

DATA INFRASTRUCTURE FOR AI

Every kind of data, turned into one secure memory that AI can use and act on

Miriel connects to the places information already lives and turns it into one memory that people and AI agents can ask questions of. It takes in data of any kind: documents, spreadsheets and databases; email, chat and calendars; images, audio and video; and live streams such as video calls, cameras and sensors, learned in real time as they happen. Answers cite their sources. Beyond answering, Miriel acts on what it knows, from sending a message to running a workflow or building a complete app. Data can be encrypted throughout, including with fully homomorphic encryption.

How it works

ANY SOURCE

- Documents, scans, drawings
- Email, chat, calendars
- Databases and business apps
- Cloud storage and code
- Meetings and calls LIVE
- Cameras, lidar, sensors LIVE



Miriel

Understand
Text, tables, images, audio and video are read, transcribed and described, including live streams as they arrive.

Remember
One store holds searchable content, the relationships between people, things and facts, and a timeline of events, all queried together.

Protect
Encryption, including FHE, per-user permissions, and deployment in your own environment if needed.



PUT TO WORK

- Answer**
Plain-language questions, with sources
- Act**
Messages, alerts, drafts and updates
- Automate**
Workflows that run on a schedule or on new data
- Build**
Context-aware apps, generated and deployed

120+ connectors for the systems organizations already use, such as Google Drive, Slack, Salesforce, GitHub, RingCentral, ServiceNow, Snowflake, Datadog, UniFi Protect cameras and MQTT sensors. They pull data in, and they do real work in those systems too: drafting and sending, posting, scheduling and updating records.

Three ways to use it

AS A PRODUCT

The Miriel workspace

A ready-to-use workspace for teams: chat sessions, a live calendar and inbox, the organization's memory and its apps. A one-box assistant, by keyboard or voice, serves anyone who wants no learning curve.

THROUGH MCP

Inside your AI tools

Miriel hosts MCP servers, so Claude, ChatGPT, GitHub Copilot, Cursor and custom agents can learn into and query the same memory as a tool.

```
# create and deploy your server
POST /api/v2/mcp/default
```

THROUGH THE API

Inside your software

One API and one credential to learn, query and build apps, so Miriel can sit behind products you already run.

```
# put data in, get answers out
POST /api/v2/learn
POST /api/v2/query/
```

What organizations use it for

Miriel serves very different jobs, because each one comes down to the same thing: many kinds of data, joined up, and put to work.

MANUFACTURING

SOP compliance and floor visibility

Cameras, lidar and equipment data are joined with the procedures that govern the work, so the floor can be asked questions.

"How many stations were welding at 2 pm, and were the pre-weld checks done?"

Sources: cameras, lidar, machine and metrology data, SOPs, drawings

SECURITY

Physical security and site awareness

Camera feeds are watched continuously. Objects and activity are found by describing them in words, with no model training, and the evidence frames come back with the answer.

"Who was in the loading bay after 10 pm, and what did they carry out?"

Sources: security cameras, access logs, incident reports

COLLABORATION

Meetings that feed everything else

Webex, Teams, Zoom and Meet calls are understood live and joined with email, documents and past meetings, then turned into follow-up work.

"What did we promise this customer last quarter, and has it shipped?"

Sources: calls, calendars, email, CRM, contracts

REGULATED KNOWLEDGE

Answers from hard documents

Drawings, scans and complex tables are read properly, and questions that need every document, such as counts and totals, are answered across all of them.

"Which parts in this package need heat treatment?"

Sources: specs, drawings, certificates, inspection records

OPERATIONS

Business processes that run themselves

Workflows watch incoming data, reconcile it against records, and bring a person in only when something does not match.

"Tell me when an invoice doesn't match our payments sheet."

Sources: email, spreadsheets, databases, finance systems

SOFTWARE

Apps built from your data, with AutoDev

Describe an app and Miriel's autonomous development system writes it, wires it to the right data, and deploys a preview to review.


"Build a dashboard of open customer commitments, updated from every call."

Sources: anything already in Miriel

From answers to action

Answering questions is the start. Because Miriel knows the context, it can act on it with the right permissions: respond with sources, draft and post to email, Slack or Webex, run workflows on a schedule or on new data, and build complete context-aware apps.

EXAMPLE: AN APP FROM ONE SENTENCE



Why Miriel

AI is only as useful as the data it can reach, and that data is large and never stops arriving. A server's logs can reach 1 GB a day. A single camera can produce 50 GB. The two usual shortcuts break down at that scale.

1 GB

a day from one server's logs

50 GB

a day from one camera

1.5 TB

a day from a 30-camera site

Give an LLM tools to search it

The model has to read through the data to find anything, so every question burns tokens scanning it, and gets slower and costlier as the data grows.

Train a custom model on it

The model is out of date as soon as new data arrives. Every new source, product or procedure means retraining and redeploying.

Miriel

Understands and indexes data as it arrives, queryable in under a second. The LLM sees only the few facts that answer the question, and new data never needs retraining.

Token-efficient

Many questions need no LLM at all. Detection, counting, matching and time-window queries run directly on what Miriel has already recorded, at zero token cost. When an LLM is needed, Miriel sends it only a compact slice of context.

EXAMPLE: TAILGATING, WITH NO LLM

COUNT(people through door) > COUNT(badge swipes) in the same time window → alert, with the video clip attached

~89x

compression: ~40M tokens of source distilled into compact facts

~20x

less storage for video, by recording only changes

15–65x

fewer tokens sent to the model per question

0 tokens

for counting, detection and time-window questions

Scalable

In Miriel's cloud, capacity grows with your data automatically. Underneath, ingestion scales out across machines with no coordination between them, so it grows in a straight line.

7,981/s

vectors ingested on one 4-GPU node, on busy video

18,175/s

on three nodes, scaling linearly

Auto

scaling in Miriel's cloud, with no hardware to plan

<1 s

from capture, such as a camera frame, to queryable

IN PRACTICE

One node keeps up with a busy site at once: for example, 30 cameras at 24 fps, plus audio feeds, door and badge access systems, other sensors and the site's documents. Video only uses capacity while something is moving, so quiet periods cost almost nothing.

Miriel · miriel.ai · info@miriel.ai

3 / 4

Secure, open and yours

Encryption and privacy

- **Encrypted storage**, with controls set per organization.
- **Fully homomorphic encryption (FHE)**: search can run over data that stays encrypted, so even the servers doing the work never see it in readable form.
- **Heightened privacy mode** keeps content out of logs and routes AI calls to a private provider that does not retain data.
- **Per-user permissions** apply to everything Miriel stores, searches and returns, so answers never cross access boundaries.

Model-agnostic

Miriel works with the leading AI model providers and with models you host yourself. You can switch models within a single conversation, even between providers, and the context carries over. The memory belongs to you, not to any model: the same memory serves every provider, and can be shared across people, teams and AI tools.

OpenAI

Anthropic

Google

AWS Bedrock

Self-hosted models

With self-hosted models, the whole system can run locally, with no outside AI service.

Runs where your data must live

Miriel cloud

The fastest start. Fully managed, with heightened privacy available per organization.

Your cloud account

Miriel deployed in your own cloud environment, under your controls.

Fully local

All of Miriel, AI models included, inside your building. Its own vision, speech and search models fit on a single modest GPU.

One memory, many ways in

Use Miriel as a finished product, plug it into the AI tools your teams already use through MCP, or build it into your own software through the API. The same secure memory sits behind all three.

Talk to us: info@miriel.ai · miriel.ai