similar services.
Logging scope	No prompt logging at the provider, or logging that sits inside the same agreement scope as the requests themselves.
Region pinning	Requests stay in U.S. regions that are explicitly covered by the agreement. Failover to non-covered regions is disabled.

Setting	What it requires
Minimum necessary	Send only the data required to answer the question. This is a HIPAA principle, not a technical convenience.
Audit trail	Every AI request that touches Protected Health Information is logged in our own audit table: who triggered it, what record it pertained to, when, and what came back.

The current state of our AI path

Today,the AI extraction inside the platform routes through OpenAI gpt-4o-mini via Vercel AI Gateway. That path is not covered by a Business Associate Agreement at our current tier and must not see Protected Health Information. Until the AWS Bedrock integration is in place, the AI features are restricted to business operations content (clinic profiles, policy text, internal team notes) and explicitly prohibited from any data that identifies a learner.

The migration to Bedrock is on the Action Items page as the prerequisite for moving Protected Health Information through the AI layer.

Allowed in This Software

This page is the day-to-day reference for what you can and cannot put into this software right now. The platform is moving through staged compliance work, and the line between allowed and prohibited will shift as each Action Item is completed. Until those are done, keep to the rules below to avoid creating a HIPAA incident on a platform that is not yet fully ready to hold Protected Health Information.

Allowed today

- Clinic legal name, DBA, address, phone, email, website.
- Clinic owner, BCBA, and RBT names in their professional capacity.
- Generic policy and procedure templates, handbook content, ICP copy, internal task lists, deliverable drafts.
- De-identified examples (all eighteen HIPAA identifiers stripped) used for training material or case studies.
- Our business's own operational data: invoices, contracts, consulting hours, internal notes.

Not allowed yet

- Any real learner first name, last name, or initials used in context.
- Session notes, progress notes, behavior intervention plans, individualized education plans, treatment plans.
- Photos of children, photos of intake forms, photos of insurance cards.
- Parent or caregiver contact information paired with the child's clinical situation.
- Medicaid IDs, commercial insurance member IDs, medical record numbers, chart numbers.
- Anything pasted into the AI extractor that contains any of the above. The current extractor routes through a provider that is not BAA-covered at our tier.

Current technical state

Where each part of the system stands today, as a quick honest snapshot:

- **Storage.** In-memory only. No encryption-at-rest control by us, no audit log, no access control. Anything saved here evaporates on server restart.

- **AI extraction.** OpenAI gpt-4o-mini via Vercel AI Gateway. No BAA at the current tier. Business operations content only.
- **Authentication.** Not yet implemented. Anyone with the URL can see everything.
- **Audit logging.** Not yet implemented. We cannot reconstruct who viewed or changed what.
- **Backups.** Not yet implemented. State is volatile.

Why these limits exist

The list of restrictions exists because the platform is mid-buildout. Each item on the "not allowed yet" list is gated on a specific Action Item being completed. Once the database is on RDS, authentication is wired through Cognito, audit logging is on, and the AI path moves to Bedrock, the corresponding restrictions can be lifted. Until then, treating the platform as if it were Protected Health Information ready would create real legal exposure.

AWS Setup Checklist

This is the step-by-step checklist to set up our AWS account for the regulated path of the platform. It is written so that someone who has never used AWS before can follow it end to end, and so that the business partner can read it cold and understand both what we are doing and why. Download the PDF version to share with the partner.

Before you start

Three things have to be in hand before opening AWS. If any of these are not yet in place, complete them first. The other todo-list items in the platform cover setting up the business email, phone, and address.

Prerequisite	Why it matters
Business email address	Use the company domain, not a personal Gmail. The AWS account is business infrastructure and should be tied to a role-style address (e.g. aws@yourdomain.com) that survives any individual leaving.
Business credit card	AWS runs a $1 verification charge that is refunded. The business card keeps AWS expenses separate from personal spending and gives the business its own credit history with AWS.
Phone number reachable during signup	AWS will text or call a verification code during account creation.

Step 1. Create the AWS account

1. Go to aws.amazon.com and click "Create an AWS Account."
2. Enter the business email and choose a strong password. Save the password in a password manager.
3. When prompted, choose the **Business** account type, not Personal. This matters for billing structure and for the Business Associate Agreement later.
4. Enter business name, address, and phone number.
5. Enter the business credit card. AWS will run a $1 verification charge and refund it.
6. Complete phone or text verification.
7. Choose the **Basic Support plan** (free) for now. We can upgrade later if we need direct AWS support engineers.

Step 2. Secure the root account immediately

The email and password you just created own the entire AWS account and all its resources. Compromise of that login is the worst-case scenario. Lock it down before doing anything else.

1. Log into the AWS console with the new root account.
2. Click the account name in the top right corner and choose **Security credentials**.
3. Enable **multi-factor authentication** on the root user. Use an authenticator app (1Password, Authy, Google Authenticator) or a hardware key.
4. Confirm that no access keys exist for the root user. If any do, delete them.
5. Sign out of the root account. From this point on, the root account is only used for tasks that explicitly require it (billing changes, account closure). Day-to-day work happens through a separate admin user.

Step 3. Set up billing alerts before anything can spike

AWS will let a bill grow indefinitely without warning unless you configure alerts. Do this on day one.

1. In the AWS console, open **Billing and Cost Management**.
2. Under **Budgets**, create a Cost budget.
3. Set the four monthly thresholds below. Each threshold sends an email when crossed.
4. Enable **Cost Anomaly Detection** for unusual spending patterns.
5. Subscribe the business email and a phone number to these alerts.

Threshold	What it means in practice
$50	Early warning. At our startup phase the entire bill should sit at or below this number, so crossing it is a signal to investigate.
$100	Approaching the upper end of the expected startup phase. Confirm the spend is from real usage, not a misconfiguration.
$250	Outside the expected startup band. Stop and audit before doing anything else.
$500	Hard alert. Treat as a possible runaway loop or unauthorized access until proven otherwise.

Step 4. Sign the AWS Business Associate Agreement

This is the legal document that puts AWS inside our HIPAA chain. Without it, we cannot route Protected Health Information through any AWS service, regardless of how secure the services themselves are.

1. In the AWS console, open **AWS Artifact**.
2. Go to the **Agreements** tab.
3. Find the **AWS Business Associate Addendum**.
4. Review the agreement (it is templated; we are accepting AWS's standard terms, not negotiating new ones).
5. Accept the agreement. The signed BAA is stored inside AWS Artifact and available for download anytime.
6. Record the signing date on our BAA Status page.

Step 5. Create an admin user (replace the root account for daily work)

1. In the AWS console, open **IAM** (Identity and Access Management).
2. Create a new user with the name "founder-admin" (or similar).
3. Attach the **AdministratorAccess** managed policy.
4. Enable multi-factor authentication on this user too.
5. Create an access key for programmatic (CLI) access. Save the access key ID and secret key in the password manager.
6. From this point on, log in through this admin user for daily work. Reserve root for billing and account-level tasks.

Step 6. Pick the AWS region

All HIPAA workloads need to live inside U.S. regions covered by the AWS BAA. We will use **us-east-1 (N. Virginia)** as the primary region. This is the largest, cheapest, and most feature-complete region. Backups can live in a secondary region (us-east-2 or us-west-2) for disaster recovery.

Step 7. Install the AWS CLI on the development machine

1. Install the AWS CLI v2 following AWS's official instructions for the operating system.
2. Run **aws configure** and paste in the access key ID, secret key, region (us-east-1), and output format (json).
3. Test with **aws sts get-caller-identity**. It should return the admin user's identity.

What is next once the account is live

Once steps 1 through 7 are complete, the AWS account is ready for actual infrastructure work. We provision the items below in order. Each one has its own configuration steps that we will work through together once the account exists.

Order	Resource	What we configure
1	RDS Aurora Postgres database	Encryption at rest, automated backups, private subnet, parameter group tuned for our workload.
2	S3 buckets for clinic file uploads	Per-clinic bucket policies, server-side encryption, versioning, lifecycle policies for cold files.
3	Bedrock model access	Request access to Anthropic Claude models in the AWS console. Pin to U.S. regions covered by the agreement.
4	Cognito user pool	Authentication, multi-factor configuration, group membership tied to clinic identity.
5	IAM roles for the application	Short-lived credentials scoped to exactly the resources the Next.js app needs to reach.
6	CloudWatch log groups	Retention configured to six years per HIPAA. Scrubbing helper installed on the application side.
7	Application wiring	Next.js application code reads from RDS, writes to S3, calls Bedrock, authenticates through Cognito, and logs through the scrubbing helper.

Sharing this with the business partner

Download this page as a PDF using the toolbar at the top. The PDF is intended to give the partner a clean, complete view of the AWS setup we are committing to. It explains both the steps and the reasoning behind each step, so the partner does not have to take any of this on faith.

Florida Specifics

Florida adds its own requirements on top of federal HIPAA. The state is a HIPAA state plus its own stack of laws and regulations, and an ABA consulting practice operating in Florida has to satisfy both layers. This page covers what stacks on top of HIPAA for Florida specifically.

HIPAA versus FIPA at a glance

Florida's breach notification law is faster than the federal one. Any breach response plan tuned for Florida residents has to be tuned to the shorter clock.

Question	HIPAA (federal)	FIPA (Florida)
Time to notify affected individuals	Within 60 days of discovery	Within 30 days of discovery
Threshold to notify a regulator	500 or more residents	500 or more Florida residents
Regulator notified	HHS Office for Civil Rights	Florida Department of Legal Affairs
Which clock applies when both do	**FIPA wins**	**Plan to 30 days**

The Florida regulatory stack

Source	What it requires
Florida Information Protection Act (FIPA)	Breach notification within 30 days of discovery. Notification of the Florida Department of Legal Affairs when 500 or more Florida residents are affected.
Agency for Health Care Administration (AHCA)	Oversees Florida health care providers, including ABA clinics on Florida Medicaid. Audits look for documented policies, dated versions, training records, and operation under the current Medicaid Provider Handbook for Behavior Analysis Services.
Florida Medicaid Provider Manual for Behavior Analysis Services	Documentation requirements stricter than commercial payors in places. Session notes, supervision documentation, parent training, and authorization paperwork have specific format and content requirements. The manual changes over time, so clinics need a documented process to know which version is in effect on a given date.

Source	What it requires
Florida Behavior Analysis Certification Board rules	Practitioner-level rules covering conduct, supervision ratios, scope of practice, and discipline. Apply at the individual clinician level. Referenced in the handbook chapters that describe BCBA and RBT responsibilities.
Florida Statutes Chapter 456	General Florida health professions law. Covers patient access to records, retention periods, and confidentiality requirements that apply alongside HIPAA. Retention requirements can extend longer than HIPAA's six years for certain record types.

What this means for the platform

Requirement	How the platform supports it
Florida-cited policy content	Every policy chapter cites Florida sources, not generic national ones. A clinic using the binder does not have to reinterpret a California policy for Florida.
Audit-ready evidence	The clinic can produce, on demand, the policy in effect on a given date, the staff who acknowledged it, and the evidence that training happened.
30-day breach response	For Protected Health Information workloads, the Incident Response Runbook is tuned to the FIPA 30-day clock rather than HIPAA's 60.
Records retention	Defaults to the longer of the applicable Florida requirement and HIPAA's six years.

For clinics in other states

As the platform expands beyond Florida, each new jurisdiction will need its own version of this page. The federal HIPAA layer stays the same; the state stack changes. We will add state-specific pages as we onboard the first clinic in each new state, rather than speculating about every state up front.

Action Items

This is the gap between the current "business operations product" and the eventual "Protected Health Information ready product." Each item is a real blocker, not a nice-to-have. Doing them in roughly this order keeps each step independently useful so the platform never sits in a half-complete state.

In order

1. Set up the AWS account and sign the BAA

The first step in the chain. The AWS Setup Checklist page walks this end to end. Without it, none of the downstream items can begin. Status: AWS account created and the AWS Business Associate Addendum signed on June 3, 2026. Account hardening (admin user, CLI access) in progress.

2. Replace the in-memory client store with RDS Postgres

Stand up the regulated database with encryption at rest, automated backups, and row-level security. Migrate the existing application data away from in-memory storage. Status: not started, blocked on Step 1.

3. Wire S3 for file uploads

Configure encrypted S3 buckets with bucket policies that enforce per-clinic isolation. Build the application upload path so files go directly to S3 via signed URLs. Status: not started, blocked on Step 1.

4. Move AI extraction to AWS Bedrock

Replace the current OpenAI gpt-4o-mini call routed through Vercel AI Gateway with an Anthropic Claude call routed through Bedrock. This is the change that unlocks Protected Health Information flowing through the AI layer. Status: not started, blocked on Step 1.

5. Add authentication and role-based access control

Stand up Cognito with multi-factor authentication and define the roles: Owner, Consultant, Clinic Admin, Clinic Staff. Bind every regulated route to a role check. Status: not started, blocked on Step 1.

6. Add audit logging for every read and write of regulated data

Append-only audit log table recording: who acted, what record, what fields, when, from where. Tamper-evident and retained for six years per HIPAA. Status: not started, blocked on Step 2.

7. Complete a HIPAA Security Risk Analysis

Required by the Security Rule. Document the threats, the controls in place, the residual risk we accept, and the actions taken to mitigate. Stored on the Annual Risk Assessment Log page. Status: not started, can begin once Steps 1 through 6 are complete enough to document accurately.

8. Publish internal Privacy and Security policies

Internal to our consulting practice as a Business Associate. Distinct from the policies we publish for clinic clients. Status: not started.

9. Set up the breach response plan tuned to FIPA

30-day notification clock for Florida residents under the Florida Information Protection Act. Codify on the Incident Response Runbook page. Status: not started.

10. Multi-tenant data isolation review

Each clinic's data must be cleanly partitioned at the database, file storage, and authentication layers. Run a deliberate review of every regulated table and bucket to confirm. Status: not started, blocked on Steps 2 through 6.

Already complete

- **Added a "do not paste Protected Health Information" guardrail to AI inputs.** Visible on the Smart Paste panel inside New Client. This kept the platform safe to use during the transition while the AI path was still routed through a non-BAA-covered provider.

How to use this list

Treat this as the canonical roadmap for getting the platform to Protected Health Information readiness. Each completed step is a real unlock that allows the platform to take on workloads it could not before. Until all ten are complete, the Allowed in This Software page sets the boundary on what can flow through the platform.