AI AGENT PLATFORM

AgentWorks

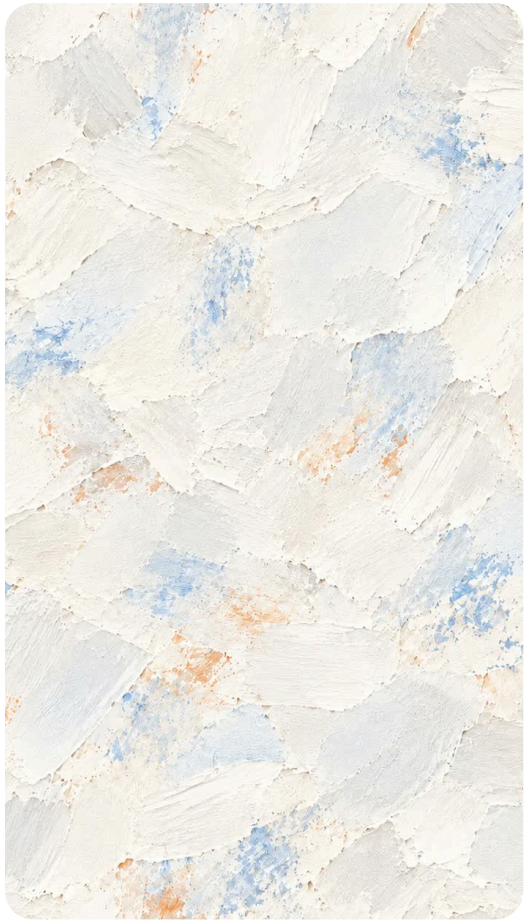
A platform to build and run AI agents

A no-code agentic AI platform that runs even on closed networks.

Build and run your own agents for your organization's work.



This brochure starts from the four concerns that security-conscious organizations and companies actually face when adopting AI. For each one, it shows how AgentWorks solves it, followed by real deployments and how to get started.



Introduction

Introduction	Four concerns about adopting AI	p.03
Platform	AgentWorks, the agent factory inside your organization	p.04
Structure	Four layers and one gateway	p.05

Four concerns and solutions

Concern 01	Network separation and policy rule out commercial AI	p.06
Concern 02	Commissioning development for every task costs too much	p.09
Concern 03	General AI doesn’t know our policies, documents or databases	p.12
Concern 04	We can’t prove who sent what to which AI	p.14

Deployments and getting started

Deployments	Different organizations, one platform	p.17
Getting started	Start small, then expand	p.19
Contact	Pilot proposal · Contact us	p.20

From security-first organizations

Four concerns that come up when adopting AI

Everyone agrees AI is needed.
The trouble comes next: policy, people, data and accountability.

Start with the concern that matters most to you.
Each one takes you to how AgentWorks solves it.



Concern 01	“Network separation and internal policy keep us from using commercial AI.”	On-site · own model →
Concern 02	“Commissioning development for every task strains budget and time.”	No-code builder →
Concern 03	“General AI doesn’t know our internal policies, documents or databases.”	Connected data →
Concern 04	“We can’t tell who sent which data to which AI.”	Two gateways →

The agent factory inside your organization

AI that doesn't stop at answering, but works inside your organization.

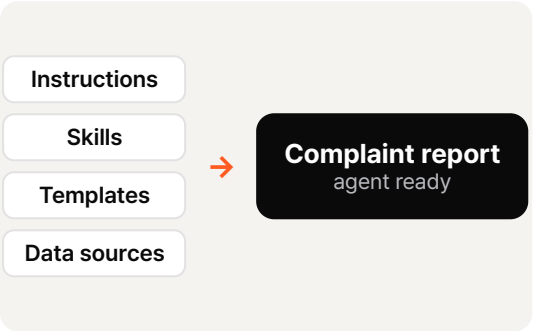
A platform to build and run AI agents

A no-code agentic AI platform that runs even on a closed network. Not a single chatbot, but a foundation for building agent after agent.

01 · Build · Run

Business teams build and run agents

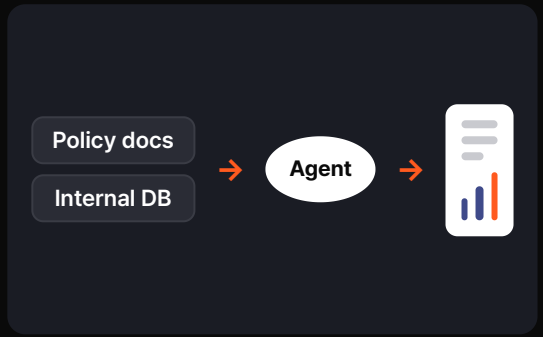
Describe the job and AI designs the task, then runs it on schedule from a queue with no one at the keyboard.



02 · Connect

Works with your own data

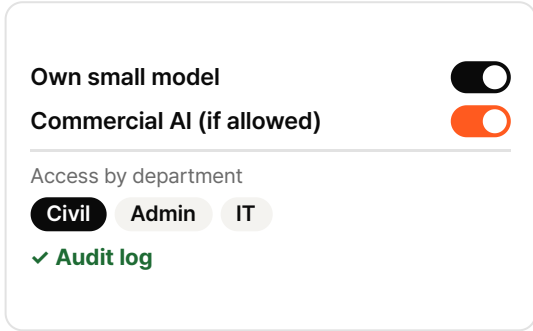
Finds evidence in internal documents and databases, and uses systems and tools to produce the deliverable.



03 · Governance · Security

Controls and records

Set models, permissions and limits, and keep usage and audit logs inside your organization.



Keep improving Usage logs and quality checks feed back into each task, and agents grow team by team.

4 ways	10	15	9	2	GS cert.
to build an agent	built-in tools	office tools	data connectors	AI call gateways	Korean public procurement

Platform structure

A platform of four layers and one gateway

User Environment
Where it is used · the entry point

Web chat staff instruct directly
Existing systems Open API · scheduled runs
New services AI inside your own app

Agent Runtime
Where it works · execution engine

Engine plan → act → check
Sub-agents split complex jobs
Schedule · queue runs unattended
Safeguards loops · time · budget · PII

Concern 04

Agent Builder
Where it is built · no-code

AI task design describe it, get the steps
Skills reusable business rules
Templates register your forms
Quality checks graded against test cases

Concern 02

Enterprise Connection
Where it connects · evidence and tools

Library knowledge from documents
Datasets ask DBs and Excel
Connectors auto-sync sources
External tools real work via MCP

Concern 03

Security · control tenant isolation · roles · sharing · delegation · audit log · login lockout

Deployment on-site (closed network) · hybrid · cloud

AI Gateway · models · keys · usage

Department gateways
Model · key management
Usage · cost records
Credit limits

Egress Gateway
One exit for external connections · connects and logs only allowed addresses · closed-network ready

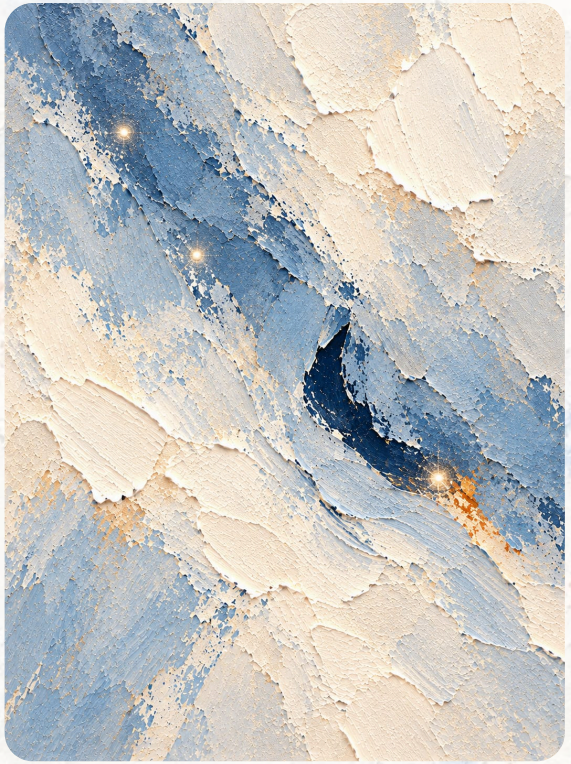
Own small model
GPU server inside your organization

Commercial AI
only if allowed, via Egress

Concern 01
Concern 04

Click a 'Concern' button to go to the concern that layer solves.

Concern 01 · Limits on external AI



Network separation and internal policy keep us from using commercial AI.

The first thing we hear from security-conscious public and financial organizations.

What actually happens

AI is needed, but the road is blocked

Network separation and security rules mean work data can't be sent to commercial AI on the internet.

Blocking it creates workarounds

Urgent work moves to personal devices and accounts, which takes it further outside the organization's control.

Building it yourself is too heavy

Bringing AI inside means preparing models, infrastructure and operations from scratch.

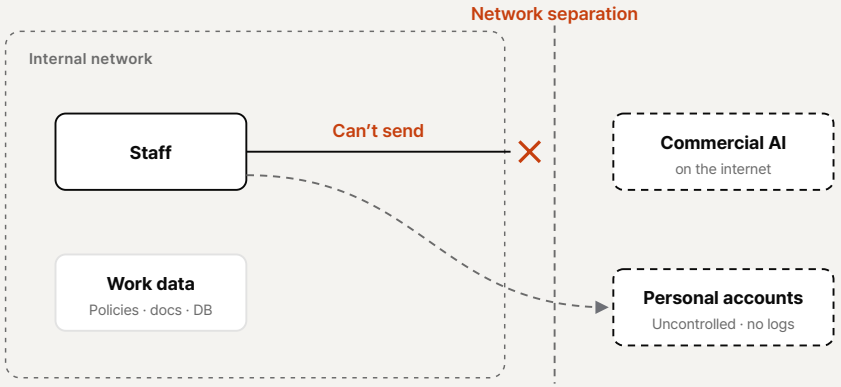
The worry underneath **"That security keeps pushing AI adoption back, again and again."**

See how AgentWorks solves it →

Solution 01 · On-site install + own AI model

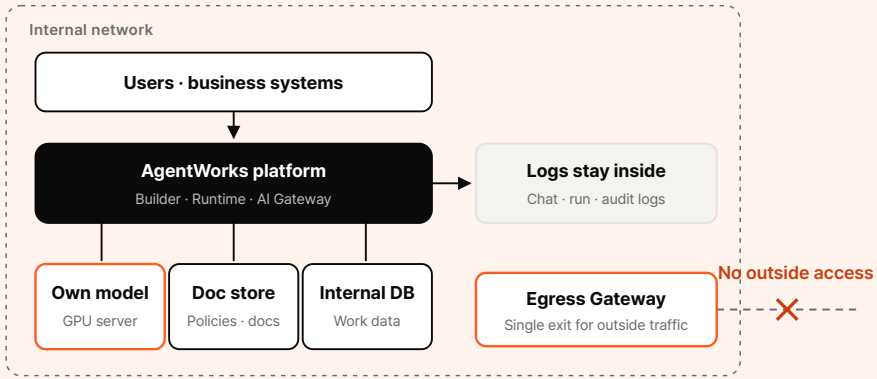
Runs without the internet, start to finish inside your organization.

Today



AI outside the internal network can't be used, and some work goes around controls.

With AgentWorks



Chat, document analysis and task runs are handled by your own model on your GPU server, and data and logs stay inside.

Your own model end to end

Chat, document analysis and task runs are handled by your own model on your GPU server. It is OpenAI-compatible, so open models connect as they are.

Data and logs stay inside

Document indexes, chats, run logs and audit logs are kept in internal storage, with data separated by organization.

Connected on the internal network

Databases, shared folders and SFTP sync through data connectors, and existing systems call agents through the Open API.

One route for outside traffic

All external connections go through the Egress Gateway. Leave the allow-list empty and it runs with no outside connection.

What you get

- ✓ Use AI for work while keeping to security rules
- ✓ On a closed network, data and logs never leave the organization
- ✓ Keep the same platform even when models change

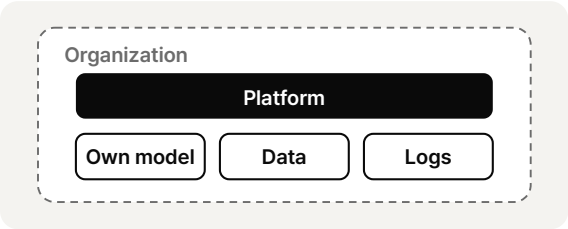
Solution 01 · Deployment for your environment

Choose how to deploy to match your security needs.

On-site (closed network)

Public · finance

- Runs on **your own model** , no internet
- All data and logs stay inside
- Empty allow-list blocks the outside



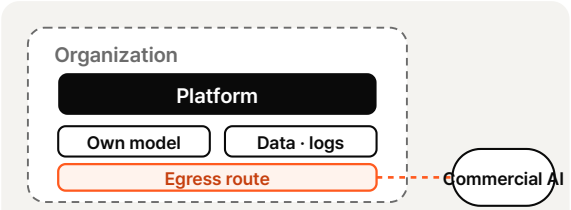
Best for
Public and financial organizations that require network separation

AI model	**Your GPU server**
Data · logs	**Inside**
Outside access	**None**

Hybrid

Split by sensitivity

- On-site + **some commercial AI**
- Sensitive work on internal models
- Commercial AI via Egress, allowed only



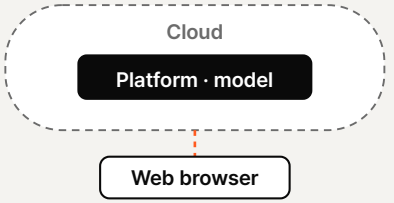
Best for
Organizations and companies with both sensitive and general work

AI model	**Own + allowed commercial AI**
Data · logs	**Inside**
Outside access	**Single Egress route**

Cloud (SaaS)

Fast pilot

- **Use right away** after sign-up
- Move on-site after the pilot
- Suits general business work



Best for
Organizations and companies that want to see results quickly

AI model	**Commercial AI**
Data · logs	**Cloud**
Outside access	**Internet**

What you get

- ✓ Pick the option that fits your security requirements
- ✓ Run sensitive and general work separately
- ✓ Pilot in the cloud, then move on-site

Concern 02 · Limited development capacity



Commissioning development for every task strains budget and time.

Something business teams and IT face together when bringing AI into their work.

What actually happens

Requests pile up, development lags

Every department has work it wants to hand over, but developers and budget are limited, so requests wait in line.

By the time it's done, the work has changed

While requirements are written up and results awaited, the way of working changes and it is commissioned again.

The people who know the work can't fix it

Those who know the work best don't code, so they have no way to fix or improve it themselves.

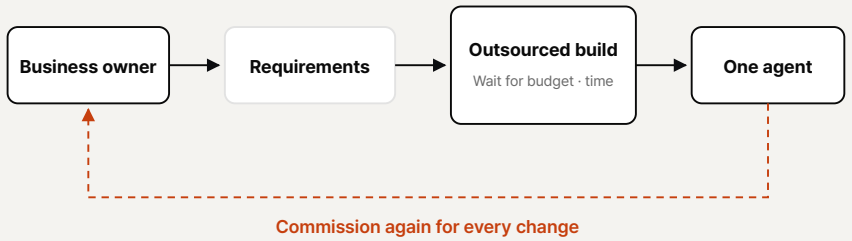
The worry underneath **"That AI adoption becomes one more development project."**

See how AgentWorks solves it →

Solution 02 · A no-code builder for business teams

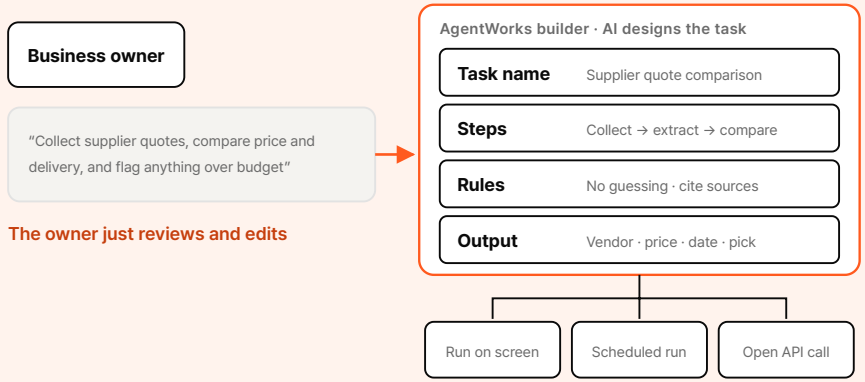
Type one sentence and it becomes an agent that works every week.

Today



Changing a single task means writing, commissioning and waiting all over again.

With AgentWorks



The owner types one sentence, AI turns it into a task, and the owner just reviews and edits.

Anyone · straight from chat

Point to a library and ask, and it becomes a guidance agent.

Business owners · describe it

Type what needs doing and AI designs the steps and sets the schedule.

Experts · package as skills

Turn writing rules and formats into skills so the whole department works to the same standard.

System owners · Open API

Call tasks from your own app or another system through the API.

What you get

- ✓ Start without a development request or outsourcing
- ✓ When the work changes, the owner fixes it right away
- ✓ The whole department works to the same standard

Solution 02 · How one is built

A complaint status report agent, built like this.

What the owner typed

“Tally last week’s complaints by type and department, list overdue and repeat complaints, and turn it into a weekly report”

Step 1	Step 2	Step 3	Step 4	Step 5
Instruct	**AI plans**	**Connect data**	**Schedule**	**Result**
Describe the result you want	query → tally → find overdue and repeats → report	Complaint dataset + handling guideline library	Runs every Monday at 09:00	Report (summary · tables · charts) · Word · PDF · email

Plan designed by AI · example

1. Query last week’s entries in the complaint dataset
2. Count by type and department
3. Find overdue items and repeat complainants
4. Check related rules in the handling guidelines
5. Write the weekly report → email the owner

Weekly complaint status Oct week 3 · auto-generated — Email sent

Received	Closed	Overdue	Avg. time	Repeats
128	112	9	2.4 days	4

By type



Needs attention

Past deadline	9
Repeat complainants	4
Possible rule breaches	0

Sources · complaint dataset · handling guidelines

You can also build: Policy Q&A that cites the clause · RFP analysis · Document summaries · Statistics in plain language · Automatic regular reports · Incident intake · root cause · First-line tech support · Minutes from audio · PII redaction

The plan and report shown are examples for illustration.

Concern 03 · Hard to use your own data



General AI doesn’t know our internal policies, documents or databases.

The first limit people notice once they try AI for real work.

What actually happens

It can’t answer from your documents

Policies and guidelines pile up as HWP and PDF files, but general AI can’t answer based on them.

No one knows where the answer came from

You can’t see which policy or clause an answer is based on, so it can’t be used as is.

Data is locked inside systems

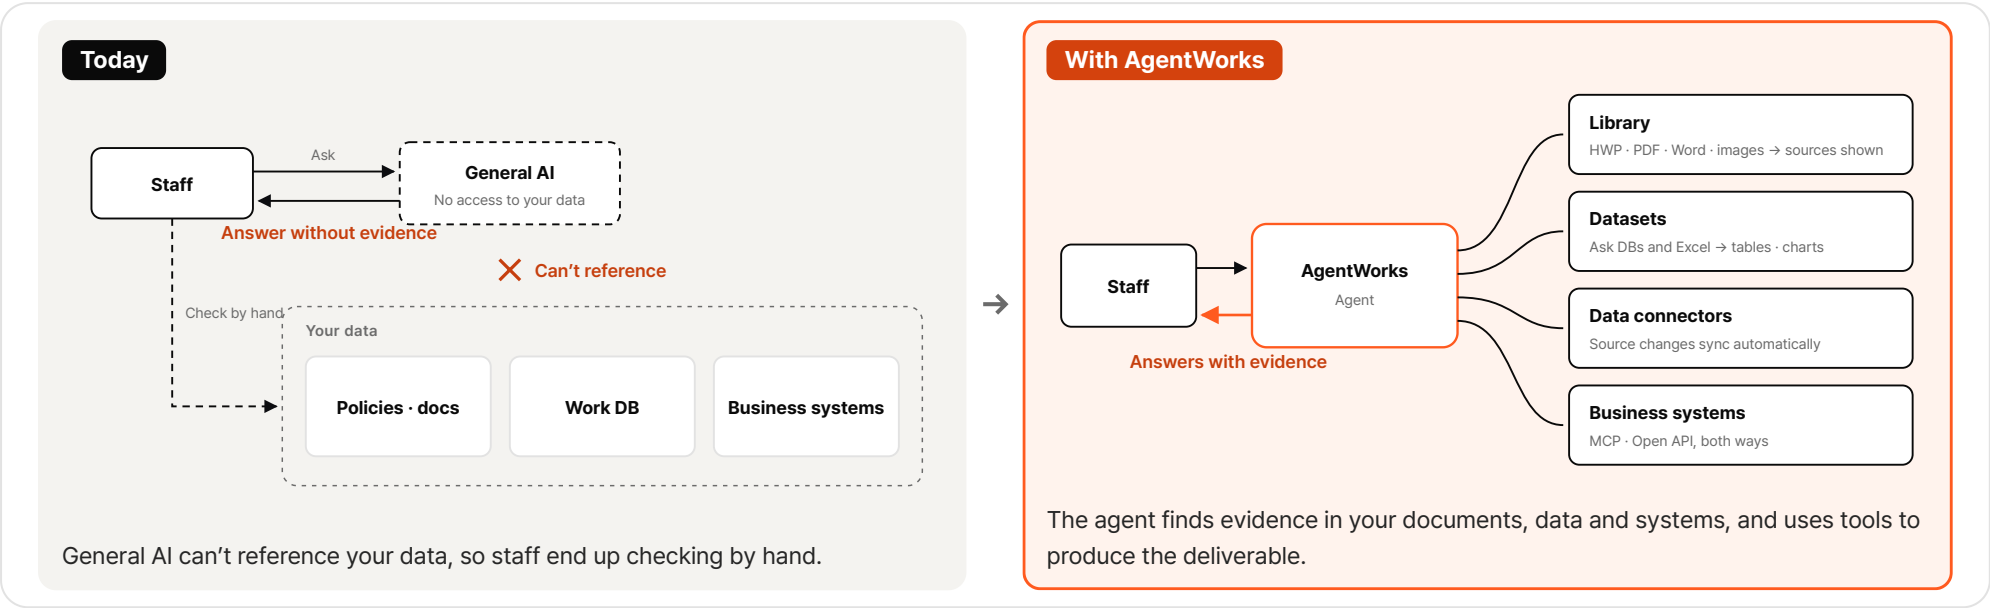
To see the numbers in a database or spreadsheet, someone has to extract and tidy them by hand.

The worry underneath “That people end up checking and fixing every AI answer anyway.”

See how AgentWorks solves it →

Solution 03 · Connected knowledge, data and systems

Your documents, databases and systems become the agent's evidence and hands.



Library · documents become knowledge

PDF, Word, HWP and image (OCR) files become searchable knowledge, and every answer shows its source documents.

Datasets · ask in plain language

Ask internal databases, Excel or CSV without SQL and get tables and charts back.

Connectors · automatic updates

Nine types of sources, including S3, SFTP, shared folders, databases and Redmine, are synced on a schedule.

External tools · real work done

Ten built-in tools plus internal tools connected over MCP handle lookups, writing and sending.

What you get

- ✓ See the source document for every answer
- ✓ Get tables and charts without knowing SQL
- ✓ Changes at the source show up automatically

Concern 04 · The burden of control and proof



We can’t tell who sent which data to which AI.

The question security, audit and management ask right before deciding.

What actually happens

Usage records are scattered

Each department uses a different AI and key, so there is no single view of usage.

Personal data goes straight in

Names and phone numbers can reach the AI as they are, with no check in between.

Costs and limits are hard to set

Without knowing which department uses how much, budgets are hard to set and keep.

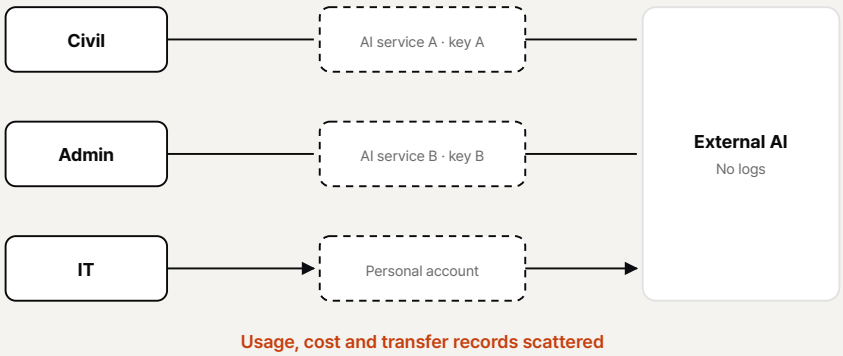
The worry underneath “Having no evidence to show when the audit comes.”

See how AgentWorks solves it →

Solution 04 · One point of control with AI Gateway

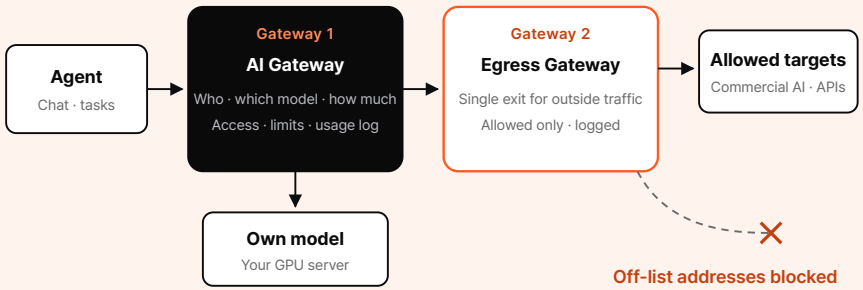
AI calls are controlled and recorded through two gateways.

Today



Each department uses its own AI and keys, so no one can see who sent what in one place.

With AgentWorks



Every call passes the same two gateways, and usage and outbound connections are logged.

What it blocks

Sending to AI services you haven't allowed, model use by departments without access, and use beyond the limit.

What it records

Usage by model and user, outbound connections, and logins and key actions in the audit log.

Department gateways

Assign models and keys per group and stop them instantly by deactivating. Set credit limits by contract, period and user.

Security by configuration

Tenant data isolation, role-based access, sharing and delegation, lockout after five failed logins, and session expiry come built in.

What you get

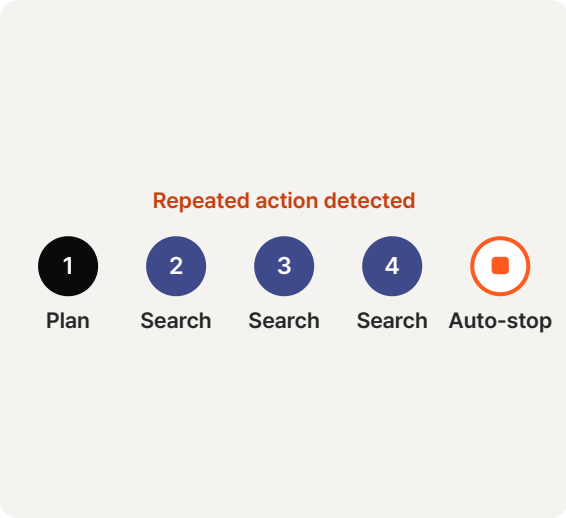
- ✓ Prove who sent what, from the records
- ✓ See usage and cost by department in one place
- ✓ Get control and evidence by configuration, without development

Solution 04 · Safe to hand work over

Even while an agent works on its own, three safeguards run alongside it.

Run safeguards

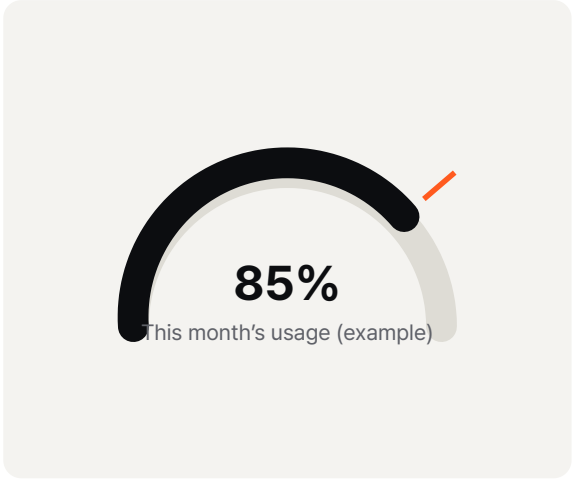
Runaway prevention



It **stops on its own** when the same action repeats or step or time limits are exceeded.

Budget · limits

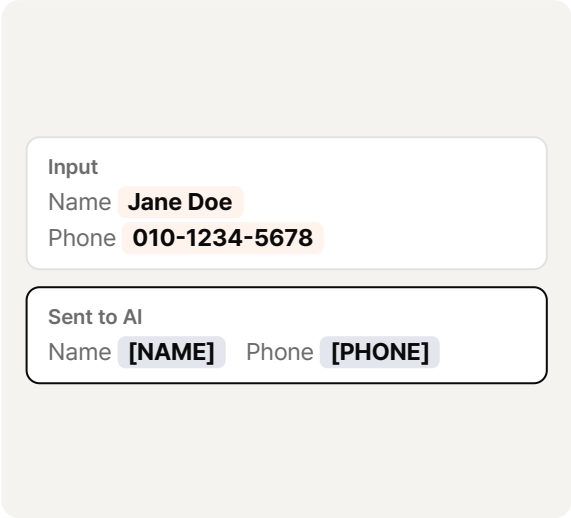
Cost ceiling



When per-task token budgets or contract, period or user limits are exceeded, **new runs are blocked**.

Pre-run check

Personal data protection



Personal data is detected at input and **blocked or redacted** before the run.

Quality you can measure

Register test cases
common questions with required and forbidden conditions

Run combinations
vary model and reasoning effort over the same cases

Grade
rule-based and AI grading decide pass or fail

Compare · improve
pass rate, score, latency and cost, before and after each change

Owner approval for irreversible actions and automatic recovery across screen changes are in development.

Deployments

Different organizations work on the same platform.

In use

Company A

AI consulting for the public sector

RFP analysis · proposal writing

Library · Canvas · Templates

Cloud

In use

Company B

Group IT services affiliate

Company-wide work AI

Chat · Library · Office tools

Cloud

Building

Agency C

Public · education

Education network incident intake · analysis

Open API · Tasks · Connectors

On-site

Pilot

Company D

Network equipment maker

Technical support center automation

Library · Tasks · PII redaction

Hybrid

Design · planned

Agency E

Public · energy

Large-scale data extraction · analysis

Datasets · Own model

On-site (closed network)

Deployment map · the higher the security need, the more on-site

Cloud
Use right after sign-up
Co. A · Co. B

Hybrid
Sensitive work inside · some commercial AI
Co. D

On-site (closed network)
Own model · no outside access
Agency C · Agency E

Lower security needs — Higher security needs

In common Different work and security settings, yet all on the same platform, building and running agents.

Names are anonymized at the customers' request. Status as of October 2026.

Case details

Build, connect, secure — in real deployments

Agency C · public · education **Building · on-site**

Education network incident intake · analysis

1,600+

incidents over four years used as evidence

1. Staff report by chat, no login needed
2. A dedicated intake app calls tasks via Open API
3. Similar cases → likely causes · check guide
4. Operator reviews the analysis and closes
5. Resolved cases feed the next analysis

Records, not memory
likely causes and checks drawn from similar cases

Deadline tracking
deadlines in business hours · early alerts

Knowledge that grows with use
added to references as soon as closed

Build · run	Dedicated app calls tasks via Open API
Connect	Four years of incident history as evidence
Governance	5 roles · regional access · IP, phone and password redaction

Company D · network equipment **Pilot · hybrid**

Technical support center automation

200,000

past issues searched and analyzed by AI

1. Classified by type and severity on arrival, owner notified
2. First-response plan from the top 3 similar issues
3. Reads logs and lists likely causes
4. Applied after an engineer reviews it
5. Results written up as reports and FAQs

People make the final call
engineers review AI output before it is applied

Existing setup kept
same portal and account system as before

Enabled without development
similar-issue search · usage control · PII redaction

Build · run	First-response and auto-classification tasks
Connect	Search across 200,000 past issues
Governance	roles · audit log · PII redaction · existing portal and accounts kept

Agency E · public · energy

Design case · planned

Large-scale analysis on a closed network

20,000+

tables · hundreds of TB · tens of thousands of requests a year

1. Analysts ask in plain language
2. Metadata picks the right tables and builds SQL
3. Syntax and query-plan checks block risky queries
4. Read-only queries straight against the source
5. Tables and charts with the SQL and sources shown

Principle 1
Data never moves

Principle 2
Metadata raises accuracy

Principle 3
Load and security control come first

Build · run	Natural-language requests become SQL analysis tasks
Connect	Queries the source directly, no data moved
Governance	closed network · own model · time and row limits · audit

How to start

Adoption starts small and grows.

1 **Consultation**
We review your work and security requirements.

2 **Pilot**
Prove the value on one or two key tasks. A closed-network pilot is also possible.

3 **Select · design tasks**
Organize the tasks and data to turn into agents.

4 **Deploy · train**
On-site installation and staff training.

5 **Expand**
Add agents department by department.

Stage 1 · Pilot

One team · 1–2 key tasks

Civil

Start with work that shows results fast, and prove outcomes and operations first.

Stage 2 · Expand

Agents added in each team

Civil Admin IT

Other departments reuse proven tasks and skills as they are, adding their own agents without development.

Stage 3 · Whole organization

The agent factory inside your organization

Civil Admin IT Audit Planning +

Agents from every department run under the same access, limits and audit rules, managed from one AI Gateway.

Departments and agent counts are examples.

Included
training · manuals · technical support

Pricing
by deployment type and scale

Purchase inquiries
VICS Technology

Public procurement (Korea)
GS certified · available via procurement

Next step

We propose a pilot on your real work.

Start with one or two tasks that show results fast, prove them inside your organization on your own model,
and plan the rollout together from the results.

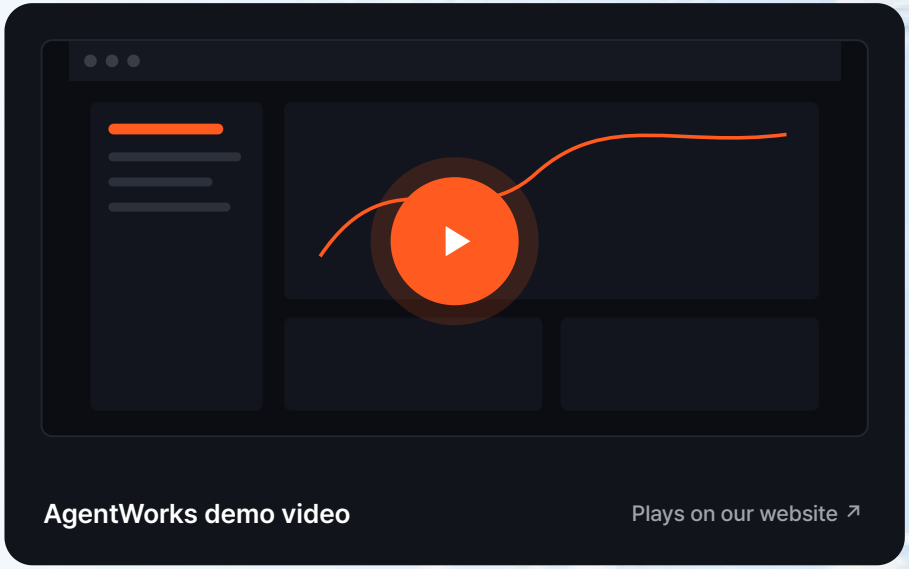
Choose key tasks
1–2 tasks that show results fast

Closed-network pilot
proven inside on your own model

Results report
impact and rollout plan

Contact us →

▶ Watch the demo



Email
sales@vicsnc.com

Website
vicsnc.com