

TRam Enterprise Technologies

TRam Studio

Put AI on the work you can't afford to get wrong.

Turn the documents and procedures you already have into agents that do the work — and keep a record of every step, so you can show anyone exactly what happened and why.

- Member, Microsoft AI Cloud Partner Program · MPN 7142971
- Built on Microsoft Azure · .NET 10
- Two granted US patents · 20+ years in regulated enterprise technology

WHAT IT IS

Your procedures, running as agents.

Six things every agent gets, without a services engagement.

01

Knowledge

Point it at your documents or your website. It answers from your material, and says so when it cannot.

02

A procedure

Ordered steps you define — read, extract, check, notify. No prompt engineering required.

03

Permissions

It may use only the tools you granted. Anything else stops the run and is recorded.

04

A record

Every run kept: what it read, what it decided, what it refused, what it cost.

05

An endpoint

Publish once and call it from any system over HTTPS with a scoped key.

06

A schedule

Run it hourly, daily, weekdays or weekly. Nobody has to remember.

ALL SIX ARE IN EVERY PLAN

It answers from your documents.

PDF, Word, Excel, CSV, JSON, text and Markdown — or your whole website.

HOW A KNOWLEDGE BASE IS BUILT

1 Upload, or paste a URL

Nine file types. A website is crawled sitemap-first.

2 Split into sections

Divided on its own headings, so a section is a real idea.

3 Indexed and searchable

Agents retrieve what matters and cite the document.

4 Kept current

Schedule a re-crawl. Unchanged pages are skipped.

WHAT THIS PREVENTS

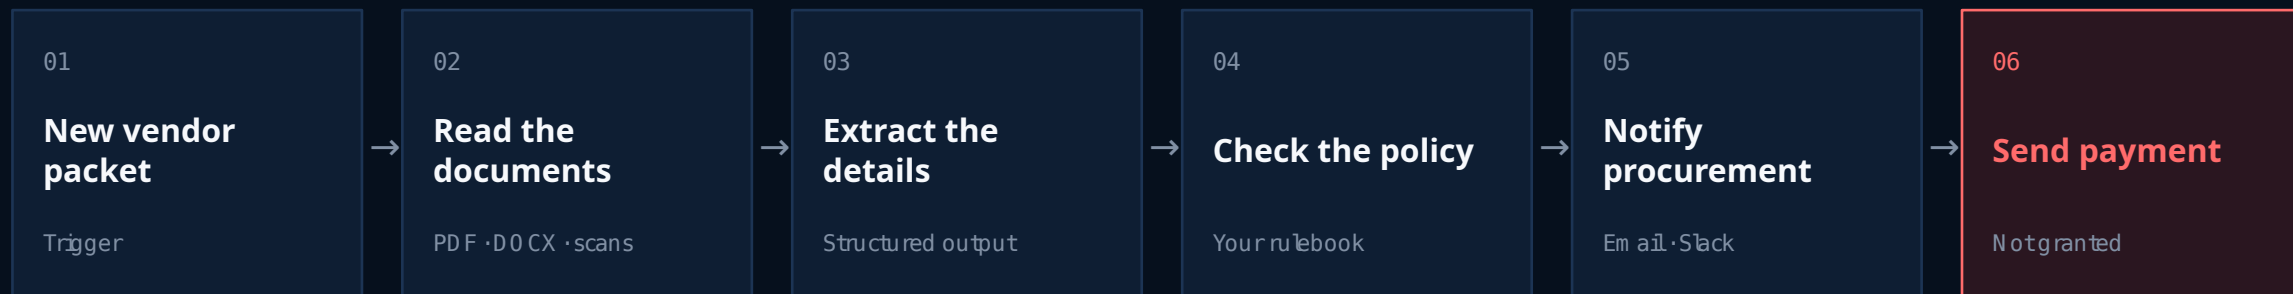
An agent that invents an answer.

When the answer is not in your material, it says so rather than filling the gap. Every answer names the document and section it came from, so a reviewer can check it.

`Grounded in: your uploaded agreement · §7.2`

The part nobody else shows you.

An agent that reaches for a tool it was not granted stops — and the stop is recorded.



WHY THIS MATTERS

Most demonstrations show the happy path. Compliance asks about the other one.

The run above stopped at step 06 because nobody granted that tool. Months later you can open that run and see exactly what it read, what it decided, and what it declined to do — and hand the record to a reviewer.

23 working agents on day one.

Clone one, point it at your documents, publish. Roughly two minutes.

Contract review

PDF in — obligations, key dates, a verdict out

Invoice extraction

Vendor, totals, line items — including scans

Compliance checklist

Present, missing or unclear, with the evidence

Document Q&A

Staff ask in plain language, get a cited answer

Resume screening

Against your criteria, with the reasoning kept

Document classifier

Sorted by type with a confidence score

Receipt to expense

Photograph in, structured expense line out

Meeting notes to actions

Raw notes in — decisions, owners, actions out

Data extraction

Any form or PDF into clean JSON

Watch it work before it ships.

Two views of the same agent — whichever one you need.

OVERVIEW

For the person answering for it

- **What this agent does** the steps, in order, including any it may not take
- **Last activity** completed or stopped, when, and how long it took
- **Typical duration** across recent runs, so you can spot the slow one
- **Recent activity** who or what started each run, and what happened
- **Reads from** which knowledge bases and systems it can reach

TECHNICAL

For the person integrating it

```
P0ST /v1/agents/42/invoke
```

```
curl -X P0ST ... /v1/agents/42/invoke \
-H "Authorization: Bearer sk_live_...." \
-d '{"inputs":{"document":"<base64>..."}}
```

Scoped keys · OpenAPI per agent · the run id and evidence link come back in the response

Your systems. Your credentials. Your model.

Registered once, referenced by name, never returned.

CONNECTIONS

Slack	post digests, alerts and answers	LIVE
Jira	read and triage tickets	LIVE
HubSpot	search CRM contacts	LIVE
SQL Server	read-only questions over your data	LIVE
Webhooks / HTTP	call any authorised service	LIVE

Secrets are encrypted at rest and write-only — a step names a connection, it never sees the credential.

MODEL CHOICE

Azure OpenAI	default
Google Gemini	selectable
Anthropic Claude	selectable
OpenAI	selectable

Choose per workspace or per agent, and bring your own key if you would rather your usage sat on your own account.

On a schedule, and supervised.

The dashboard opens on what needs attention, not on what is fine.

NEEDS ATTENTION

- **Vendor Onboarding failed** 2h ago
Could not read the supplier packet
- **Invoice Desk has been running a long time** 3h ago
Started 3h ago and has not finished
- **Policy library did not sync** 9h ago
The site could not be read last night

When nothing is wrong it says so — silence would be ambiguous.

SCHEDULES, IN PLAIN ENGLISH

Every weekday at 6:00 AM ET

next tomorrow at 6:00 AM · last run completed

Every Monday at 9:30 AM ET

next Monday at 9:30 AM · last run completed

Every day at 3:00 AM ET

re-reads the policy site · last run completed

Hourly, daily, weekdays or weekly. Overlap-safe, and the schedule advances even when a run fails.

Team-ready, six languages, spend capped.

Roles and workspaces

Owners control money and credentials. Admins run keys and connections. Members build and run. Every workspace is isolated — agents, keys, documents, activity.

Spend that cannot surprise you

Usage is metered per run and enforced before a model is called. At the cap the API returns a refusal, not a larger invoice.

Six languages

English, Spanish, French, German, Chinese and Arabic — with the entire layout flipped for right-to-left.

Your data stays yours

Customer content is never used to train foundation models and is never sold. Documents and run records live in your workspace.

HOW IT FITS TOGETHER

One platform, end to end.

WHERE WORK ARRIVES

Console · your systems over HTTPS · schedules

THE GATEWAY

Scoped keys · per-agent auth · tenant isolation · spend cap enforced before the model is called

THE RUN ENGINE

Resolves the procedure and runs its steps in order — read a document, reason, reach a system, return a result

KNOWLEDGE AND MODELS

Your indexed documents and websites · Azure OpenAI, Gemini, Claude or OpenAI

REACH

Vaulted connections — Slack, Jira, HubSpot, SQL Server, webhooks

THE RECORD

Every run, every step, cost and status — exportable as evidence

HOW YOU BUY IT

A subscription that cannot surprise you.

Every plan includes every capability. Plans differ by usage, not by features.

Trial

Free · 14 days

Everything, with a starter allowance.

- All 23 templates
- Full agent builder
- Live API endpoint
- No card required

Starter

Monthly

For one team automating a few tasks.

- Everything in Trial
- Higher usage allowance
- Scheduled runs
- Email support

Growth

Monthly

For several teams and real volume.

- Everything in Starter
- Larger allowance
- More workspaces
- Priority support

Enterprise

Talk to us

Your own terms, your own key.

- Bring your own model key
- Custom limits
- Security review
- Named contact

PRICES ON THE LISTING · NO SETUP FEE · CANCEL AT THE END OF ANY PERIOD

Bring us one tedious task.

On a single 20-minute call we will turn it into a working agent, run it against your own document, and show you the record it produced. If it is not useful you will know inside the call.

[Book a live build](#)

[Start a free trial](#)

studio.tramenterprise.com · engage@tramenterprise.com · 440.554.4850

TRam Enterprise Technologies LLC · Member, Microsoft AI Cloud Partner Program · M PN 7142971