Clinitect.ai — The One-Page Security Promise

Clinitect.ai · SECURITY-ONE-PAGER.md

For doctors and clinic admins evaluating whether it's safe to use.

The core promise

Your patient's name, DOB, MRN, and every other identifier are replaced with placeholders before the AI ever sees them. The AI model receives text like `Patient [NAME_001] ([ID_001]) reports headache x3 days.` — never the real patient data. The placeholders are swapped back to real values on our server *after* the AI responds, and the mapping that makes that possible is destroyed within seconds.

This is the mechanic. Everything else below is just the detail behind it.

What you need to know in 30 seconds

✅ HIPAA compliant?	**Yes.** Every piece of our infrastructure (AWS Bedrock, AWS Comprehend Medical, AWS Textract, DynamoDB, S3, Cognito) is covered by Amazon's Business Associate Agreement.
✅ Patient data leaves AWS?	**Never.** Not to OpenAI, not to ChatGPT, not to any vendor. It stays entirely inside Amazon's BAA-covered perimeter.
✅ Patient data stored in our database?	**No.** We do not keep a copy of the transcripts or the notes you generate.
✅ The AI model sees real patient data?	**No.** The model sees placeholder tokens only.
✅ Used to train any AI model?	**Never.** Contractually prohibited by AWS.
✅ Encrypted in transit and at rest?	**Yes.** TLS 1.3 on the wire, KMS-managed AES-256 on disk.
✅ Does a BAA get signed?	**Yes.** Clinitect provides a BAA to every customer.

The visual — what happens when you paste a note

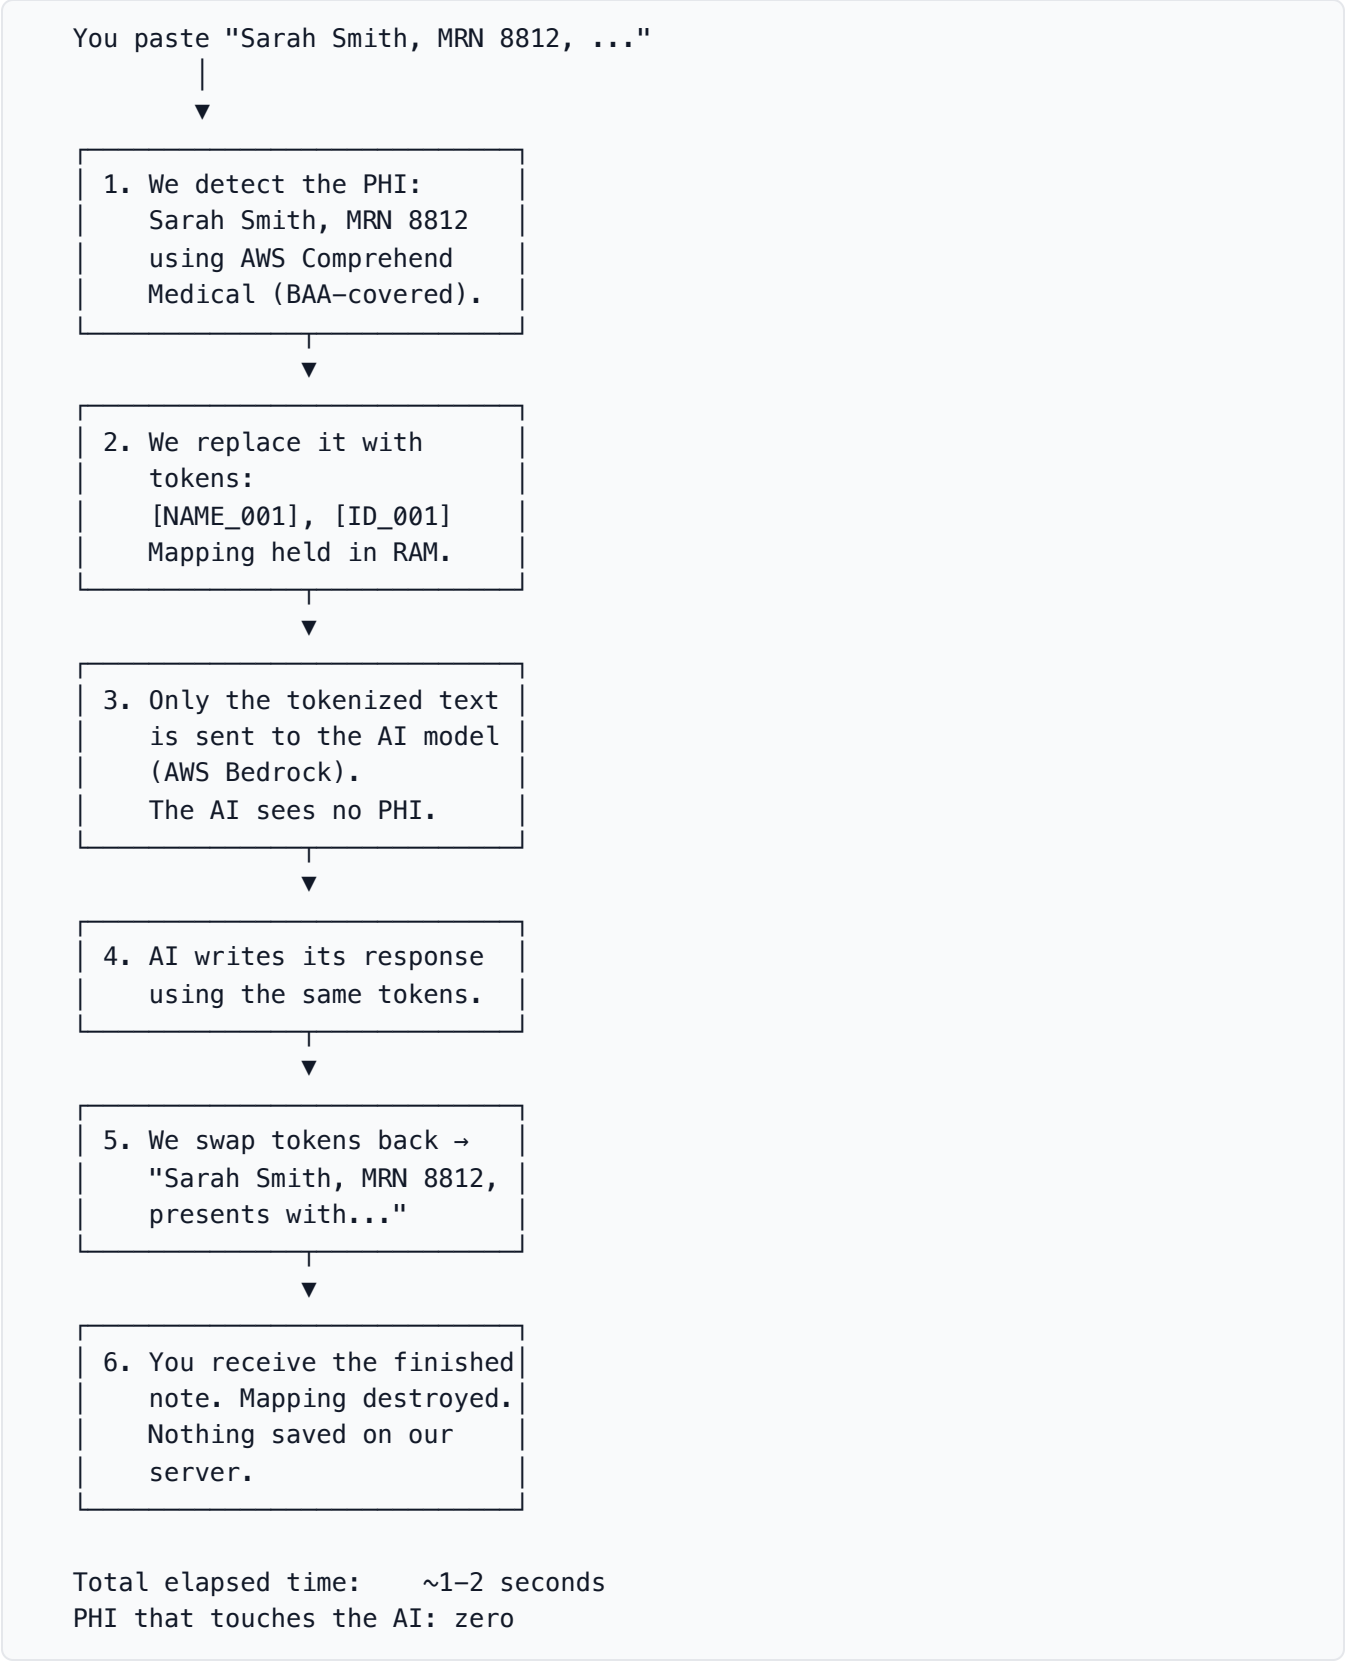

Why this is different from pasting into ChatGPT

	Pasting into ChatGPT	Clinitect.ai
AI provider has signed a BAA with you?	❌ Not on consumer tier	✅ AWS BAA covers everything
AI sees patient name, DOB, MRN?	✅ Yes	❌ Never (only tokens)
Transcript stored on their servers?	✅ 30 days default, forever if you saved the conversation	❌ No copy kept
Used to train their AI?	Depends on your tier	Never (contractual AWS prohibition)
Audit trail you can produce for a regulator?	❌ No	✅ Yes
Shared infrastructure with non-healthcare users?	✅ Yes	❌ Dedicated per-tenant AWS account (enterprise)

The three sentences a compliance officer will want

> *"Clinitect runs entirely on AWS and is covered by AWS's Business Associate Agreement. Protected health information is identified and replaced with placeholder tokens before reaching the language model, and the mapping between tokens and real PHI is held in memory only for the duration of a single API request. We do not retain transcripts, do not persist the token map, do not share data with third-party AI vendors, and do not use customer data for model training."*

If a buyer wants deeper detail

Point them to the full document: **Clinitect Security & HIPAA Overview** — it has the complete architecture, the full FAQ, and an honest roadmap of what we're still working on (SOC 2 audit, penetration test) so nobody feels like we're hiding anything.

For questions: contact hello@clinitect.ai.