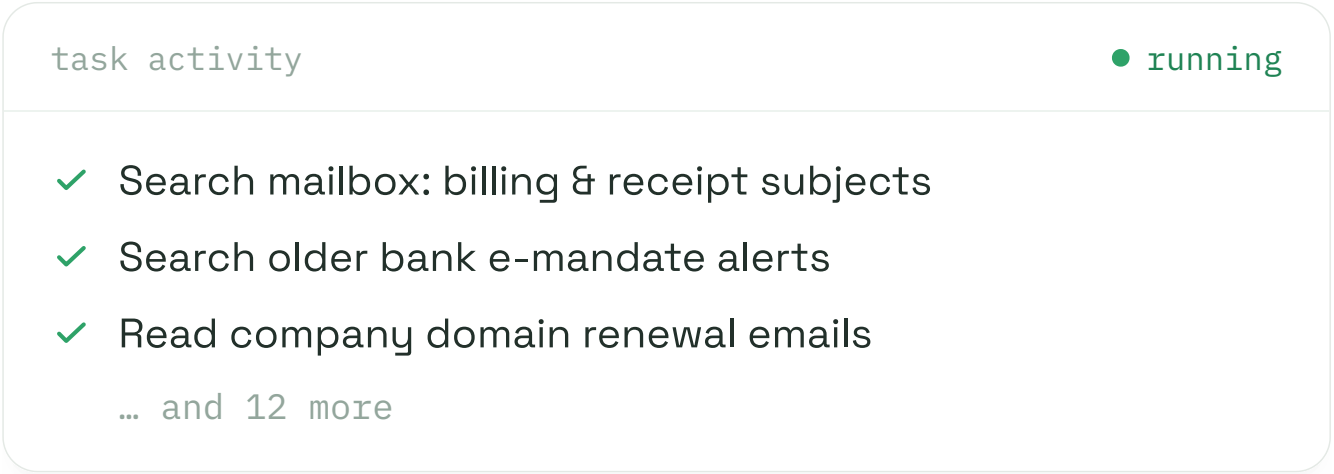
I asked an AI chatbot to find my subscriptions.

It went through my entire mailbox.



What it did, on its own:

01	SEARCH	Billing & receipt subjects
02	SEARCH	Recurring charge & mandate phrases
03	SEARCH	App store & streaming receipts
04	SEARCH	Software & cloud billing emails
05	SEARCH	Older bank e-mandate alerts
06	SEARCH	Receipts & renewals, late 2025 → mid 2026
07	SEARCH	Streaming, delivery & app store receipts
08	READ	App store, cloud storage & bank threads
09	SEARCH	Apple, Netflix, Prime & OpenAI receipts
10	SEARCH	YouTube, telecom & food-delivery billing
11	READ	Bank, developer, domain & delivery threads
12	SEARCH	Recent emails, Oct–Nov 2026
13	READ	Telecom bill & domain auto-renewal
14	SEARCH	Recent bank e-mandates
15	READ	Company domain receipts

15
mailbox searches & reads

1
broad OAuth consent

0
per-action approvals

Recreated from the Gemini app's Task activity.
Sender names generalised.

One consent. Unlimited reach.

A non-deterministic system decided for itself what to search, what to open and what to read.

WHAT THE TASK NEEDED	WHAT THE GRANT ALLOWED
✓ Billing emails only	— Every sender, every thread
✓ Last 12 months	— Your entire history
✓ Sender, amount, date	— Full message content
✓ Read-only, once	— Until you revoke it

Same prompt tomorrow? A different path through your data.

Now picture it inside the enterprise.

01

Broad grants

Gmail, Drive, CRM and ERP opened up through one OAuth consent.

02

No tool-level limits

Nothing stops which tools the agent calls or which data it reaches.

03

No attribution

No record of which user, agent or model touched which data.

04

Drift, unchecked

A prompt or plan slightly off course, and nothing in the way.

> We'd never give a new hire "read everything" on day one. Why do we give it to AI agents?

Granular OAuth control for every MCP server.

AI agents & assistants

Chat · copilots · developer tools · autonomous agents

SecuriX AI gateway

ENFORCED

Identity · per-tool policy · DLP · audit · MCP governance

Existing & managed MCP servers

Mail · docs · CRM · ERP · enterprise data & tools

- ✓ Policy boundaries per agent, tool and data source
- ✓ PII redacted before it reaches the model
- ✓ Tool calls and data access governed at the action layer
- ✓ Every action attributed to user, agent, model and tool

AI agents are going to touch enterprise data. The only question is how:

—	✓
Allow all	Guardrails you define

Book a demo at securix.app →

Built for BFSI, insurance, healthcare and enterprise IT.