Protect Your Data from AI Agents Running in the Wild

Policy-Based Access Control for Agentic AI

Prevent AI Agents from Exposing Data

Agentic AI systems introduce a new level of operational power—and risk. They connect to enterprise systems, retrieve data, and generate outcomes on behalf of users—**often without built-in controls.**

These agents can access sensetive Information such as MNPI, financial forecasts, M&A plans, or clinical results, that spans multiple applications, models, and data layers.

Without unified policy enforcement across the AI flow, they can unintentionally combine or expose sensitive data, breaching compliance barriers and increasing insider risk.

The result: growing compliance exposure, audit gaps, and loss of control as AI adoption scales across the enterprise.

The growing attack surface

- **Unauthorized Access:** Agentic AI acts on behalf of users or systems with insufficient enforcement of contextual identity and permissions, leading to unauthorized access to sensitive data or restricted actions.
- **Data Exposure:** Gaps in access control can lead to regulatory and compliance failures, increasing the risk of fines, legal exposure, and lasting reputational damage.
- **Lack of Auditability:** Multi-step reasoning, external calls (APIs, Tools, Data), make it difficult to trace and verify what services and data were accessed.

Why Policy Management for Agentic AI?

Build AI With Identity-First Security
Ensure every AI action is bound to the user's identity and entitlements. Policy enforcement is tightly integrated with identity context, ensuring that AI agents act only within the permissions of the user or system they represent.

Minimize Risk with Dynamic, Context-Aware Enforcement
Reduce data exposure, misuse, and unauthorized actions through real-time controls. Policies adapt to runtime context, such as user intent, data sensitivity, and agent behavior, mitigating threats like privilege escalation and tool misuse.

Centralize Control Across the Full AI Workflow
Manage access decisions at every stage: prompt, data retrieval, generation, and response. A unified policy engine governs how data is queried, combined, and shared across applications and data sources, maintaining consistent and secure behavior.

The only enterprise-ready AI access policy management solution readable in plain business language.
Security teams see the code. Auditors see the logic. Comprehensive logging ensures transparent, verifiable data access control across AI environments.

Policy Management for Agentic AI: The Enterprise-Grade End-to-End Solution

DISCOVER → **MANAGE** → **AUTHORIZE**

DISCOVER
Discover the tools and data sources to protect

MANAGE
Define policies to control access throughout the AI flow

AUTHORIZE
Build Agentic AI systems with integrated security controls

Control the Entire Agentic AI Flow with an Intuitive Policy Builder

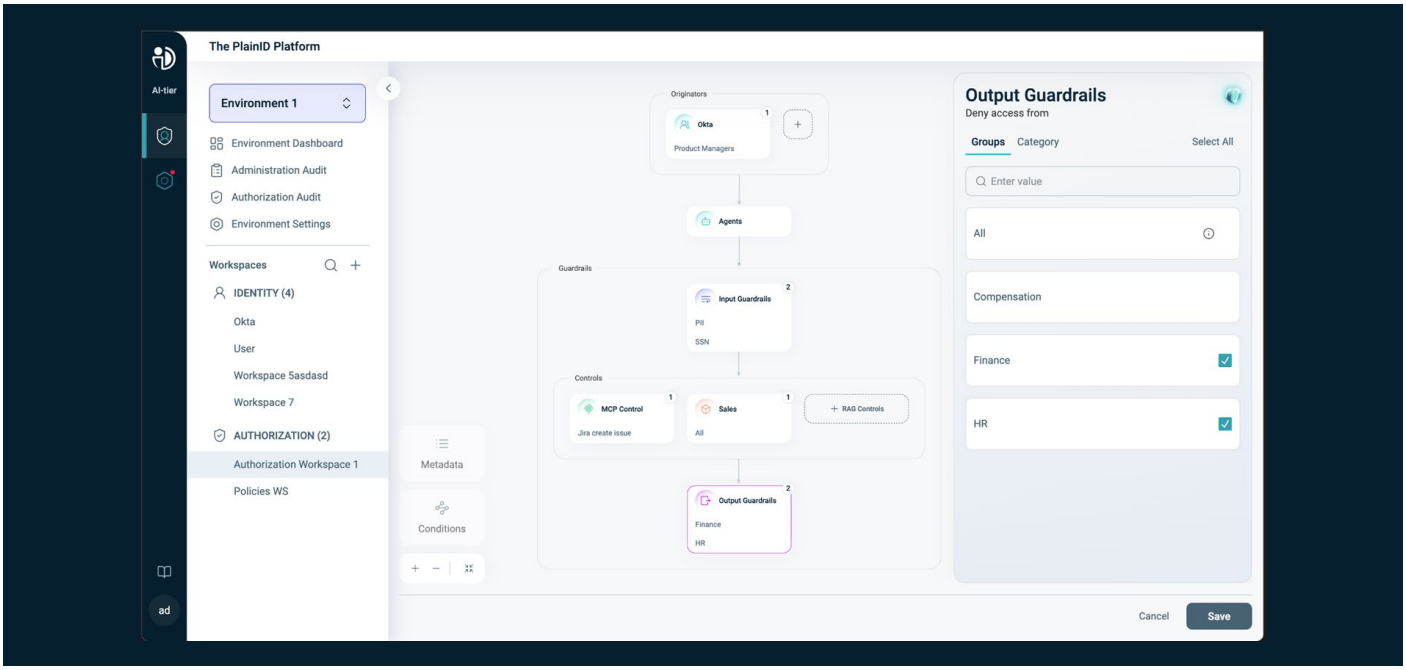

- **Control the Prompt**
  Security should always be implemented at the nearest gate possible, if the user is not authorized for the topic, why start the process? Prevent unauthorized attempts to extract sensitive data and reduce exposure risks before retrieval.
- **Control the Data**
  Apply filters to data retrieval, unstructured and structured, based on authorized topics and user data. Determine who can access what and when in real time based on the identity and context of access. Prevent retrieval of unauthorized documents to ensure access only to authorized and relevant content.
- **Control the Tools**
  With MCP it's easier for agents to access and utilize services and tools. Control access to those tools based on the human and agent identity provide the right context to the access.
- **Control the Response**
  Mask and filter data from generated responses by ensuring that AI-generated responses displayed to users align with their permissions. This prevents the LLM from exposing unauthorized insights, keeping response output secure, compliant, and controlled.

A Simple, Guided Policy Experience

AI that does the heavy lifting

- **Natural-Language Policy Creation:** Define and refine access rules in plain language.
- **AI Recommendations & Data Mapping:** Automate connections across datasets.
- **Intuitive Canvas Interface:** Visualize, auto-fill and adjust policies effortlessly.

Technical Insight

Identity-aware Access Control Prevents AI Overreach

AI Agents function as NHIs, but what data they can access should be tied to the user they serve, ensuring permissions align.

Limiting AI access based on the user's identity (e.g. role, attributes) ensures NHIs never exceed intended permissions.

About PlainID

PlainID is the world's leading provider of enterprise Authorization, helping enterprises address the complex challenges of Identity Security. The PlainID Platform allows you to discover, manage, and authorize access control policies for enterprise applications and data. Our solution is architected to protect against identity-centric security threats powered by Policy-Based Access Control (PBAC). Visit PlainID.com for more information.