

How can agencies use AI to assemble health records for review?

Software gathers a person's records, puts them in order, and links each entry to its source for a reviewer, with masking, access control, and logging built in.

By Linda Apsley, Managing Partner & CEO · Published September 23, 2026 · [strataedge.ai](#)

Agencies can use AI to gather a person's treatment records from separate systems, match them to the right person, put them in order, and link every entry to its source document so a reviewer can check it. The reviewer reads the assembled record and makes the decision. Because the records contain protected health information, masking, access approval, and logging have to be designed in from the start. We built this for the Defense Health Agency (DHA), as described below.

What we built for the Defense Health Agency

For DHA, under contract HT003826PE001, we built a multi-cloud data ecosystem and a proof of concept on it. Pipelines bring health records from separate systems into governed data products with full lineage. AI agents analyze the records and answer questions across documents, including questions about what is missing. Service members review, correct, and certify their assembled treatment record in a portal. It runs at a scale of one million records, and every value traces back to its source. The proof of concept is complete.

Why assembly takes so long by hand

A person's treatment history can sit in several places: older and newer electronic health record systems, records from outside providers, and scanned documents. A reviewer assembling it by hand requests records from each source, waits for them, checks that each document belongs to the right person, removes duplicates, and puts events in order. Only then can the review start. Missing records are easy to overlook, because nothing in the file shows that they are missing.

What the software does and what people decide

Record assembly software gathers records from each approved source, matches them to the person, removes duplicates, and orders them by date. AI reads unstructured material, such as clinical notes and scanned documents, and pulls out dates, providers, and events. Every entry in the assembled record links to the document it came from, so the reviewer can open the source and check it.

The software should also report what it could not do: sources that did not respond, documents it could not confidently match to the person, and gaps in the timeline. Those go to a person. Decisions about the person's care, claim, or case stay with the reviewer.

Security controls for health records

Health records need these controls, and the agency should be able to see evidence that each one works:

- **Masking.** Analysts, testers, and engineers usually do not need names or identifiers to do their work. Mask them by default, and record who can unmask them and who approved each exception.
- **Access.** Give each role only the records it needs, approve access through the agency's process, and remove access when the work ends.
- **Logging.** Record who viewed or exported each record, every unmasking, and every request that sent record data to a model. Keep the log where security staff and auditors can read it.
- **Model hosting.** Run models inside the agency's approved boundary, and list every service that receives record data. A cloud provider's authorization covers its own services. The application built on them needs its own security review and authority to operate.

Starting with a small proof of concept

Start with one type of review and a fixed set of sources. Agree on what reviewers need to see and how you will judge the assembled records. Then test the software against records that reviewers have already assembled by hand, and have them review the differences. At StrataEdge, a proof of concept usually takes four to eight weeks.

See [our healthcare work](#), [government capabilities](#), and [typical engagements](#).

[Talk with Linda about your agency's records](#)