sentra

AI Governance Starts with the Data, Not the Model

Yair Cohen, CPO & Co-Founder | Nov 17th, 2026

Your AI stack has three layers.
AI Data Readiness layer is missing.

Sentra answers the most important AI question:
What can AI see, what can AI do, and how do you continuously govern it?

01

AI Infrastructure Layer

Models, pipelines, GPUs, deployment.
The layer that got funded first.

02

AI Governance Layer

Guardrails for AI behavior.
Blind to the data underneath.

THE MISSING LAYER

AI Data Readiness Layer (Trust Layer)

Continuous discovery, classification, hygiene, and access governance—at petabyte scale, in your environment.
The missing layer AI has been running without.

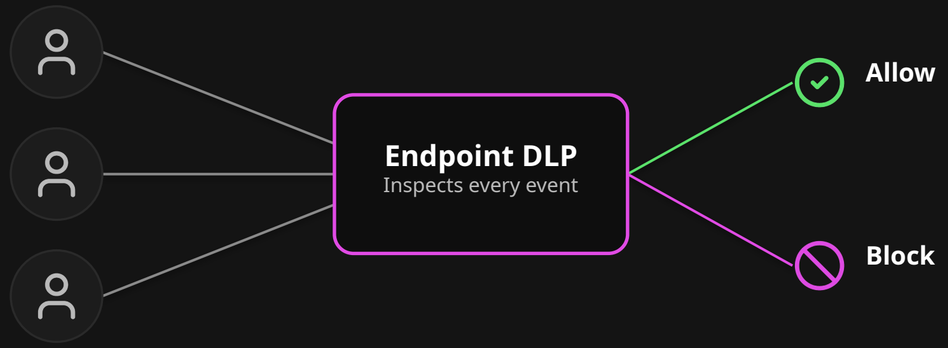
The agentic era changes everything

Agents multiply data interactions. Security must shift from policing each event to governing access itself.

UNTIL NOW

Reactive data security

Built for human-scale data access

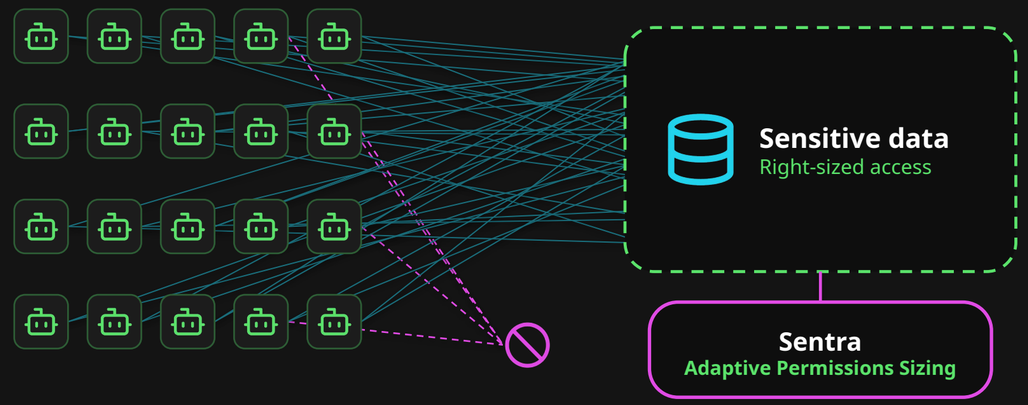


- Humans access data one interaction at a time
- Volume is low enough to inspect every event
- **Decide in the moment: allow or block**

THE AGENTIC ERA

Proactive data security

Built for machine-scale data access

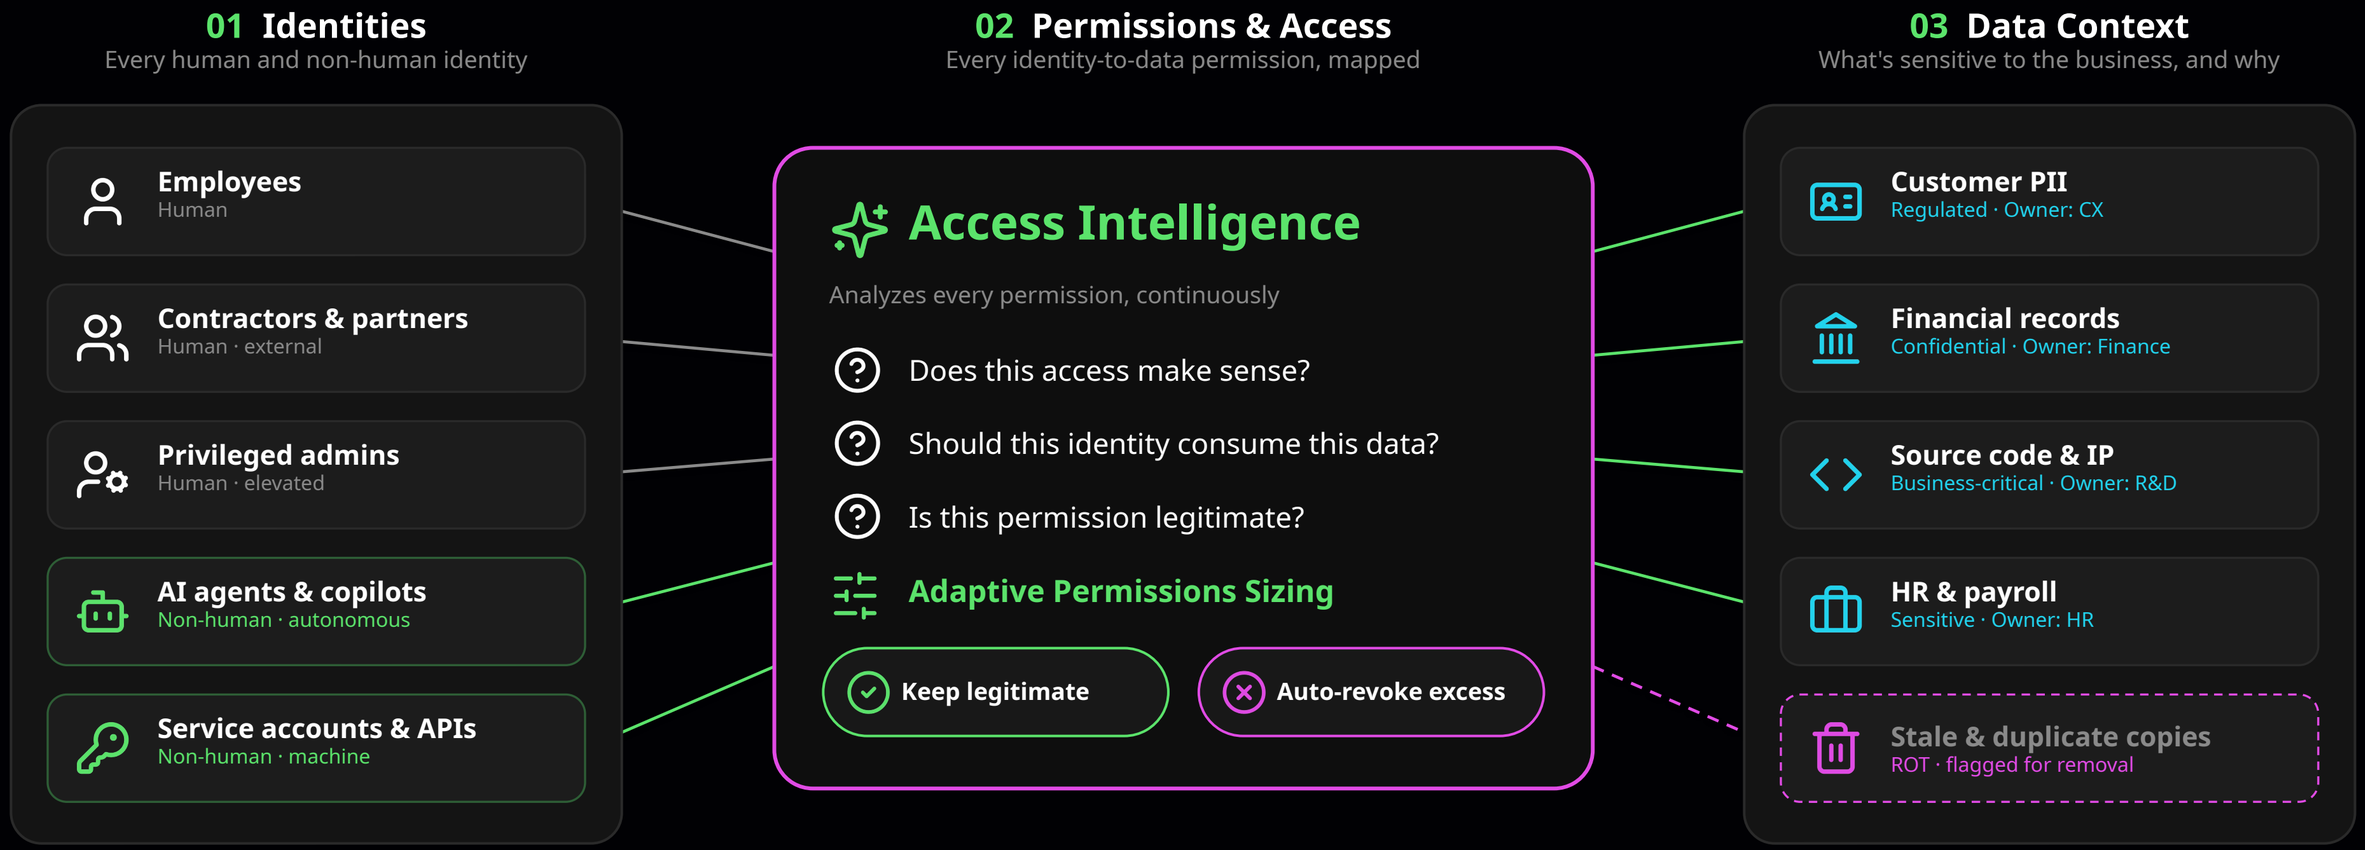


- Agents multiply interactions by orders of magnitude
- Inspecting each one at machine speed doesn't scale
- **Right-size what agents can reach, before they act**

Gartner predicts the agents reaching your data will multiply 10,000× in three years. Source: Gartner, "Gartner Identifies Six Steps to Manage AI Agent Sprawl," April 28, 2026.

Securing data for the agentic enterprise

Know every identity. Understand every dataset. Make sure every access is legitimate.



Proactive, not reactive: Sentra continuously right-sizes what every human and AI agent can do with sensitive data, before it is used.

Agentic Data Security

Proactive approach: every human and AI agent has exactly the access it needs, and no more.

01

Every agent, accounted for

Discover every AI agent, copilot and MCP server across cloud, SaaS and endpoints, with its identity and the sensitive data it can reach.

Roadmap track: AI Security + Identities

02

Access that has to make sense

For every human and agent, Sentra answers: should this identity have this access? Every verdict is explained.

Roadmap track: Access Intelligence

03

Fixes you can trust to run on their own

Remove excess access automatically across SaaS, cloud and data platforms, with dry run, blast radius and rollback built in.

Roadmap track: Automatic Remediation

04

Protection where AI is used

Follow sensitive data into prompts, browsers and devices, and use those signals to prioritize risk in the cloud.

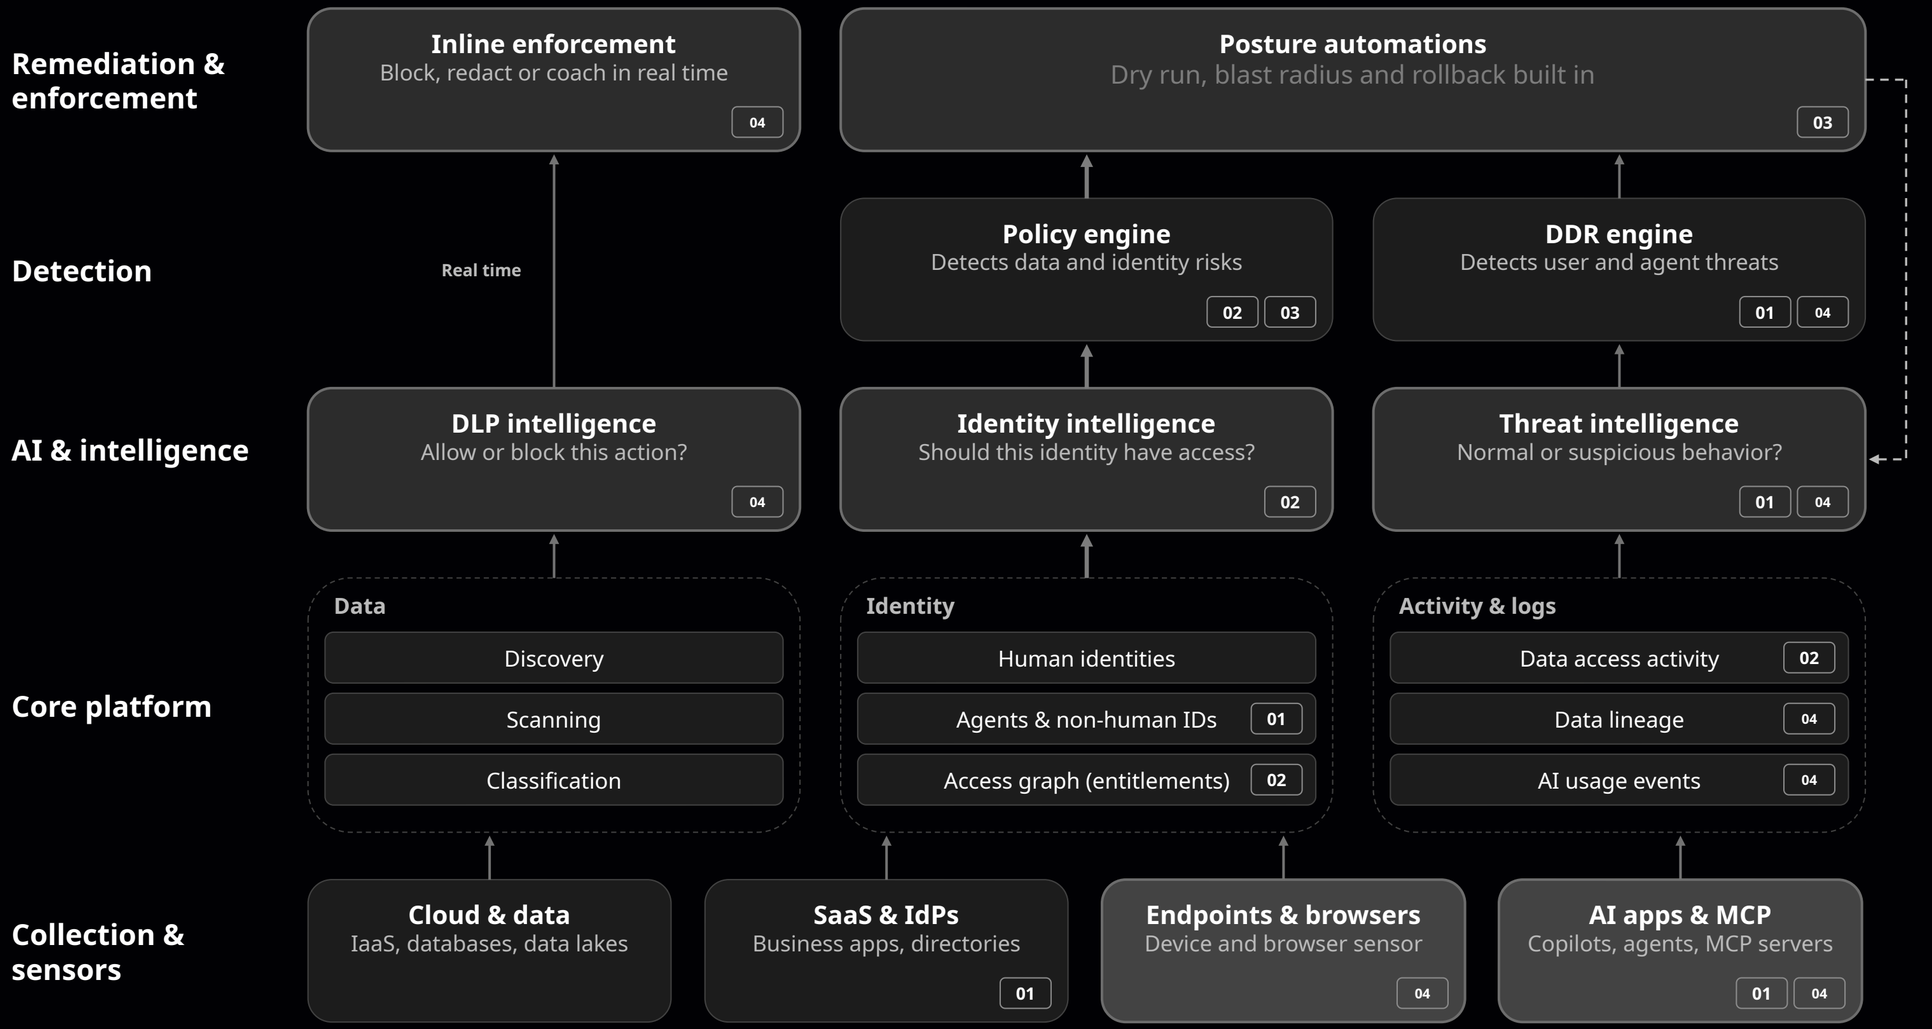
Roadmap track: Endpoint DLP

All four converge on Adaptive Permissions Sizing

Continuous right-sizing of access for humans and agents, on the same Sentra engine you already have deployed.

Sentra Platform Architecture

Five layers, three new intelligence engines and one learning loop: what we build to deliver the four themes.



Roadmap themes

- 01 Every agent, accounted for
- 02 Access that has to make sense
- 03 Fixes you can trust to run on their own
- 04 Protection where AI is used

sentra

The Sentra platform: one engine, built for agentic AI

01 REQUIREMENT
Know the data AI can reach
What data can AI reach, and what is it?
Continuous discovery · Business-context classification

02 REQUIREMENT
Govern who and what can access it
Should it reach that data, and what can it do with it?
Reachability · Agent least privilege · Auto-remediation

03 REQUIREMENT
Keep it governed over time
How do we stay governed as everything changes?
Continuous state · ROT removal · Audit-ready evidence

AI Security
Inventory every agent, copilot and model, and the sensitive data each one can reach.

Privacy Operations
Know whose personal data AI touches and where, with audit evidence for privacy and AI regulation.

Data and Access Governance
Map what every human and agent can reach, and turn access sprawl into least privilege.

Insider Risk Management
Treat agents as insiders: catch misconfigured, compromised or misaligned agents early.

Next-Gen DLP
Right-size agent permissions, apply AI-native guardrails, and monitor usage and prompts.

Cyber Resilience
Keep AI-critical data backed up and encrypted, and recover fast when agents alter or delete it.

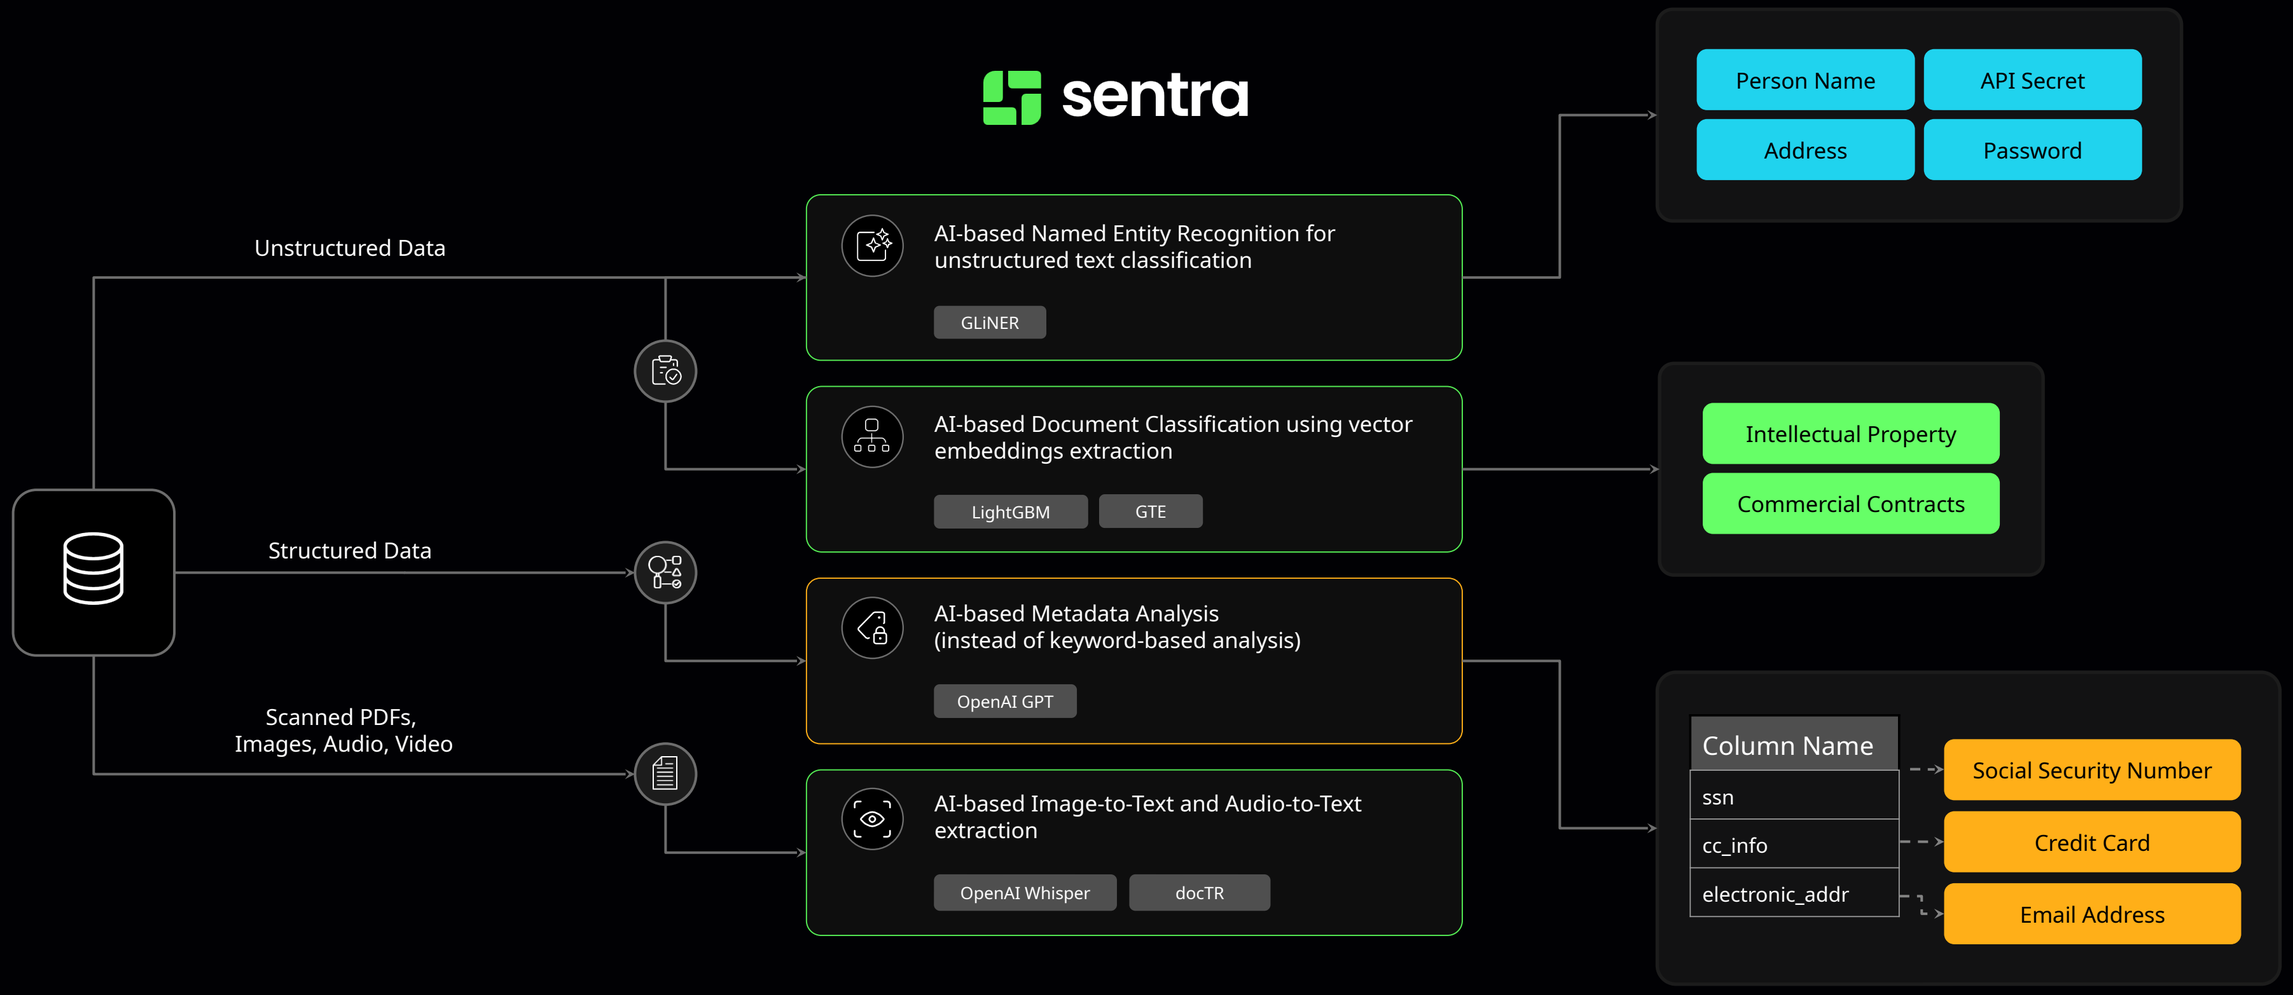
Storage & Cost Optimization
Retire stale and duplicate data before agents find it, cutting exposure and storage spend.

Sentra Engine
Live today

Discovery & classification · Business context · Identity-to-data reachability · Detection & remediation · In-environment, agentless

Every module runs on the same engine and data model, so nothing deployed has to be rediscovered.

AI-Powered Classification



Why Customers Choose Sentra

Agent Access Visibility

Map each agent, its knowledge bases and the sensitive data it can reach. Copilot Studio, Bedrock, ChatGPT Enterprise.

AI-Powered Classification

250+ classifiers and AI models find the data only you have, with context like customer vs employee. 98%+ accurate.

Data Stays in Your Environment

Scanners run in your cloud with read-only roles. Only metadata leaves, so the data behind your agents is never copied.

Comprehensive Coverage

M365, Slack, Snowflake, S3, on-prem shares and unmanaged databases: every store your agents pull from.

Petabyte-Scale Speed & Cost

Smart sampling and incremental scans classify petabytes in days, at a price that holds from POC to production.

Shrink the Blast Radius

Find exposed API keys, tokens and private keys before an agent or attacker does. Know what to rotate in minutes.

Trusted by

Nestlé · Expedia · SoFi · Munich RE · XPO Logistics · MARQETA · PENNYMAC · lyft · Global-e · ADP · instacart · RIVIAN

Gartner Peer Insights
Highest Recommended DSPM Platform

Gartner Peer Insights™

4.9 ★★★★★

By Sentra in Data Security Posture Management (DSPM)

98.5%

Recommendation Rate