Checklist: Choosing the Right Policy Management Platform for your Agentic AI Enterprise Systems

As agentic AI systems are granted more autonomy, they introduce significant security risks like data exposure and privilege escalation. Traditional, role-based access controls were not designed for these dynamic, multi-step workflows, creating critical security and compliance gaps. The foundation for controlling these risks is an enterprise-grade, purpose-built access control platform.

This checklist outlines the critical capabilities required to secure your AI workflows – from prompt to response.

Business Capabilities

These capabilities focus on the governance, risk, and compliance aspects essential for CISOs, CIOs, and business leaders.

Enterprise-Grade Scalability:
Handle the performance and complexity demands of large-scale enterprise environments.

Delegated Management & SoD:
Distribute policy management across departments while enforcing separation of duties.

Visibility & Auditability:
Full traceability of the entire agentic AI flow and interactions to satisfy compliance and governance needs.

Lifecycle Management:
Versioning, testing, and controlled deployment of policies from sandbox to production.

Unified Policy Management for Business & Technical Teams:
Enable non-technical stakeholders to define and approve policies using a business-readable language, while providing developers Policy-as-Code and low-code framework support.

Policy Simulation & Investigation:
Simulate "what-if" scenarios to analyze how changes in user or resource attributes will affect an AI agent's access rights, and investigate the full reasoning behind any authorization decision.

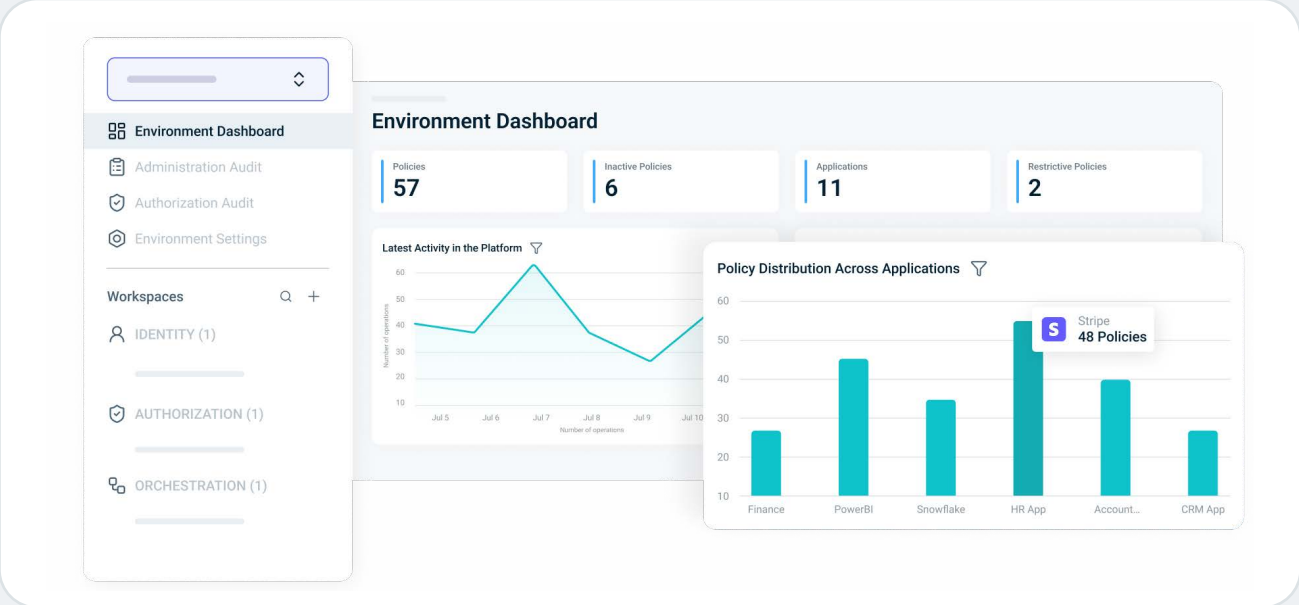

Technical Capabilities

These capabilities focus on the specific security controls, architecture, and enforcement mechanisms critical for Security Architects and AI Developers.

- **Control Access Across the Full AI Flow Using One Platform:** Govern controls at the earliest possible point and across all control points—from prompt and data retrieval to tool usage and response generation.
- **Identity-Aware Enforcement for All Actors:** Tie every action within the AI workflow to a verified identity and its entitlements, with full context for both human users and non-human AI agents (NHIs) to build a Zero Trust foundation.
- **Dynamic Decisioning:** Real-time policy evaluation that factors in identity, attributes, context, and events — no reliance on static permissions.
- **Authorization Sandbox:** Build, test, and simulate policies for AI agents in an isolated sandbox environment before deploying them to production.
- **Connectors for Identity Sources:** Integrate with multiple existing identity sources, including LDAP, SQL, SCIM, and IAM/IGA systems, to leverage your current identity infrastructure for AI authorization.
- **Multi-Level Granular Permissions:** Implement least-privilege access with controls ranging from coarse-grained to fine-grained, ensuring users and AI agents can only access the precise data and tools required for their task.
- **Context-Based Conditions:** Enforce policies based on a rich set of conditions, including time of day, date, user location (source IP), and authentication method, to add layers of dynamic control to AI agent actions.
- **Proactive Data Filtering:** Prevent data exposure by filtering sensitive structured and unstructured data before it is retrieved by the AI Agent.
- **Universal Enforcement:** Consistent controls across AI pipelines, APIs, microservices, data platforms, and application layers.
- **Dynamic Response Masking:** Automatically redact or mask sensitive information in AI-generated responses in real-time based on entitlements, preventing data leakage through model outputs.

Securing Your AI-Powered Future

Choosing the right policy management platform is a strategic decision to enable responsible, secure, and compliant AI adoption at scale. By ensuring your solution provides the comprehensive business and technical capabilities outlined in this checklist, your organization can confidently innovate while maintaining a robust governance and security posture.

Ready to build your AI systems on a foundation of Zero Trust?

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.**

plainID