# Policy Management for Agentic AI

Identity-aware Access Controls for AI Systems.

## Now is the time to adopt AI **responsibly**

Agentic AI systems present a new frontier of operational power. They act autonomously, accessing vast amounts of sensitive data, invoking tools, and delivering outcomes on behalf of users. Without purpose-built controls, they pose serious security, compliance, and governance risks.

**Misalignment between engineering and security teams delays AI deployment and increases governance friction.**

The growing attack surface:

- **Privilege Escalation** - 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** - Regulatory & Compliance Failures due to gaps in access control increase the risk of non-compliance, fines, and legal consequences. Additionally, Breaches erode trust and brand value leading to reputation Damage.
- **Lack of Auditability** - Multi-step reasoning, external calls (APIs, Tools, Data), make it difficult to trace and verify what services and data was accessed.

**Policy Management, reduce risk, prevent sensitive data exposure, ensure explainability, and meet compliance obligations in AI-driven environments**

## Why Policy Management for Agentic AI?

**Build AI with 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 the entire AI flow, from the questions users ask to the services and data they access, enabling consistent and secure behavior.

**Enable Auditing and Observability**
Gain full visibility into agent behavior and policy decisions. Comprehensive logging and policy event tracking provide the transparency needed for compliance, explainability, and forensic investigation in hybrid or distributed AI environments.

## The path to Responsible Agentic AI

**DISCOVER** → **MANAGE** → **AUTHORIZE**

Connect and map the data sources

Define policies to control access throughout the AI flow

Build Agentic AI systems with integrated security controls

Ready to use libraries for:

Integration hub

Ready to use Authorizers for controlling access to APIs and Data

snowflake, Google BigQuery, SQL Server, okta, Microsoft Entra ID, envoy, apigee, Azure API Management, spring, SDK, Open Policy Agent, databricks, PostgreSQL, amazon REDSHIFT, Auth0, PingIdentity, Kong, Istio, Amazon API Gateway, MuleSoft, WSO2, OpenID

## Control the Agentic AI flow with Policy Management

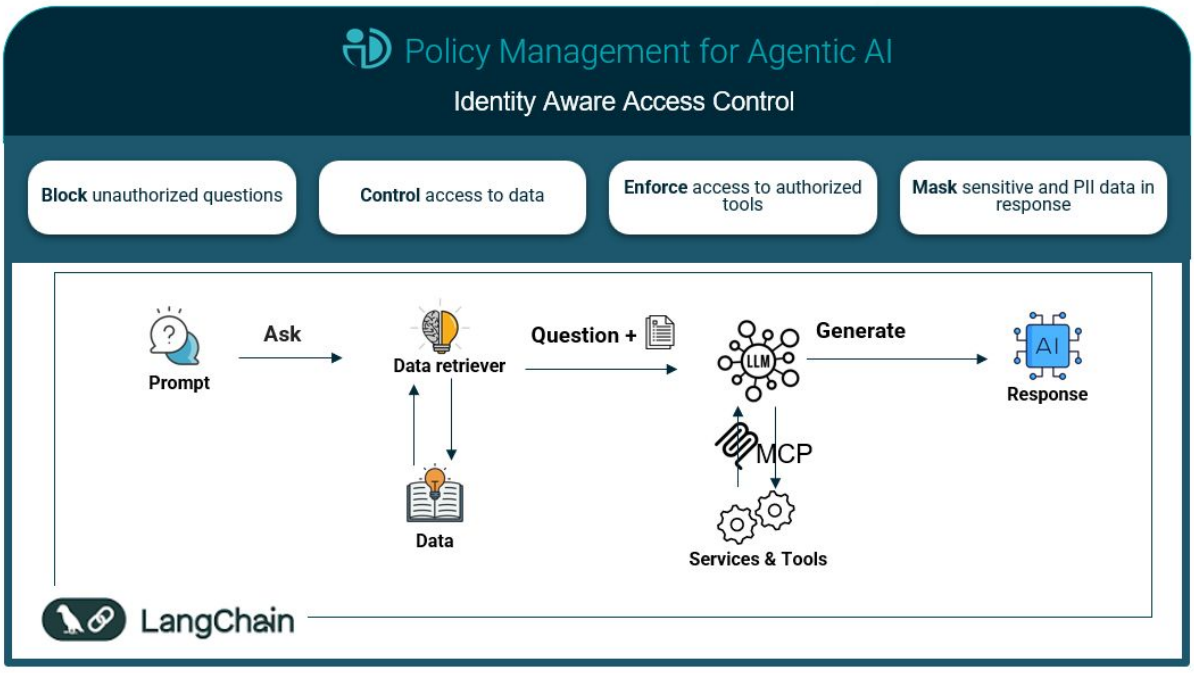


- **Control the prompt**
  Security should always be implemented at the nearest gate possible, if the the user is not authorized for the topic, why start the process? Prevent unauthorized attempts to extract sensitive data and reduce exposure risks before retrieval
- **Control the data**
  Apply filters 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 document to ensure security and provide focus on authorize and relevant content.
- **Control the tools** (beta)
  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.

### Technical Insight

**Identity-aware Access Control Prevents AI Overreach**

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

Limiting AI access based on the user's identity (e.g. role, attributes, etc.) ensures NHIs do not 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.