Clinitect.ai — Security & HIPAA Overview

Clinitect.ai · SECURITY-AND-HIPAA.md

Audience: clinicians, clinic administrators, compliance officers, and hospital IT teams evaluating whether Clinitect.ai is safe to use with protected health information (PHI).

Purpose of this document: answer every question a reasonable buyer will ask about how we protect patient data, in plain language that both doctors and IT teams can understand and verify.

TL;DR — what makes Clinitect safe

1. **The AI model never sees real patient data.** Names, dates of birth, MRNs, and every other identifier are replaced with placeholders like `[NAME_001]` *before* any text is sent to the language model. The model only sees the placeholders.
2. **Patient data never leaves AWS.** Every service we use (Bedrock, Comprehend Medical, Textract, DynamoDB, S3, Cognito) is covered by Amazon's signed Business Associate Agreement (BAA). No third-party AI vendor, no OpenAI, no Anthropic direct, no Google — nothing PHI touches exits AWS.
3. **We don't keep a copy.** The mapping that connects `[NAME_001]` back to "John Smith" lives in computer memory for the life of a single request (~1 second) and is then destroyed. Nothing is written to our database or logs.
4. **Encrypted everywhere, end-to-end.** HTTPS in transit. AES-256 or KMS-managed encryption at rest on every data store. No exceptions.
5. **Customer controls you can actually verify.** AWS CloudTrail records every API call. Every row in every database is tagged with the user ID that created it; access is enforced server-side on every read and write.

The data flow in one picture

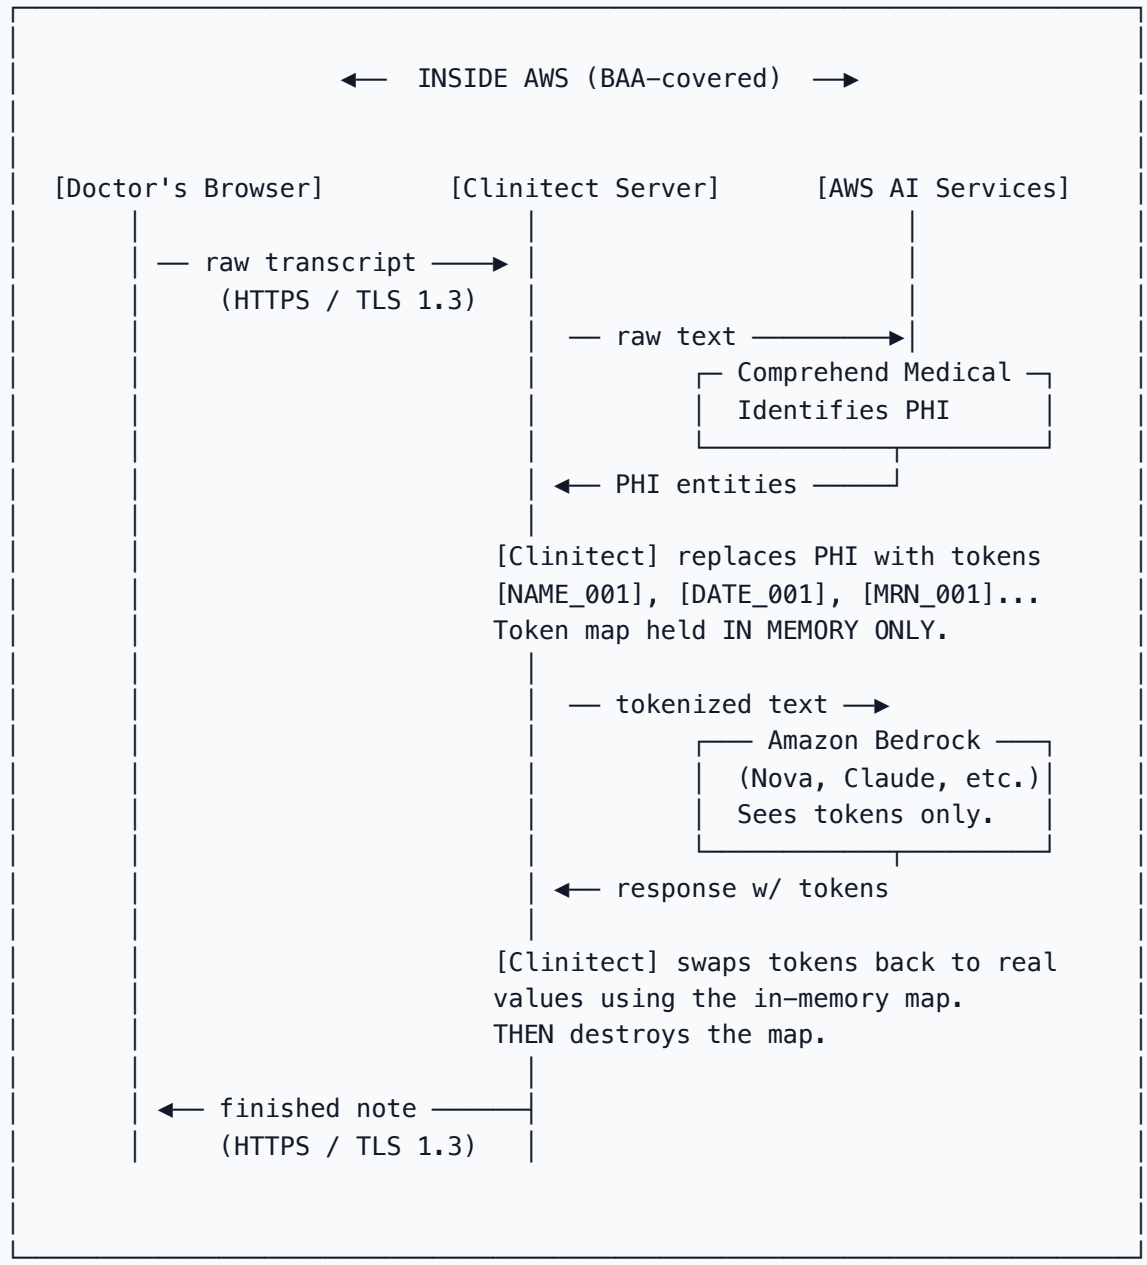

The single most important thing to see in this diagram: the AI model column receives tokenized text. It never receives real patient data.

Who sees real patient data — a plain-English breakdown

This is the table we recommend showing to a clinical skeptic:

System or party	Sees the patient's name / DOB / MRN?	How long?	Can they use it?
Doctor's browser	Yes (they typed it)	As long as the tab is open	Normal clinical use
Clinitect server	Yes — **briefly, in memory only**	~1 second per request	Processes, tokenizes, destroys
AWS Comprehend Medical	Yes (identifies PHI)	Up to 24 hrs for AWS's internal debugging	Cannot train models on it (contractual)
AWS Textract *(only if an image / PDF is uploaded)*	Yes (extracts text)	Up to 24 hrs for AWS's debugging	Cannot train models on it (contractual)
The AI model itself (Nova / Claude / Llama via Bedrock)	❌ **Never**	n/a	Sees `[NAME_001]` only
Model providers (Amazon, Anthropic, Meta)	❌ **Never**	n/a	Bedrock contractually prohibits
OpenAI, ChatGPT, or any third party	❌ **Never**	n/a	Not in the architecture
Clinitect's database	❌ **Never**	n/a	Nothing PHI-bearing is persisted here
Clinitect's logs	❌ **Never**	n/a	Logs record *counts* ("12 PHI entities detected") not values
The public internet	❌ **Never**	n/a	All traffic is TLS-encrypted; internal AWS traffic stays on AWS's backbone

HIPAA compliance — the bullet list a compliance officer wants

Business Associate Agreements

- **AWS Business Associate Agreement (BAA)** — automatic coverage for every AWS service we use. AWS's BAA is a standard published agreement; no negotiation required.

- **Services covered by the BAA we rely on:** Amazon Bedrock, Amazon Comprehend Medical, Amazon Textract, Amazon S3, Amazon DynamoDB, Amazon Cognito, Amazon CloudWatch, Amazon KMS, AWS IAM.
- **When Clinitect signs a customer agreement, we provide our own BAA downstream** so the customer is legally covered in the chain.

Encryption

- **In transit:** TLS 1.3 via Let's Encrypt certificates. HTTP is not available — all port-80 traffic 301-redirects to HTTPS automatically.
- **At rest:** AES-256 server-side encryption on every data store. DynamoDB tables use AWS KMS-managed keys. S3 bucket uses AES-256 SSE.
- **In memory:** the token map lives only in process memory for the lifetime of a single API request; it is garbage-collected and zeroed immediately after.

Access control

- **Customer authentication via AWS Cognito** (email + password with 8-character minimum, OR Google sign-in). Session tokens expire in 60 minutes; refresh tokens are revocable.
- **Every database row is scoped to a specific user ID.** Reads and writes are enforced with server-side conditional checks — a user presenting a valid token for account A cannot read or modify account B's data, even by constructing a request by hand.
- **Administrative access** to our AWS account is restricted to named individuals with IAM credentials; all access is logged to AWS CloudTrail.

Audit & accountability

- **Every API call is logged.** AWS CloudTrail captures every administrative action on our AWS resources. Application-level logs record request metadata (timestamps, user ID, action) — **never the PHI content itself**.
- **Logs are tamper-evident.** CloudTrail logs are stored in S3 with versioning and object-lock options available for enterprise customers.

Breach handling

- **Breach notification procedures** follow the HIPAA Breach Notification Rule (45 CFR §§ 164.400–414). In the event of a confirmed incident, notification is provided within 60 days to affected individuals and covered entities.
- **Incident response plan** includes: detection via AWS GuardDuty (planned), containment via IAM role revocation, forensics via CloudTrail log review, and notification via contractual channels.

Data retention & deletion

- **Token maps:** held in memory only for a single request, then destroyed.
- **Conversation state** (for multi-turn refinement): stored in DynamoDB with a hard **1-hour time-to-live** from last activity. After that the row is auto-deleted by AWS — no human action, no background job.

- **Saved notes** (for customers who opt to create an account and save to a project): stored in DynamoDB with KMS encryption, retained for the life of the account, and deleted on customer request within 30 days.
- **Right to deletion:** customers can delete any individual note or entire projects at any time via the UI. Full account deletion is available by written request.

Training on customer data

- **We do not train any model on customer data. Ever.**
- **Our AI provider (AWS Bedrock) is contractually prohibited** from using customer prompts or completions to train any Amazon or partner models. This is written into AWS's service terms.
- **Model providers (Amazon for Nova, Meta for Llama, Anthropic for Claude when enabled) contractually cannot see customer data.** Bedrock is the intermediary; the providers receive no prompts and no responses.

What we store vs. what we do NOT store

We DO store:

- Your account information (email, name, and the fact that you have an account) — required for you to sign in and see your own saved work.
- Notes you explicitly save to a project — encrypted at rest, scoped to your account only, deletable at any time.
- Usage counters for rate-limiting (by IP address for anonymous demos, by user ID for signed-in users) — no PHI content, just counts.
- Request audit logs — timestamp, user ID, action, response code. No PHI.

We do NOT store:

- **Raw patient transcripts** — they are processed and returned; no copy kept server-side.
- **The mapping between `[NAME_001]` and "John Smith"** — memory-only, destroyed after each request.
- **The de-identified transcript after the response is returned** — discarded with the rest of request state.
- **Any log line containing PHI** — our logging explicitly redacts entity values and logs only counts.
- **Anything received from AWS Comprehend Medical or Textract** beyond the immediate request — we use the response, return it to the user, and let it fall out of memory.

Frequently Asked Questions — the answers to memorize

Q: Is this HIPAA-compliant?

Yes. Every service in our stack is covered by AWS's Business Associate Agreement. We sign downstream BAAs with customers. Our architecture ensures that no unauthorized third party (including AI model providers) ever sees PHI.

Q: Who is your AI provider? Have they signed a BAA?

Our AI provider is **Amazon Web Services (AWS)** via the Bedrock service. AWS has signed a Business Associate Agreement that covers Bedrock. Within Bedrock, the specific model used (Amazon Nova, Claude, Llama) receives tokenized text only — no actual PHI — so the question of the model provider's BAA doesn't arise in the first place.

Q: Does OpenAI or ChatGPT touch my data?

No. Clinitect does not use OpenAI, ChatGPT, or any non-AWS LLM service.

Q: What happens if I paste a note with a name into Clinitect? Where does that name go?

The name travels over HTTPS to our server. Our server runs AWS Comprehend Medical to detect it, replaces it with a placeholder like `[NAME_001]` before sending anything to the language model, receives the model's response, swaps the placeholder back with the real name, and returns it to your browser. The memory structure that holds that mapping is destroyed within a second of the response being sent. The name is never written to our database or logs.

Q: What about AWS retaining data for 24 hours for debugging? Doesn't that break the "no storage" claim?

This is a nuance worth explaining clearly: AWS's published policy allows internal retention of Comprehend Medical and Textract inputs for up to 24 hours solely for service-quality debugging. That data is covered by the same Business Associate Agreement as the rest of AWS, is contractually prohibited from being used for model training, and can be further restricted with the AWS AI Services Opt-Out Policy at the organization level. To be fully precise: "patient data never leaves AWS" is true; "patient data is not stored anywhere" is almost true — it may briefly exist in AWS-internal debug storage for up to 24 hours, fully covered by the BAA.

Q: What happens if an employee of yours tries to peek at a patient's note?

They can't unless they first acquire an administrative AWS IAM credential (which a very small number of named engineers have). Every access to our AWS resources is logged to AWS CloudTrail, which is tamper-evident. The database is encrypted at rest with AWS-managed keys — reading it requires both AWS credentials and the KMS key permission, both of which are audit-logged.

Q: What if AWS gets breached?

AWS operates under SOC 2, HIPAA, HITRUST, ISO 27001, and FedRAMP certifications. A breach of AWS's core infrastructure at the scale that would compromise a specific DynamoDB table has not

occurred in the service's history and would be a global-news event. In the unlikely scenario, AWS's BAA defines the notification chain; Clinitect would in turn notify affected customers within 60 days per the HIPAA Breach Notification Rule.

Q: Can I see an audit log of everyone who accessed my data?

Not yet via self-service in the current UI. On request, we can export the relevant AWS CloudTrail events for your account. Audit-log self-service via the UI is on the enterprise feature roadmap.

Q: Is this different from just pasting into ChatGPT?

Yes, in every way that matters. - ChatGPT's standard consumer product **does not sign BAAs** and **may retain conversations for 30 days** (or longer if the user has memory enabled). Clinitect stores no copy of the transcript and is fully BAA-covered. - ChatGPT receives the patient's name, DOB, MRN, and every other identifier in the clear. Clinitect's AI model receives only tokens. - If you paste a patient note into ChatGPT and the user's account is later compromised, the patient's chart is in the attacker's search history. With Clinitect, there is nothing for an attacker to exfiltrate — we do not keep the notes.

Q: What happens if I delete my account?

All project data (saved notes, conversation history, usage counters) is deleted from our database. Aggregate counters (e.g., "how many notes did this user save in total?") are retained in de-identified form for business metrics. Audit logs are retained for the time period required for compliance.

Q: Can I self-host Clinitect on my own AWS account?

Yes — enterprise customers can deploy Clinitect inside their own AWS organization. In that configuration, even the brief 24-hour AWS Comprehend debug window happens in the customer's own account under their own BAA, and Clinitect (the company) has no access to any patient data. Contact us for enterprise deployment pricing.

Architecture summary for the technical reviewer

For a compliance officer or IT reviewer who wants the one-page architectural summary:

- **Application stack:** Python Flask backend + Next.js frontend deployed on a single EC2 instance in a private VPC segment. Nginx terminates TLS; all traffic below nginx is local-loopback only.
- **Identity:** AWS Cognito user pool with mandatory password policy or federated Google OAuth. All API endpoints that expose user data verify a Cognito-issued JWT on every request.
- **PHI detection:** AWS Comprehend Medical `DetectPHI` API, supplemented by regex matchers for structured identifiers (SSN, MRN, NPI, DOB, insurance ID). False-positive filter removes over-detection of common clinical terminology.
- **Tokenization:** Deterministic per-request token map — same text within a single session always produces the same token. Token format: `[TYPE_NNN]` (e.g., `[NAME_001]`, `[DATE_003]`).
- **LLM call:** AWS Bedrock `Converse` API, model-agnostic (currently Amazon Nova family + Meta Llama 3.3; Anthropic Claude enabled where account subscription permits). Only tokenized content is

transmitted.

- **Re-identification:** Token substitution on the response, with fallback for hallucinated tokens (replaced with `(not specified in transcript)`). Markdown stripping for plain-text output.
- **Data stores:**
- `clinitect-conversations` (DynamoDB, KMS-encrypted, 1-hour TTL, never contains raw PHI)
- `clinitect-projects` and `clinitect-notes` (DynamoDB, KMS-encrypted, user-scoped)
- `clinitect-rate-limits` (DynamoDB, IP- or user-keyed counts, 24-hour TTL)
- `clinitect-metrics` (DynamoDB, aggregate counters only, no PHI or user-identifying content)
- `clinitect-ai-uploads` (S3, AES-256 SSE, public-access-blocked, currently unused in the application path)
- **Permissions:** Single IAM role on the EC2 instance with least-privilege scoped to specific DynamoDB tables and specific Bedrock/Comprehend Medical/Textract actions.
- **Secrets:** Loaded from environment variables populated by AWS Secrets Manager in production. No static AWS access keys on disk.
- **Transport:** TLS 1.3 with HTTP Strict Transport Security.
- **Monitoring:** CloudWatch Logs for application events; CloudTrail for AWS API calls. GuardDuty integration planned.

What we are still improving

In the interest of full transparency — here's the honest punch-list of what we are working on, not what we are hiding. A reviewer who reads only the marketing version of a vendor's security story will be more suspicious, not less.

- **Formal SOC 2 Type II audit:** planned for the next 12 months. Not yet completed.
- **Third-party penetration test:** planned for the next 6 months. Not yet conducted.
- **AWS WAF and rate limiting at the edge:** planned, not yet deployed (we rate-limit at the application layer today).
- **Customer-facing audit log self-service:** planned, not yet available.
- **Enterprise single-tenant deployment mode:** available on request; automated provisioning not yet built.

None of these gaps compromise the core HIPAA controls described above. They represent the maturity path of a young healthcare startup that intends to earn hospital-grade trust over time, not a set of show-stoppers for a pilot with a single clinic today.

How to use this document in a sales conversation

A suggested order of operations when a prospective customer's compliance person asks about security:

1. **Lead with the diagram.** Show them the one-picture data flow. Compliance officers appreciate visual confirmation.
2. **Hand them "Who sees real patient data — a plain-English breakdown"** — the table. It answers the most common question (who sees PHI) in the most concrete way.
3. **If they ask about AWS:** point to the BAA section and mention Bedrock/Comprehend Medical/Textract/DynamoDB are all covered by the AWS BAA.
4. **If they ask about storage:** walk through "What we store vs. what we do NOT store."
5. **If they ask about the difference from ChatGPT:** the FAQ section has the answer written out.
6. **If they ask what's on the roadmap:** share the "What we are still improving" section. Honesty builds more trust than fake completeness.

If a customer's compliance officer wants a live walkthrough of the architecture, an AWS IAM policy review, or sample CloudTrail events, Clinitect's engineering team can provide that on a call.

Document version: 1.0 — prepared April 2026. For questions: contact hello@clinitect.ai.
