Control AI Agents. Control What They Access, Do, and Expose

Continuous, Real-time Authorization of Agentic AI

End-to-end solution
Discover, manage, authorize

Enforce all identities
Human and non-human

Scale AI responsibly
With zero standing privilege

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.

Without unified policy enforcement across the AI flow, agents function with "standing privileges," retaining access that can be reused or escalated, can unintentionally combine or expose sensitive data.

The result: potential leakage of sensitive data, such as MNPI or financial forecasts, can lead to 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?

Static controls cannot keep pace as agentic AI operates continuously at machine speed across the enterprise environment. Because context shifts with every request, fixed permissions and standing privileges create immediate security gaps.

- **Build AI with Identity-First Security**
  Ensure every AI action—whether by a human or a Non-Human Identity (NHI)—is strictly bound to permissions.
- **Minimize Risk with Dynamic Enforcement**
  Policies adapt to runtime context, mitigating threats like privilege escalation and tool misuse in real-time.
- **Centralize Control Across the Full AI Workflow**
  Manage access decisions at every stage: prompt, data retrieval, generation, and response

> Through 2029, over 50% of successful cybersecurity attacks against AI agents will exploit access control issues.

Gartner, How to Secure Custom-Built AI Agents, Dionisio Zumerle, Jeremy D'Hoinne, 11 June 2025 GARTNER is a registered trade mark and service mark of Gartner, Inc. and/ or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

PlainID is Trusted by Fortune 2000 Companies to scale AI and secure access for millions of identities

Pre-retrieval Filters

Enforce granular filters on RAG workflows before data is retrieved or exposed

Granular MCP Tools Control

Govern exactly which MCP tools, parameters, and data agents can access

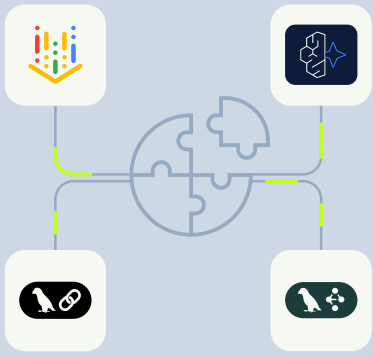

Flexible Integration Model

Embeds natively into leading AI and agent development frameworks

The only end-to-end authorization platform that governs access across the full agentic AI flow at scale

- **Input Guardrail**
  Enforce prompt authorization so agents only ask within approved scope. Block attempts to extract sensitive data before retrieval.
- **Data Retrieval Guardrail**
  Apply policy-based guardrails to manage who can access which data, in real time, by identity and context. Prevent retrieval of unauthorized documents.
- **MCP Tools Guardrail**
  Govern which services and tools agents can invoke. Context-aware access ensures only the right identities can use the right tools.
- **Output Guardrail**
  Mask and filter generated responses to avoid exposing sensitive or unauthorized insights. Keep output secure, compliant, and controlled.

Technical Insight: Applying Intent-Based access control with dynamic guardrails

Intent is the why access is made, Zero standing privileges (ZSP) is how it is enforced.

An agent with permanent access isn't autonomous; it's unsupervised. PlainID enforces Zero Standing Privileges by granting access Just-in-Time (JIT), strictly scoped to the agent's immediate intent. This ensures that even if an agent is compromised, it never holds "always-on" keys to sensitive data.

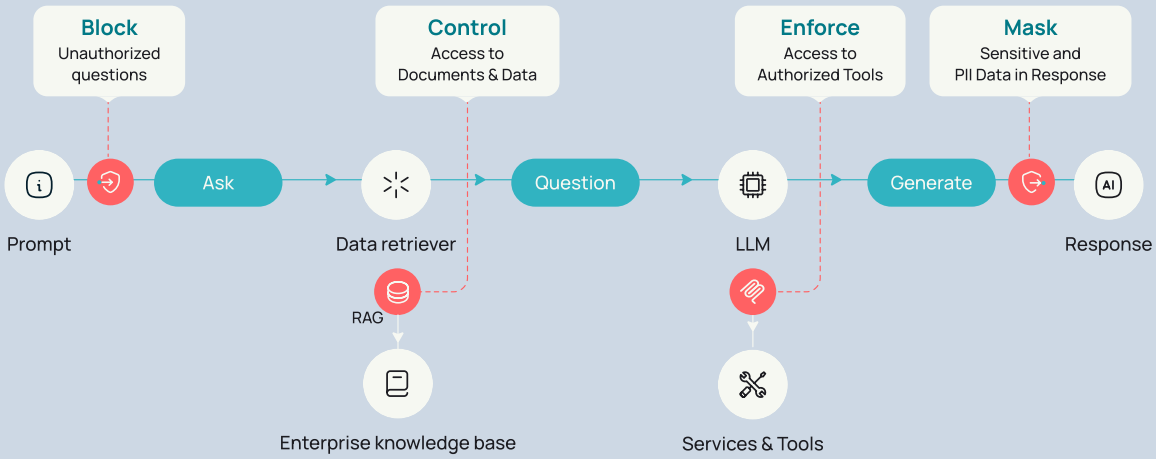

Why Enterprise Leaders Choose PlainID vs. Point Solutions

Capability area	Generic security tools/ AI security point solutions	plainID
Coverage Scope	Partial, point solutions	True end-to-end control across the AI flow: Prompt → Data → Tools → Output
Authorization Lifecycle	Fragmented or incomplete	Full lifecycle: Discover → Manage → Authorize
Decision Enforcement	Mostly reactive	Real-time, dynamic authorization. Driven by context and intent
Identity & Accountability	Limited visibility	Every action tied to human & non-human identities, with clear ownership and scope
Prevention Model	After-the-fact response	Proactive enforcement before data is retrieved or exposed, with built-in output masking
Audit & Explainability	Raw or technical logs	The only solution with business-readable authorization decisions and full audit trail

Beyond Access Controls: Enterprise-Grade Capabilities

PlainID provides a comprehensive framework to operationalize responsible AI. It enables you to discover tools and data to protect, manage policies centrally, and authorize actions to all identities across the environment, in real time.

Discover	Manage	Authorize
Data Source Visibility Gain insight into what data and APIs your agents are attempting to access.	**Identity and Agent aware controls** Manage access to Humans and Agents and how they interact.	**Control for all Identities** Human, NHI and Agent Identity are considered in the enforcement flow.
RAG Discovery & Classifier Connects to Vector DBs (e.g., Pinecone) to list and enrich existing classifications and metadata for policy use.	**Investigation View** Graph-based visibility into any object in the agentic flow and all its related identities, datasets and tools across the enterprise technology stack.	**Guardrails Across the full AI Flow** Enforcement across input and output guardrails, data retrieval controls, and tool invocation control.
Tools Discovery & Classifier Automatically discovers the MCP tools that the organization intends to use and classifies them according to categories of usage.	**AI-led, guided policy experience** Describe intent in natural language; our no-code Policy Builder interface interprets and automates policy creation, pulling the right "building blocks" and suggests related actions and components.	**Flexible Integration Model** Works across all major and emerging AI development and agent frameworks, with its enforcement module fully compatible with the enterprise's technology stack.
	Full audit and visibility Gain visibility into access decision logs with reasoning data for audits and regulators.	**Built-in Output Masking** Derived from PlainID's proven data authorization expertise.
		Granular Tools Control Enforce which MCP servers and tools can be used by the agentic flow, extend control to the data utilized by the tool.
		Beyond Access Controls Granular data-level filters and constraints, including parameter governance, applied before data is retrieved or exposed.

Secure AI access now

Scale your agentic AI responsibly with the market-leading authorization platform

Request a demo

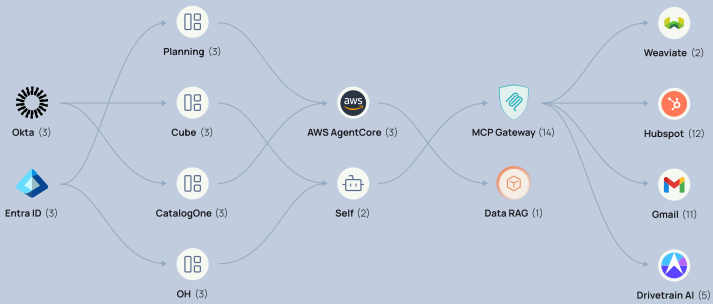

About PlainID

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.