LangProtect

Securing Employee AI Usage Across Your Organisation

Introduction

Your employees are using AI every day. Most of them without telling IT.

Right now, someone on your team is pasting a client contract into ChatGPT to get a summary. Someone in finance is asking Claude to draft a report that includes account figures. A developer is dropping source code into Gemini to debug a function. None of them think they are doing anything wrong. None of them have been told what happens to the data once it leaves their browser

And your security team has no idea any of it is happening.

This is not a prediction about where AI adoption is headed. This is the current state of most organizations, including yours. Employees have discovered that AI tools make them dramatically more productive. They are not waiting for IT governance to catch up. They are not reading the acceptable use policy before they paste. They are working, and the tools are helping, and the data is flowing into public LLM endpoints that your security stack cannot see.

The problem is not that employees are using AI. The problem is that every organization has a growing gap between the AI policy written in a document and the AI reality happening on every laptop, every day.

Your DLP tool monitors file transfers and email attachments. It does not read prompts. Your CASB solution tracks which applications employees access. It does not see what employees type into those applications. Your firewall knows ChatGPT exists and may even allow it. It has no idea whether the prompt sent to ChatGPT contained a patient record, a trade secret, or a client's social security number.

The result: organizations are approving AI adoption without any visibility into what that adoption actually looks like, or what it is costing them in data exposure.

LangProtect Guardia was built to close that gap.

Why What You Have Isn't Enough

The instinct of most security teams when confronted with unmonitored AI usage is to reach for existing tools. Extend the DLP policy. Add ChatGPT to the CASB catalog. Block unauthorized AI domains at the firewall. These responses are understandable. They are also insufficient, not because the tools are poorly implemented, but because they were designed for a fundamentally different threat surface.

Understanding why requires a precise look at where existing controls operate and where the AI prompt layer sits in relation to them.

> Traditional security tools were not built for AI interactions. They were built for files, emails, and network traffic. The prompt is none of those things.

THE DLP BLIND SPOT

Data Loss Prevention solutions operate on structured data, documents, spreadsheets, email attachments, endpoint file activity. They are pattern-matching engines trained to recognize specific data types in specific formats: a sixteen-digit sequence that matches a credit card pattern, a nine-digit sequence that matches a Social Security Number in a file transfer.

A prompt is none of these things. When an employee types “Here is our Q3 patient census data, please summarize the key trends”, into ChatGPT and pastes a table of patient records directly into the text field, that interaction does not trigger a DLP alert. There is no file. There is no attachment. There is no structured transfer event. There is only a browser session and a text input, both of which traditional DLP was never designed to inspect.

> Organizations using enterprise DLP solutions have, on average, less than 10% visibility into the sensitive data flowing through employee AI interactions.
>
> *Source: LangProtect platform data*

THE CASB LIMITATION

Cloud Access Security Broker solutions were a significant advancement in SaaS governance. They gave security teams visibility into which cloud applications employees were accessing and the ability to enforce access policies at the application level. For a generation of SaaS tools, this was sufficient.

AI tools break the CASB model in two critical ways.

First, the interaction surface is conversational. CASB solutions track application access events, login, file upload, download, share. They do not capture the semantic content of what happens inside a session. A CASB tool can tell you that an employee accessed ChatGPT for forty-five minutes on a Tuesday. It cannot tell you that the employee spent those forty-five minutes uploading client case files and requesting contract analysis.

Second, the application catalog is expanding faster than governance can track. New AI tools launch continuously. Employees discover and adopt them independently. By the time an AI application is reviewed, risk-rated, and added to a CASB policy, it has already been in active use across the organization for weeks.

In a representative 30-day deployment of LangProtect Guardia across a mid-size enterprise, the Shadow AI discovery dashboard identified the following employee AI tool usage:

- ChatGPT.com — 90% of AI interactions
- Claude.ai — 6% of AI interactions
- Perplexity.ai — 1% of AI interactions
- Chat.mistral.ai — 1% of AI interactions
- Gemini.google.com — 2% of AI interactions

The security team had formally approved zero of these tools for use with sensitive business data.

THE FIREWALL CANNOT READ A PROMPT

Network-level controls, firewalls, secure web gateways, DNS filtering, operate at the connection layer. They can allow or block access to a domain. They cannot inspect the payload of an encrypted HTTPS session, which means they have no visibility into the content of what is being sent to any AI platform.

Blocking ChatGPT at the firewall is a blunt instrument that eliminates productivity without eliminating risk, employees immediately route to mobile devices, personal hotspots, or the dozens of alternative AI platforms that are not yet on the block list. It is governance by restriction rather than governance by control.

The organizations that have attempted to solve this problem through domain blocking consistently report the same outcome: employees find alternatives, shadow AI usage increases, and the security team has less visibility than before the block was implemented.

37% of all employee AI prompts analyzed by LangProtect Guardia contained sensitive data, including PII, PHI, financial records, and confidential business information.
In the same dataset: 1,425 sensitive prompts detected across 3,893 total interactions. 1,263 blocked before reaching the LLM. 162 exposed after user override.

THE COMPLIANCE ACCOUNTABILITY GAP

Beyond the technical limitations of existing tools, there is a compliance accountability problem that DLP, CASB, and firewall solutions cannot solve even in combination: the absence of an auditable record of AI interactions.

When a HIPAA auditor asks an organization to demonstrate that patient health information was never transmitted to a public AI platform, the honest answer, in most organizations today, is that they cannot prove it either way. There is no log. There is no record of what was typed into ChatGPT by which employee on which date. The interaction happened in a browser session that no existing security control captured.

This is not a hypothetical audit scenario. HIPAA enforcement actions increasingly include questions about AI tool usage. The HHS Office for Civil Rights has explicitly noted that covered entities must evaluate whether their use of AI tools involving PHI meets the requirements of the HIPAA Security Rule. An organization without AI interaction logs has no defensible answer.

The three questions every CISO should be able to answer, and currently cannot without a purpose-built AI governance tool:

1. Which AI tools are my employees using right now, including tools IT has not approved?
2. Has any sensitive data, patient records, client files, financial information, source code been sent to a public LLM in the last 30 days?
3. If a regulator asked me to produce an audit log of all AI interactions involving protected data, how long would it take me to produce it?

If the answer to question three is "we do not have that log," the gap is not in policy. It is in infrastructure.

The Lead-up to the Solution

DLP sees files. CASB sees application access. Firewalls see connections. None of them see the prompt. And the prompt is where the data leaves.

Organizations that rely on these tools alone to govern employee AI usage are not governing AI usage. They are governing the infrastructure around it while the exposure happens in plain sight, inside an encrypted browser session, in natural language, at the speed of a keyboard.

The gap requires a tool built specifically for the AI interaction layer, one that operates at the point where the employee types, inspects the content of what is being sent, enforces policy in real time, and produces a verifiable record of every interaction involving sensitive data.

That is what LangProtect Guardia is built to do.

What is LangProtect Guardia?

LangProtect Guardia is a browser-level AI governance platform built specifically for enterprise employee AI interactions. It operates at the one layer that every existing security tool misses, the point between the employee's keyboard and the AI tool's input field, inspecting every prompt, every file, and every response in real time before sensitive data has the opportunity to leave the organization's control.

Guardia does not require network reconfiguration, endpoint agent deployment, or changes to the AI tools employees are already using. It installs at the browser level, becomes active immediately, and operates invisibly for employees who are following policy, surfacing only when a prompt or file contains content that requires intervention.

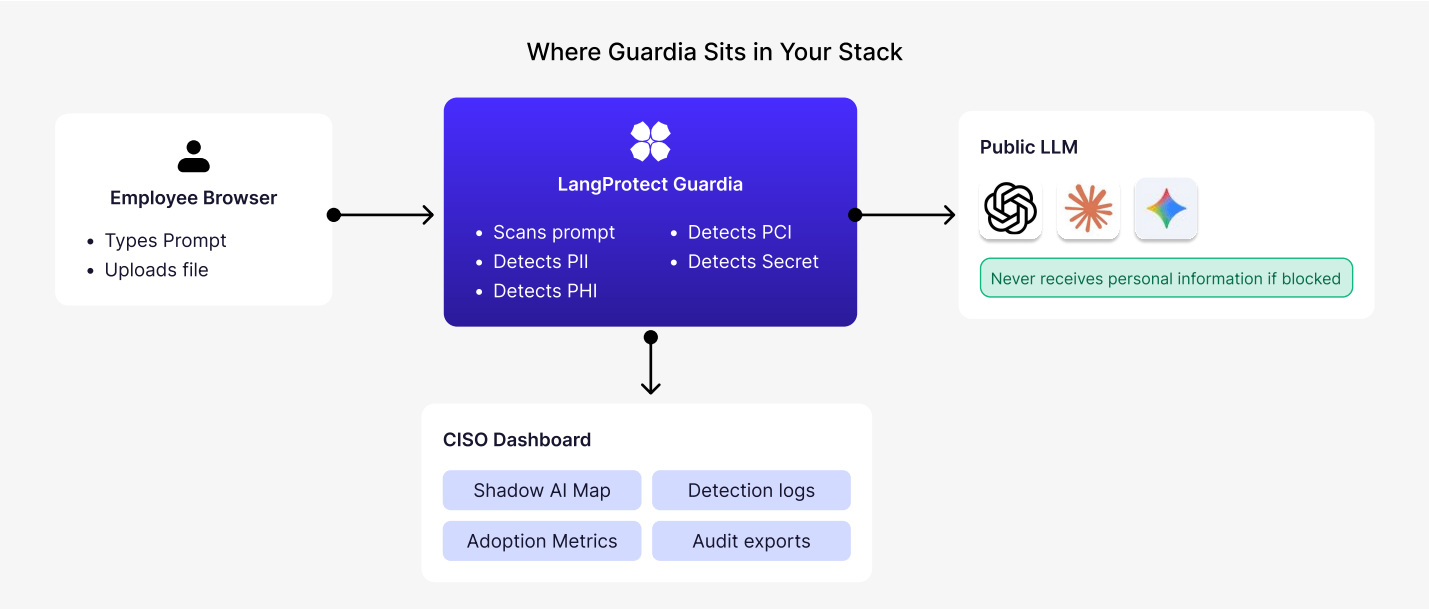

The architecture is intentionally lightweight. Guardia intercepts the interaction at the browser layer, applies its detection and policy engine in real time, and either allows the prompt to proceed, warns the employee, or blocks the transmission entirely, all before the LLM endpoint receives any content. The employee workflow is uninterrupted for compliant interactions. The security team gains complete visibility into every interaction that matters.

Guardia at a Glance - Platform Coverage

Guardia covers thousands of AI applications including ChatGPT, Claude, Gemini, Microsoft Copilot, Perplexity, Mistral, and every emerging AI tool employees discover independently. Coverage extends across all browser-based AI interactions; no tool is outside its scope simply because it is new or unsanctioned.

LangProtect Guardia Dashbaord

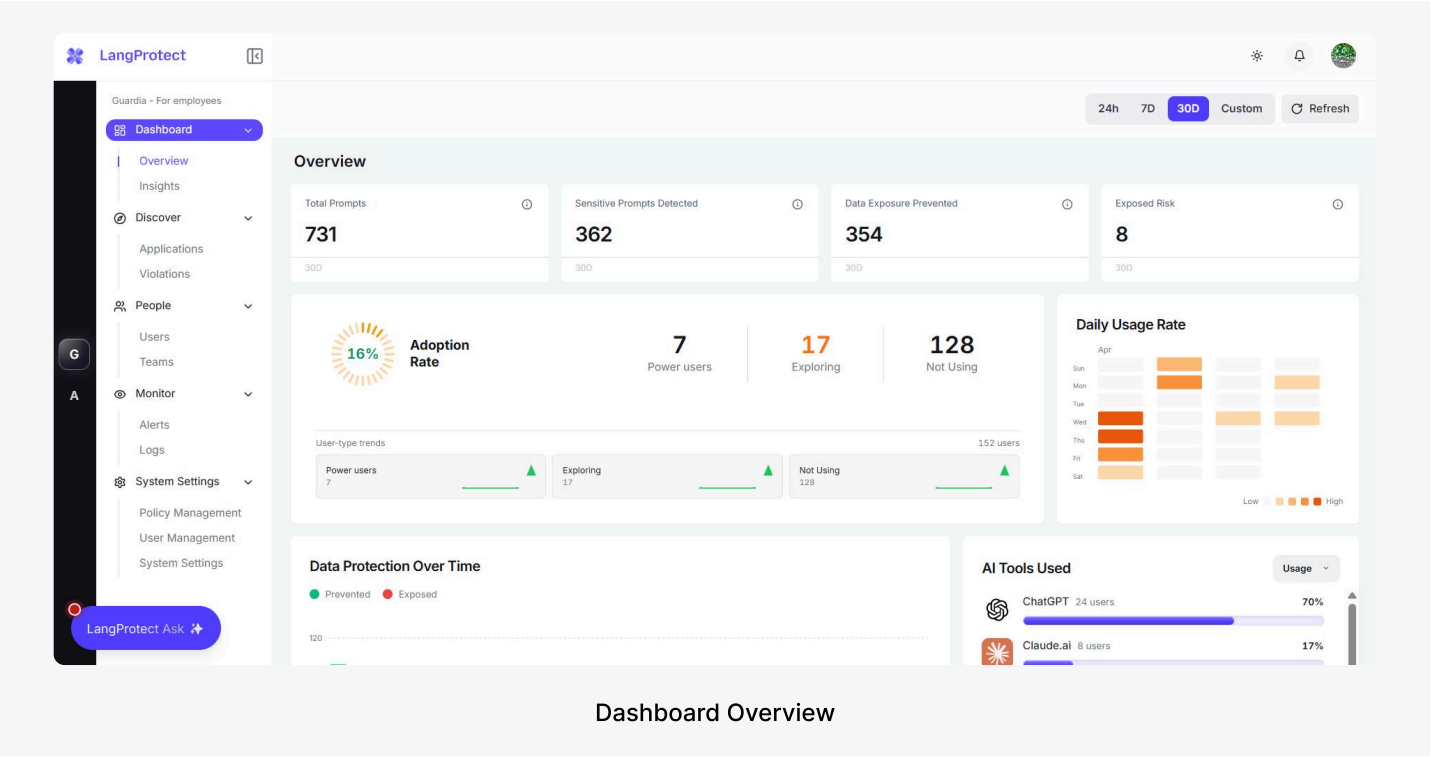

Dashboard Overview

In a single 30-day deployment window, Guardia identified that more than one in three employee AI prompts contained sensitive data the organization had no prior visibility into. The majority were blocked automatically. A fraction proceeded after employee override, every one of them logged, timestamped, and available for compliance review.

This is not an edge case. This is the baseline state of enterprise AI adoption without a purpose-built governance layer.

Four Capabilities That Give You Control

01 Shadow AI Discovery

Find Every AI Tool Your Employees Are Using, including the Ones IT Has Never Seen

What It Does

Shadow AI Discovery gives security and compliance teams a complete, real-time map of every AI application in active use across the organization, including tools that were never submitted for IT review, never approved for use with business data, and never appeared on a CASB catalog or acceptable use policy.

Every AI tool is automatically risk-rated based on its data-handling practices, compliance certifications, geographic data residency, and model-training policies. The dashboard presents this as an actionable inventory: which tools are in use, by how many employees, at what frequency, and at what risk level.

How It Works

Guardia monitors browser-level AI traffic across the organization in real time when an employee accesses any AI application, sanctioned or unsanctioned.

Guardia logs the interaction, identifies the tool, classifies its risk profile, and adds it to the Shadow AI inventory. New tools appear in the dashboard automatically. There is no manual catalog to maintain and no configuration required when a new AI platform enters the market.

What the CISO Sees

The Shadow AI panel displays every identified tool alongside its usage volume, active user count, compliance certification status, and a risk recommendation. For tools classified as high risk, Guardia surfaces a recommended action, including the option to block the application entirely and redirect current users to a lower-risk alternative.

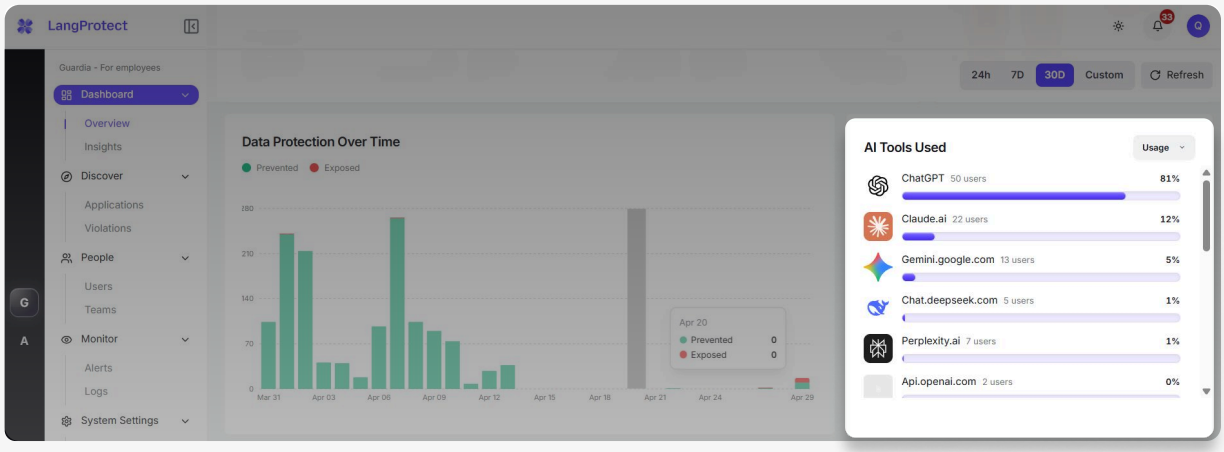

Shadow AI Applications Panel

02 Real-Time Data Protection

Detect and Block Sensitive Data Before It Reaches Any LLM

What It Does

Real-Time Data Protection is Guardia's core enforcement capability. It inspects every prompt and every file an employee attempts to send to an AI tool, identifies sensitive content across a comprehensive range of data classifications, and enforces the organization's data protection policy before the LLM endpoint receives anything.

Detection covers all major sensitive data categories: personally identifiable information (PII), protected health information (PHI), payment card industry data (PCI-DSS), API keys and credentials, proprietary source code, confidential business documents, and custom-defined sensitive content categories specific to the organization's environment.

How It Works

When an employee submits a prompt or uploads a file, Guardia's detection engine analyzes the content in real time against its full scanner suite. Detection happens at the browser layer before the request is transmitted. If sensitive content is identified, Guardia executes the configured policy action:

- **BLOCK** - The prompt or file is stopped. The LLM receives nothing. The employee receives a policy notification explaining what was detected and why the interaction was blocked.
- **WARN** - The employee is alerted that sensitive content has been detected and given the option to proceed or revise the prompt. The decision and outcome are logged regardless of the employee's choice.
- **ALLOW WITH LOGGING** - For lower-sensitivity content categories, the interaction proceeds while being logged for compliance review.
- **REDACT** - Automatically remove or transform sensitive entities before the prompt is sent to the AI. Enabling either redact mode disables Warning and Block Users.

Every action (block, warn, or allow) is recorded with a full interaction log including the user identity, the tool accessed, the content classification, the scanner confidence score, and a timestamp.

What the CISO Sees

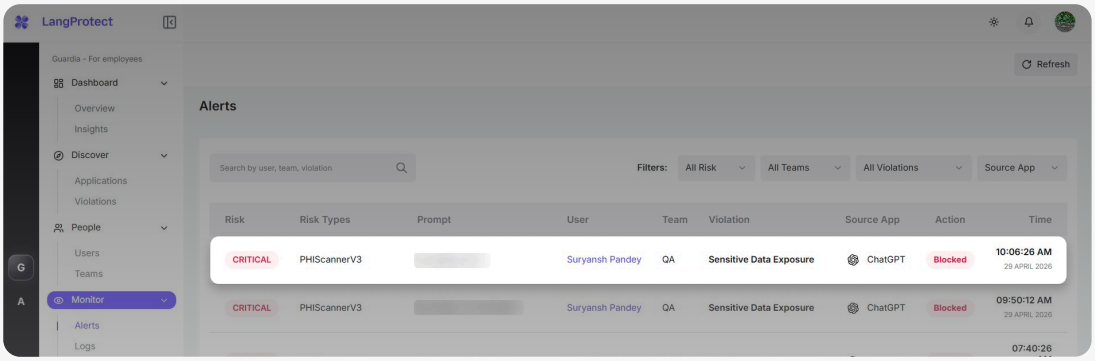

Threat Log View

03 Org-Wide AI Adoption Metrics

Know Whether Your AI Rollout Is Actually Working, and Where the Risk Is Concentrated

What It Does

Most organizations have invested significantly in AI tool procurement, licensing, and employee training, and have no reliable way to measure whether those investments are translating into actual adoption. Guardia's Adoption Metrics dashboard answers the question that no other security tool is positioned to answer: is our AI rollout working, and where is the usage concentrated?

The dashboard segments the entire employee population into three behavioral categories: Power Users, Exploring, and Not using, and tracks movement between categories over time. It maps daily usage patterns, identifies which departments and individuals are driving adoption, and surfaces where high-volume AI usage correlates with elevated data risk.

How It Works

Guardia aggregates interaction data across the full employee population in real time, building individual and group-level usage profiles without capturing personally identifiable prompt content for general analytics purposes. Adoption metrics are calculated at the organizational, departmental, and individual level, with configurable time windows for trend analysis.

The combination of adoption data and risk data in a single dashboard is what distinguishes this capability from basic usage analytics. A department with high AI adoption and high sensitive data detection rates requires a different governance response than a department with high adoption and low risk. Guardia makes that distinction visible and actionable.

What the CISO Sees

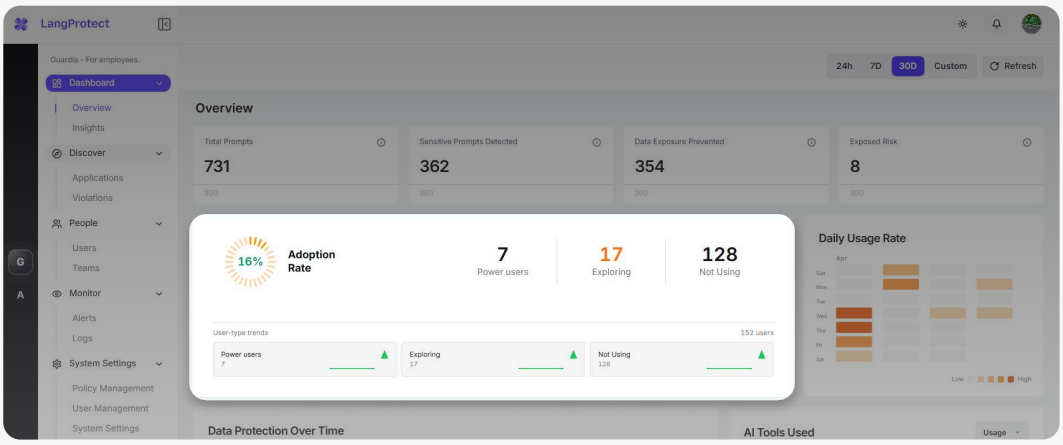

Adoption Rate Panel

WHY ADOPTION METRICS BELONG IN A SECURITY PLATFORM

Security teams are increasingly accountable not just for what AI usage is blocked, but for whether the organization's AI investment is delivering value safely.

Guardia answers both questions from a single dashboard: which employees are using AI, at what volume, with what risk profile, and whether the data protection layer is keeping exposure within acceptable bounds.

The CISO and the Chief AI Officer get the same source of truth.

04 Audit Trail for Regulators

Every AI Interaction Involving Sensitive Data, Logged, Exportable, and Regulator-Ready

What It Does

Guardia maintains a complete, immutable audit log of every AI interaction that involves sensitive data detection, including the user identity, the AI tool accessed, the content classification, the scanner that triggered, the confidence score, the action taken, and the precise timestamp. This log is searchable, filterable, and exportable in formats suitable for regulatory submission, internal compliance review, and legal hold.

For organizations subject to HIPAA, SOC2, GDPR, or PCI-DSS, the audit trail is not a reporting feature, it is the compliance infrastructure. It transforms AI governance from a policy document into a verifiable, defensible record.

How It Works

Every interaction that triggers a Guardia scanner generates an immutable log entry automatically. No configuration is required to activate logging, it is on by default for every detection event. Logs are stored with full interaction context and can be filtered by user, date range, AI tool, data classification, scanner type, and action taken.

Export is available in PDF and CSV formats. A compliance officer preparing for a HIPAA audit, a SOC2 review, or an internal security assessment can pull a filtered, formatted report of every AI interaction involving protected data in under two minutes.

What the CISO Sees

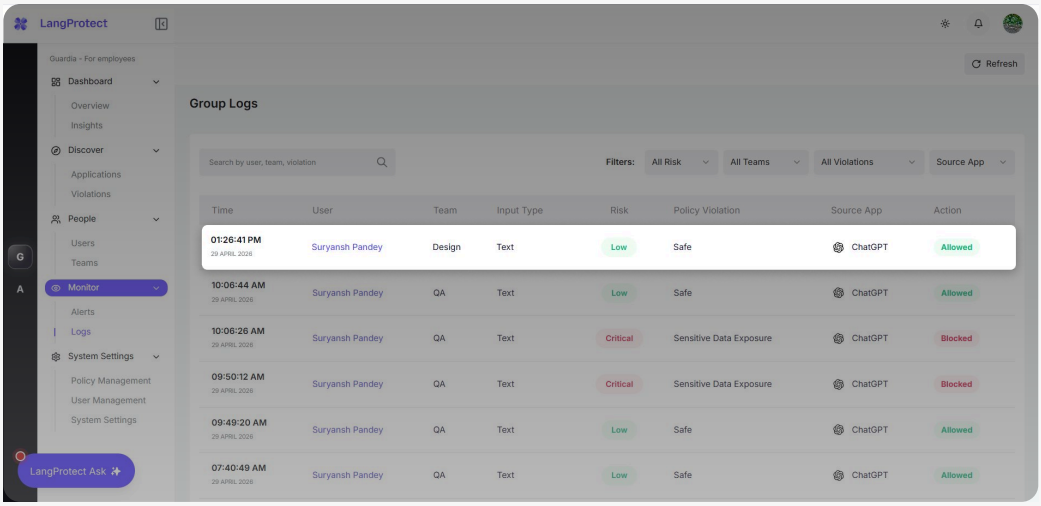

Log Overview

When a regulator asks 'show me every time patient data was sent to an AI tool in the last 90 days,' this is the screen you open.

- Filter by: Healthcare PHI Protection scanner.
- Filter by: Date range.
- Export to PDF.

The answer that used to take weeks to reconstruct, or could not be produced at all, is now a two-minute compliance task.

05 Guardia Insights

38 Live Signals That Turn Enforcement Data Into Governance Intelligence

What It Does

Guardia Insights transforms every detection event, enforcement action, and usage signal across the organization into 38 continuously updated intelligence signals across 15 risk categories, covering everything from individual behavior anomalies and prompt injection attempts to boardroom-ready risk scores and regulatory readiness ratings. It answers the question no log can: what does all of this actually mean for the organization?

How It Works

Insights are generated automatically from live Guardia enforcement and usage data. No configuration required. Every insight card displays a severity rating, a live metric value, and the specific data driving it. Results are filterable by time window: Last 24h, Last 7 Days, Last 30 Days, and by category.

The 15 insight categories cover the full governance picture:

- **AI Visibility & Discovery**: Shadow AI detected, new tools added, approved vs actual usage gap
- **Data Exposure & Security Risk**: % prompts with PII / PHI / source code, top tools receiving sensitive data, team exposure heatmap
- **User & Behavior Intelligence**: Abnormal usage spikes, repeated sensitive prompts by user, power vs casual user segmentation
- **AI Security Threats**: Prompt injection and jailbreak attempts, policy-violating outputs, model vulnerability alerts
- **Governance & Compliance**: Audit-ready logs, policy violations by tool and team, Regulatory Readiness Score
- **Cost & Spend Optimization**: Unused licenses, redundant tools, cost per team, free vs paid tool leakage
- **ROI & Productivity**: High vs low ROI use cases, time saved vs cost incurred, task acceleration signals
- **Adoption & Maturity**: AI Maturity Score, governed vs ungoverned usage %, adoption curve by department
- **Tool & Model Intelligence**: Model preference trends, hallucination rate, model switching behavior
- **Board-Level Risk**: AI Risk Heatmap, Near-Miss & Incident Timeline, AI vs Human Boundary Index

What the CISO Sees

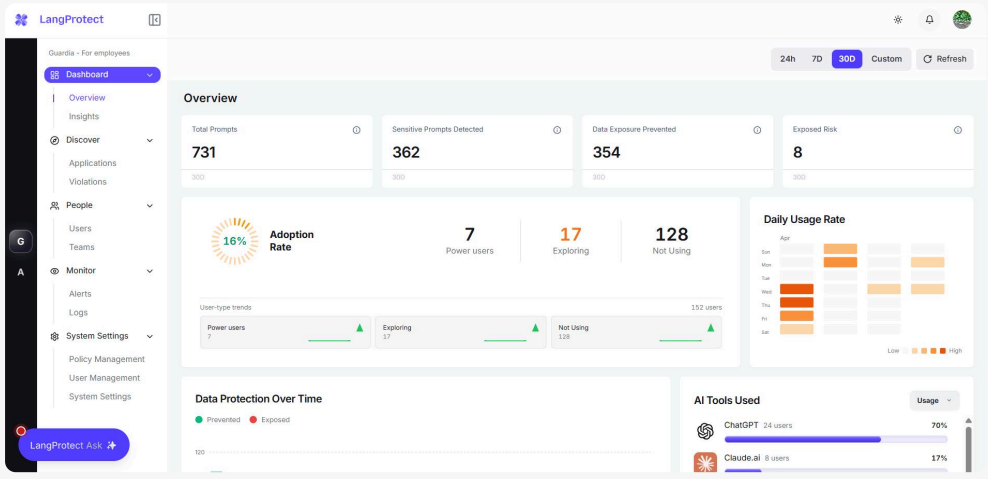

Dashboard Overview

What GUARDIA Looks Like in Practice

Guardia in a 30-Day Enterprise Deployment

The numbers below are drawn directly from LangProtect Guardia platform data. This is what the dashboard shows from day one of deployment.

Day One - The Organisation-Wide Picture

Within hours of deployment, the Guardia dashboard surfaces the first complete view of the organization's actual AI usage; which tools, which employees, what data, and what risk.

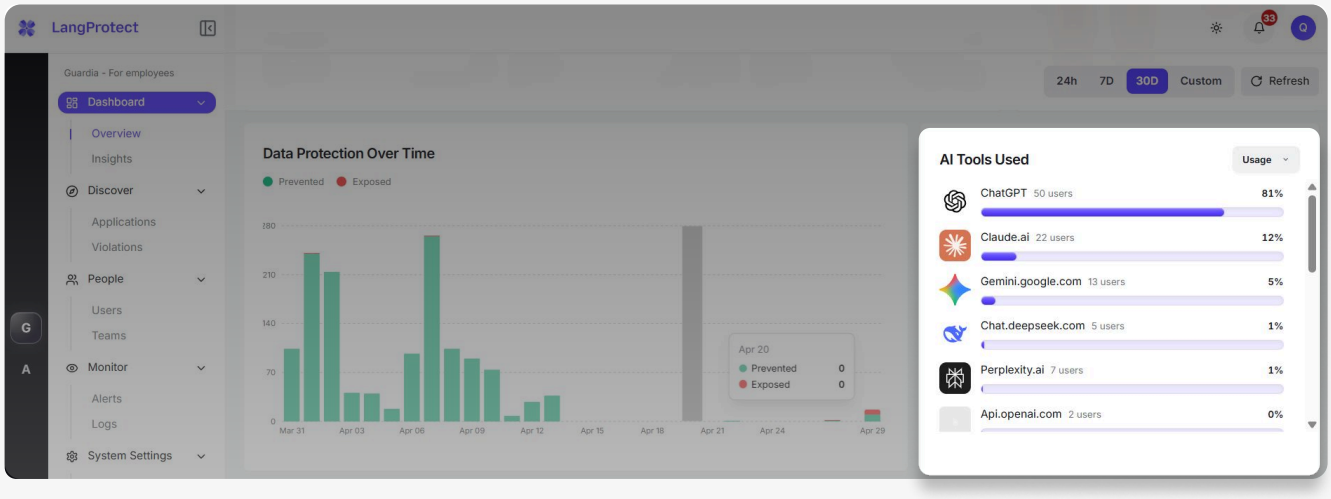

Shadow AI Applications Panel

WHAT THE CISO SEES IN THE FIRST HOUR

90% of all employee AI traffic is flowing through ChatGPT, a platform with no HIPAA BAA, no enterprise data residency guarantee, and no organizational oversight until Guardia was deployed. This is the baseline state of most enterprises today. Guardia makes it visible. Immediately.

Week Two — Drilling Into the Highest-Risk Application

The compliance team drills into ChatGPT.com, classified at Risk Level: Critical. Guardia surfaces the full interaction picture for that single application.

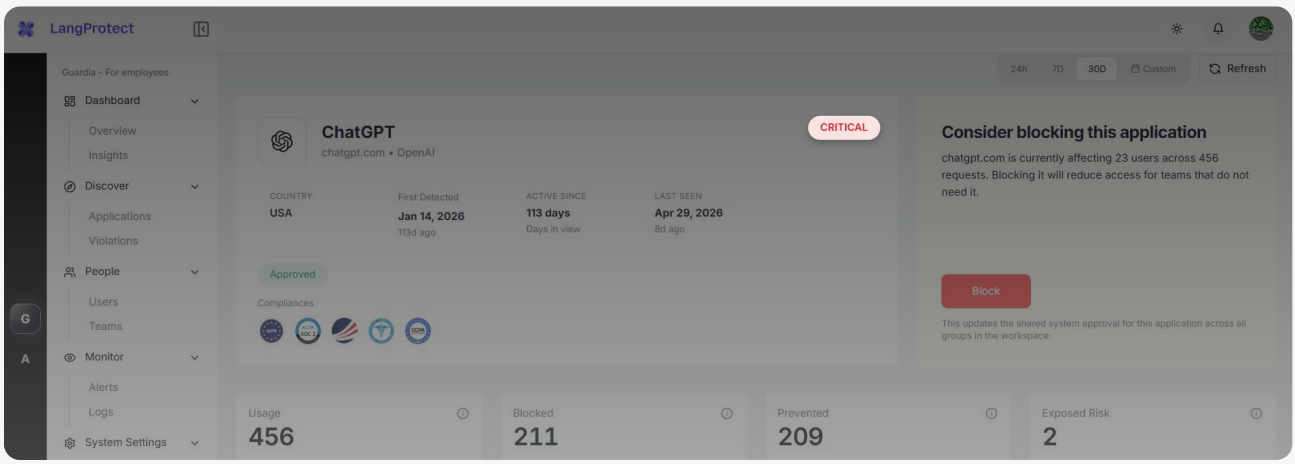

Application Detail View

Week Three — A Detection Event in Detail

Every Guardia detection generates a complete log entry. The following is a representative event from this deployment period.

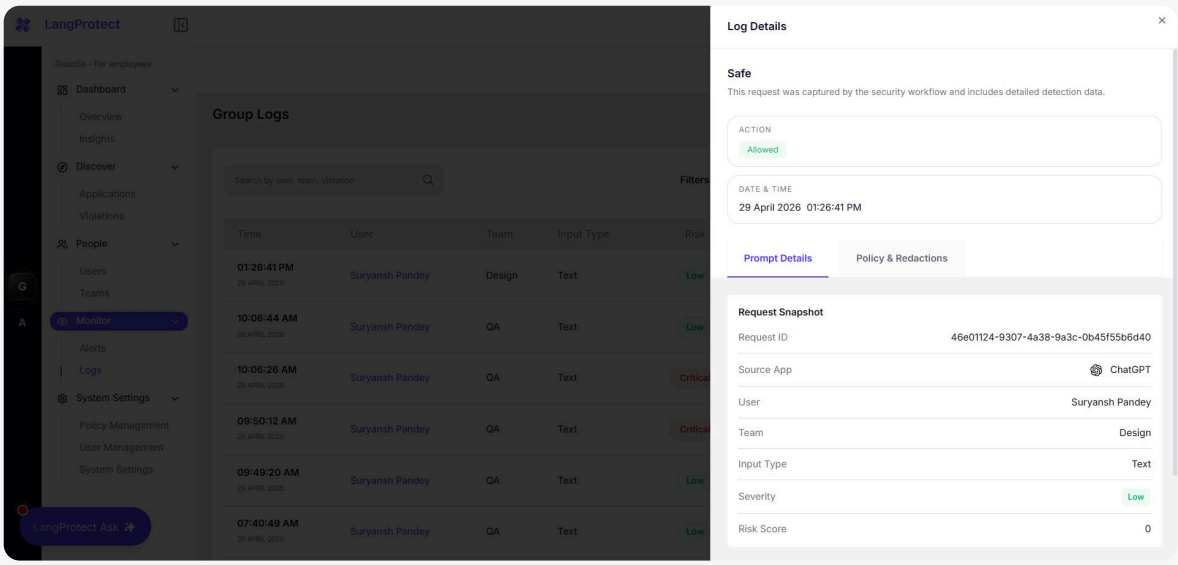

Without Guardia, this interaction completes. A patient record is processed by an external LLM with no HIPAA controls, no audit log, and no organizational awareness.
With Guardia, it is blocked, logged, and available for regulatory review in under 2 minutes.

Compliance Coverage

Guardia does not replace your compliance program. It gives your compliance program the evidence it has always been missing.

The four capabilities covered in Section 4 map directly to the specific requirements of the regulatory frameworks most relevant to US enterprise organizations. The capability map below is not a legal opinion. Every organization should validate applicability with qualified legal and compliance counsel.

GUARDIA CAPABILITY - REGULATORY MAPPING

HIPAA

Key Requirement Covered entities must ensure that protected health information is not disclosed to unauthorized systems or third-party platforms without a signed Business Associate Agreement. Audit controls must be maintained for all ePHI access and transmission events.

Guardia Capability Real-Time Data Protection + Audit Trail

How It Is Met Guardia's Healthcare PHI Protection scanner detects patient names, identifiers, insurance records, clinical data, and financial health information in every employee prompt and file before it reaches any LLM endpoint. Matching interactions are blocked automatically. **Every detection event** — blocked or overridden, is logged with the user identity, timestamp, tool accessed, scanner confidence score, and action taken. The full log is exportable in PDF or CSV format for auditor submission in under two minutes.

SOC2 TYPE II

Key Requirement Organizations must demonstrate logical access controls over systems that handle sensitive data, maintain documented evidence of policy enforcement, and provide auditors with a complete record of how data access is monitored and controlled across the environment.

Guardia Capability Shadow AI Discovery + Audit Trail

How It Is Met Guardia's Shadow AI Discovery dashboard provides a complete, real-time inventory of every AI application accessed by employees, including unsanctioned tools that have never been submitted for IT review. Each tool is risk-rated and logged. The policy enforcement record captures every block, warn, and allow event across the full monitoring period. Both the application inventory and the enforcement log are available for direct submission to SOC2 auditors without additional preparation.

GDPR

Key Requirement Personal data must not be transferred to processors or third-party platforms without adequate contractual and technical controls in place. Organizations must maintain records of processing activities under Article 30 and be able to demonstrate accountability for all personal data handling across their environment.

Guardia Capability Real-Time Data Protection + Audit Trail

How It Is Met Guardia's PII detection scanners identify personal data: names, email addresses, identification numbers, financial data, and other GDPR-defined personal information, in employee prompts before transmission to any external AI platform. Blocked interactions are logged with full detail. The audit export provides a structured record of every PII detection event, the tool involved, the action taken, and the outcome, suitable for Article 30 records of processing activities and regulatory accountability documentation.

PCI-DSS

Key Requirement Cardholder data must be protected at all points of transmission. Organizations must maintain audit logs of all access to systems that handle payment card data and demonstrate that transmission to unauthorized external systems is actively prevented and monitored.

Guardia Capability Real-Time Data Protection + Audit Trail

How It Is Met Guardia's Financial DLP Protection scanner detects payment card numbers, account data, and associated financial identifiers in employee prompts and files. Interactions containing cardholder data are blocked before reaching any external LLM. Every financial data detection event is logged with full interaction detail: user, tool, timestamp, scanner, confidence score, and action, and is available for PCI audit submission in exportable format.

Compliance Readiness — What Guardia Produces For Each Framework

Framework	Document Guardia Can Produce	Time to Generate
HIPAA	PHI detection log: all events, all users, filtered by date range	GPT-4o, GPT-4 Turbo, GPT-3.5 Turbo
HIPAA	Blocked transmission record with scanner confidence scores	Claude 3.5 Sonnet, Claude 3 Opus, Claude 3 Haiku
SOC2 Type II	AI application inventory with risk classifications	Llama 3, Llama 2
SOC2 Type II	Policy enforcement log: all block and warn events	Gemini 1.5 Pro, Gemini 1.5 Flash
GDPR	PII interaction log: detections, blocks, and exposures	Mistral Large, Mistral 7B
GDPR	Data flow summary by AI tool and data classification	Mistral Large, Mistral 7B
PCI-DSS	Financial data detection log with action taken per event	Custom and self-hosted models

Guardia is not a compliance certification and does not constitute legal advice on regulatory requirements. It is a technical control layer that produces the detection records, enforcement logs, and audit evidence that compliance programs require.

The difference between an organization that can answer a HIPAA audit question in two minutes and one that cannot is not the quality of their policy documentation. It is whether the underlying infrastructure captures the evidence at the moment the interaction occurs.

Guardia captures it. Every time.

Book a LangProtect Guardia Demo

See the Shadow AI discovery dashboard, the real-time detection engine, and the compliance audit trail with your organization's use case as the starting point.

BOOK A DEMO >

SEE FULL CAPABILITIES >

ABOUT LANGPROTECT

LangProtect is an enterprise AI security platform built to secure the full spectrum of organizational AI activity across employees, applications, and autonomous agents.

Guardia	Armor	Vector
Latency Employee AI interactions	Latency AI applications and LLM APIs	Latency AI agents and MCP connections
Primary Use Case Shadow AI discovery, data leakage prevention, compliance audit trail	Primary Use Case Prompt injection defense, PII redaction, runtime policy enforcement	Primary Use Case Agentic workflow governance, unauthorized tool call prevention

LangProtect deploys at the browser and API layer, no network reconfiguration, no endpoint agent complexity, no disruption to existing workflows. Detection runs in under 50ms. Audit logs are available from the moment of deployment.