sentra

Solution Brief

The Agentic Data Security Platform

Know what sensitive data every AI agent can reach, and right-size that access before it is used.

AI agents are becoming one of the main ways enterprise data gets accessed. Gartner expects the average global Fortune 500 enterprise to run over 150,000 agents by 2028, up from fewer than 15 in 2025. Each agent acts on its own, chains actions across tools, and inherits permissions that were originally granted to people. As a result, years of data sprawl, stale copies and excessive access are now reachable by software that runs without a person reviewing each step.

Existing controls were not designed for this. DLP inspects individual interactions and allows or blocks them, which works when traffic comes from people. Inventorying agents one by one falls behind for a similar reason: new agents are deployed faster than teams can catalog them. Neither approach shows which data an agent can actually reach, which is where the risk sits.

Sentra Agentic Data Security Platform

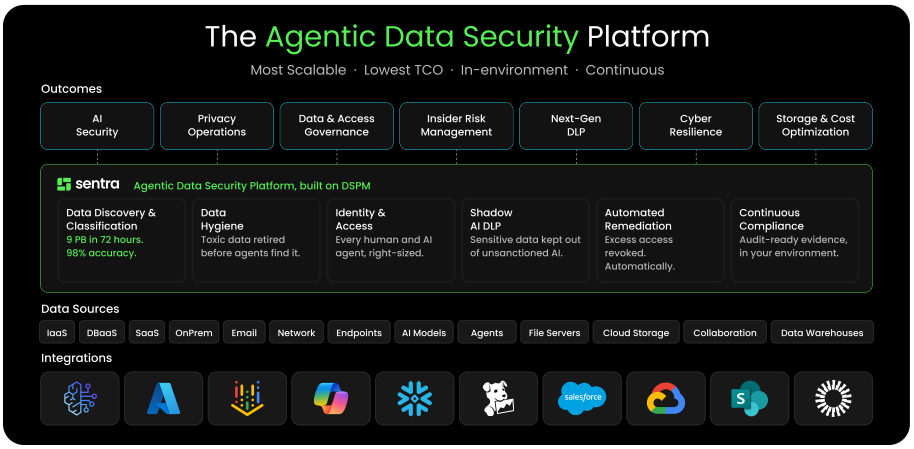

Sentra is the agentic data security platform built on DSPM. It discovers and classifies sensitive data across cloud, SaaS, databases, data warehouses, on-premises systems and AI environments. It then maps which human and AI identities can reach each data store and reduces excess access before it is used, through Adaptive Permissions Sizing.

Because Sentra maps access starting from the data, a new agent is in scope as soon as it is granted permissions. It does not need to be registered or tracked separately.

Sentra's scanners are read-only, managed by Sentra, and run inside the customer's own cloud account. Discovery, classification and analysis happen in the customer environment. Only metadata is sent to Sentra.

Sentra answers three questions, continuously:

01

What sensitive data can each agent reach?

02

Should it be able to reach it?

03

What can it do with that access?

Solution Highlights

Your data never leaves your environment.

Sentra-managed, read-only scanners run inside your own cloud account, and only metadata returns to Sentra. There's no copy of your data to protect and no third-party processing review, so security and legal approvals move faster.

Know what every AI agent can reach.

Every copilot, agent, service account and person is mapped to the sensitive data it can reach, including inherited, unused and externally shared access. Excess access is removed before an agent uses it.

Labels you can trust to drive agent permissions.

Classification runs at 98% accuracy with under 1% false positives, with business context on every finding. Native labels feed Purview, Google Workspace, Snowflake and Databricks, so your existing DLP, IAM and AI guardrails enforce on the same signal.

Petabyte coverage without the overhead.

Sentra has scanned 9 PB in under 72 hours, and scanners scale to zero when idle. Incremental scans read only what changed, with no appliances to deploy and no scanning infrastructure for your team to run.

© Sentra 2026. All rights reserved. Confidential.

Key Capabilities

The Sentra Agentic Data Security Platform, built on DSPM

Data Discovery and Classification
See every data store, AI surface and sensitive record before an agent does.

Data Hygiene
Retire redundant, obsolete and toxic data before agents find it.

Identity & Access
Right-size what every human and AI agent can reach, before they act.

Shadow AI DLP
Keep sensitive data out of unsanctioned AI apps and agents.

Automated Remediation
Controls that act, not just alert. Excess access is revoked automatically.

Continuous Compliance
Audit-ready evidence that keeps pace with AI adoption.

Data Discovery and Classification

See every data store, AI surface and sensitive record before an agent does. Agentless discovery and classification runs across cloud storage, SaaS, collaboration platforms, databases, warehouses, on-premises systems and AI environments. Nothing runs on production hosts and nothing sits in the path of your data. Sentra keeps a live picture of the estate rather than a point-in-time report, including abandoned, duplicated and shadow data that no inventory reflects. More than 250 out-of-the-box classifiers across more than 130 file formats, including images, audio and video, run on scanners inside your environment, and customer data never reaches an external LLM. Smart sampling groups similar objects so petabyte-scale stores are classified in days while every user-written document is still opened. Discovery covers AI agents too: Sentra inventories agents, copilots and models, and the sensitive data each one can reach, including RAG and training sources. Each asset arrives with its sensitivity, business context (employee or customer, residency, real or synthetic), its security configuration, and every human and AI agent that can reach it. Context, not patterns: 98% accuracy with under 1% false positives.

Built for the Agentic Era

AI agents covered.

Microsoft Copilot Studio, Amazon Bedrock and ChatGPT Enterprise, plus the data behind them in Microsoft 365, Snowflake, Databricks and more.

Access that follows the data.

Entitlements track lineage from S3 into a Bedrock knowledge base, or from Snowflake into Databricks.

Surgical remediation.

Sentra removes one external user, disables one public link, or quarantines one file, without broad revocations that break the business.

Compliance evidence built in.

Evidence for GDPR, HIPAA, CCPA and the EU AI Act is generated inside your environment.

© Sentra 2026. All rights reserved. Confidential.

Data Hygiene

Retire redundant, obsolete and toxic data before agents find it. Duplicate, stale and unnecessarily exposed data is identified across the estate, including stores no inventory reflects. Duplicate and obsolete data can be deleted outright, high-risk files quarantined to a location users cannot reach, or the finding routed to the owning team with policy context attached. Deleting data nobody needs is the cheapest control: it shrinks what any agent can reach and lowers storage and processing cost.

Identity & Access

Right-size what every human and AI agent can reach, before they act. Sentra maps sensitive data to every employee, contractor, service account, application and AI agent that can reach it, and evaluates the real access path rather than a permission list. Entitlements follow the data through lineage as it moves, from Snowflake to Databricks or from S3 into a Bedrock knowledge base. Adaptive Permissions Sizing keeps legitimate access, removes the excess and keeps re-sizing as data, agents and usage change, so least privilege becomes continuous rather than a quarterly review.

Shadow AI DLP

Keep sensitive data out of unsanctioned AI. Sentra discovers the AI apps and agents in use across the enterprise and stops sensitive data from being shared with them. Instead of adding an inline checkpoint that cannot keep up with machine-speed traffic, Sentra right-sizes agent permissions ahead of time and feeds accurate labels to the native guardrails of the AI platforms and the DLP you already run. Agent usage and prompts are monitored, so misuse shows up early.

Automated Remediation

Controls that act, not just alert. Sensitivity labels are applied automatically across Microsoft 365 through Purview Information Protection, Google Workspace, cloud storage tagging, Snowflake, Databricks, Salesforce and data catalog tools, so existing DLP, CASB and SSE, IAM, AI gateway, ITSM and cloud controls enforce on the same signal. Direct action is surgical: remove a specific external user, disable a public link, revoke over-privileged access or quarantine sensitive data. High-confidence violations are fixed automatically, while ambiguous cases route to human review.

Continuous Compliance

Audit-ready evidence that keeps pace with AI adoption. Scanning runs on an ongoing basis rather than as a periodic audit, so new service accounts, agents, connectors and data stores enter the picture as they appear. Evidence is tied to datasets, classification decisions, access relationships and AI identity mappings, covers GDPR, HIPAA, CCPA and the EU AI Act, and is generated inside the customer environment. No data leaves, and no third-party processing attestation is required.

Solution Benefits

Say Yes to AI, Faster.

Copilot and agent rollouts clear security review faster because you already know what each agent can reach.

Fewer Agent-Driven Exposures.

Excess access is removed at the data layer before an agent can use it.

Lower Cloud and Storage Spend.

Retiring redundant, obsolete and toxic (ROT) data cuts storage and processing costs and shrinks what agents can reach.

Less Operational Burden.

Fully managed scanners and selective auto-remediation free your team from running infrastructure and chasing low-risk alerts.

Audit-Ready by Default.

Continuous evidence for GDPR, HIPAA, CCPA and the EU AI Act replaces audit-by-audit scrambles.

Sovereignty Without Exceptions.

In-region classification meets data residency requirements by design.

© Sentra 2026. All rights reserved. Confidential.

Built to run inside your environment

Sentra scales scanner count to the size of the environment and returns it to zero when idle, so no compute cost persists between scans. Incremental scanning rescans only what changed, smart clustering groups similar objects before sampling, and partial file reads sample only the portion of a large file needed to classify it. The price you see in the POC is the price at production volume. Scanning is indirect by design, working from snapshots, restored replicas and API-mediated reads rather than production systems. Classification runs in the region where the data resides, so sovereignty requirements are met by architecture, and data is never moved between regions.

One platform for agentic data security

Sentra provides a single platform, built on DSPM, for discovering, understanding, governing and protecting the data that every human and AI agent can reach, with:

- **A single live inventory of every data store, AI agent and AI surface, and of the sensitive data inside each one.**
- **A single access map covering every human, service account, application and AI agent identity.**
- **A single classification signal that the DLP, IAM, AI guardrails and cloud controls you already own can enforce on.**

About Sentra

Sentra is the agentic data security platform. It shows what sensitive data every human and AI agent can reach and right-sizes that access before it is used. Built on DSPM, it runs inside your own environment, so your data never moves and your team never runs scanning infrastructure.

Proven. Trusted. Recommended.

Certified and independently tested.
SOC 2 Type II and ISO 27001, with third-party penetration testing every six months and automated vulnerability testing on an ongoing basis.

Built for enterprise scale.
9 PB scanned in under 72 hours. 98% classification accuracy with under 1% false positives.

Your data never leaves.
Discovery, classification and analysis run inside your own environment.

Don't run after your agents.
Run after your data.

Schedule a Demo

Gartner Peer Insights Recommended DSPM Platform

Gartner Peer Insights℠

4.9 ★★★★★
By Sentra in Data Security Posture Management (DSPM)

98.5%
Recommendation Rate

Trusted by

lyft · Glovo · septerna · GlobalE · Valenz Health · Renaissance · PENNYMAC · SoFi

© Sentra 2026. All rights reserved. Confidential.