
# The Market-Leading Runtime Authorization Platform

Unified Policy Fabric, delivering **Zero Standing Privilege (ZSP)** at scale.

# Trusted by Global Fortune 500 Enterprises

PlainID secures mission-critical access for the world's most demanding organizations across banking, retail, technology, and manufacturing sectors.


**35M+**
Identities Managed

**100K+**
Data sources protected

**<1 MIN**
Policy Update

**100M+**
Authorization Decision per Day

# The Security Decision Paradigm Shift - From Static to Intent

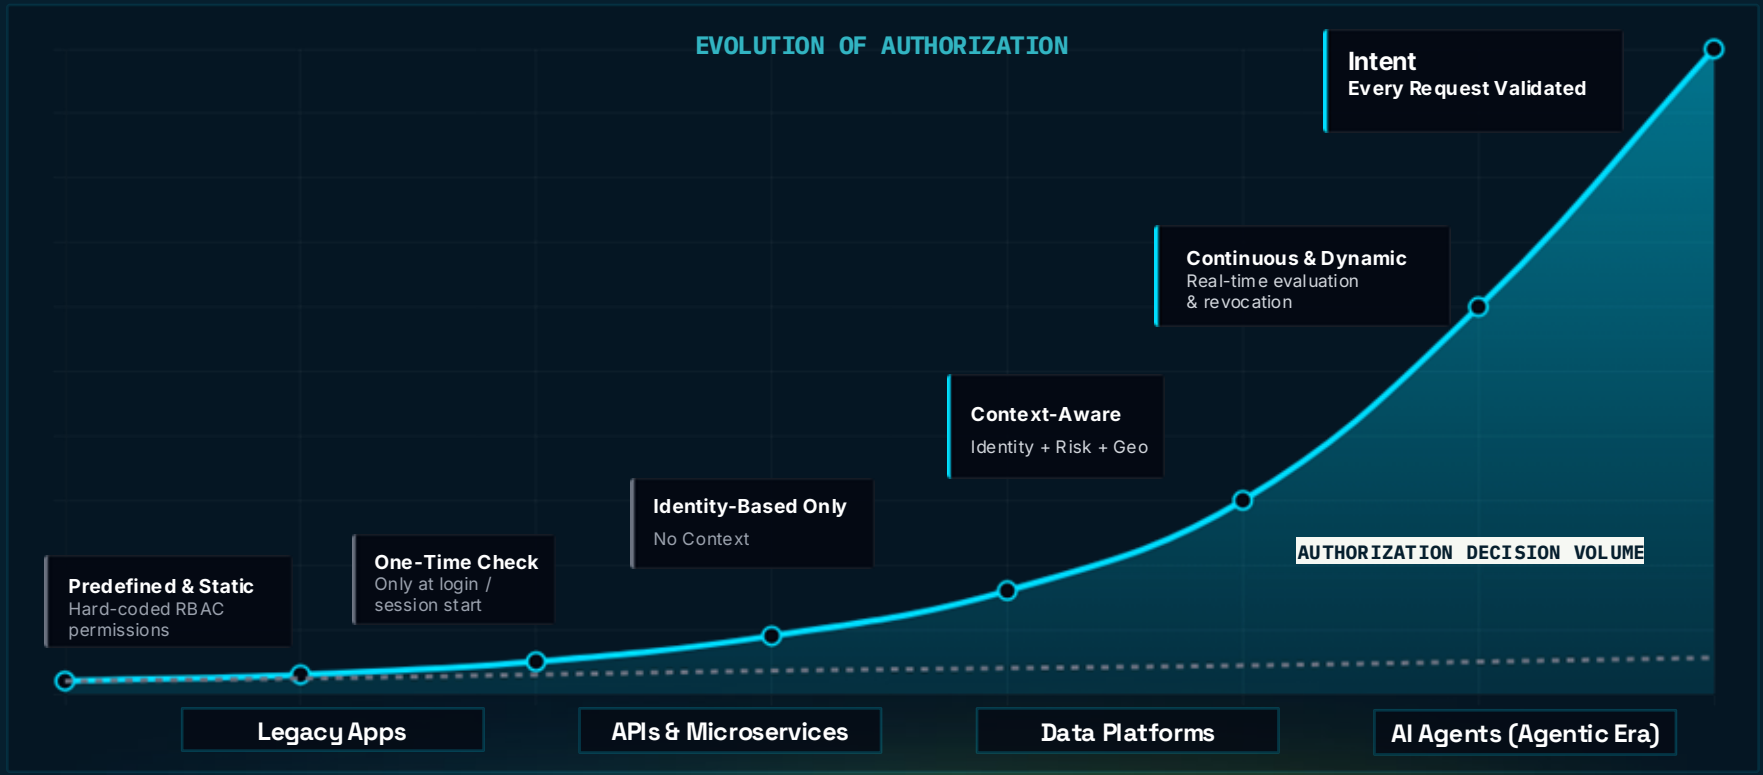

# Evolving Identity Threat in the AI Era

For applications, APIs, Data and AI agents

## Identity Landscape Exploded

Machine identities outnumber humans 109:1. APIs, microservices, copilots - all acting as identities

## Standing Privileges Are the Attack Surface

Over-privileged identities expanding blast radius. IAM knows who you are. Nothing controls what you're allowed to do

## Sensitive Data Exposure

User, apps and AI agents ingest and process vast amounts of enterprise data.

## Visibility≠ Control

Modern security requires continuous, real-time enforcement of what every identity can do.

# DISCOVER → MANAGE → ENFORCE

## The PlainID Runtime Platform

VISIBILITY | CONTROL | CONSISTENCY | STANDARDIZATION

### Central Policy Management

- Application Access Control
- API & Microservices Access Control
- Data Access Control
- Agentic AI Access Control

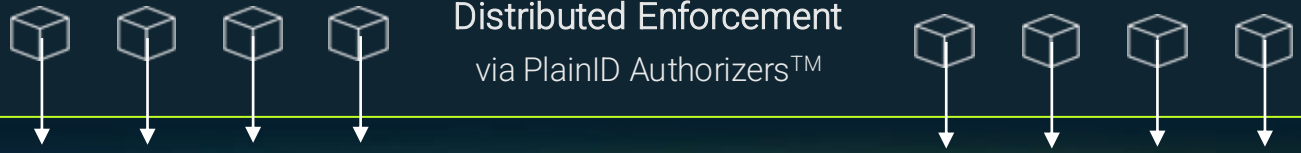

Distributed Enforcement

via PlainID Authorizers™

# PlainID Technology Ecosystem

Connects technologies, applications, and microservices with consistent policy enforcement across the enterprise , on-premise, cloud, or distributed architectures.

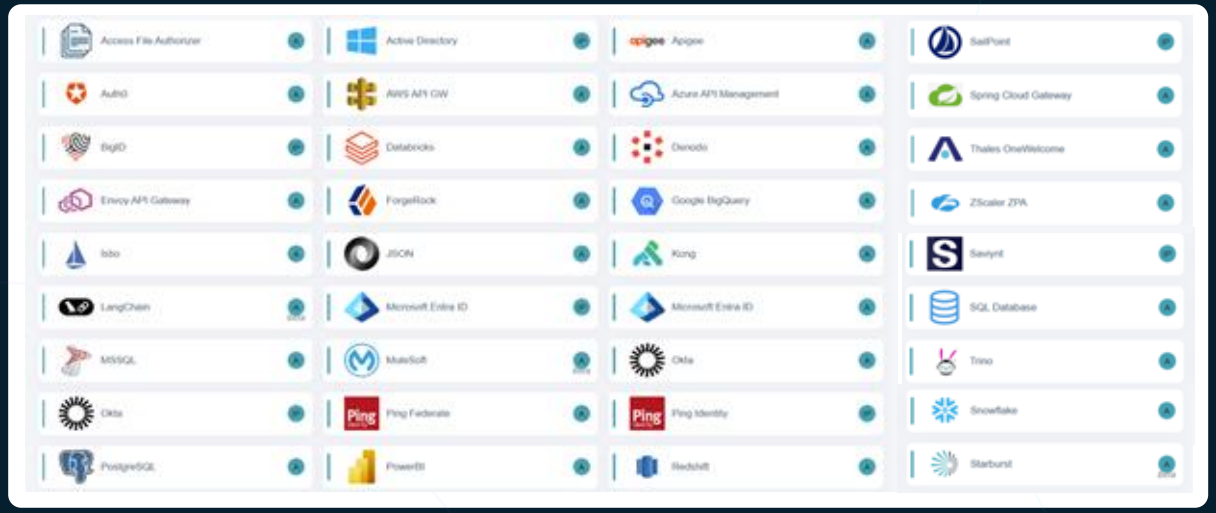


Leading standards

OpenID

OIDC

OpenID

AuthZEN

Open Policy Agent

REGO

SCIM

MCP tools

**Out-of-the-Box Integrations**
Pre-Built Connectivity

**Flexible SDK**
Developer First

**Accelerated Deployment**
Time to Value

# How it works?

## Who can do what (and on what) and when?

# Discover, Manage, & Enforce

## Agentic AI

**Discover** and classify data and MCP tools

**Manage** policies with a guided experience, supported with an AI assistant

**Enforce** runtime authorization across the full agentic flow

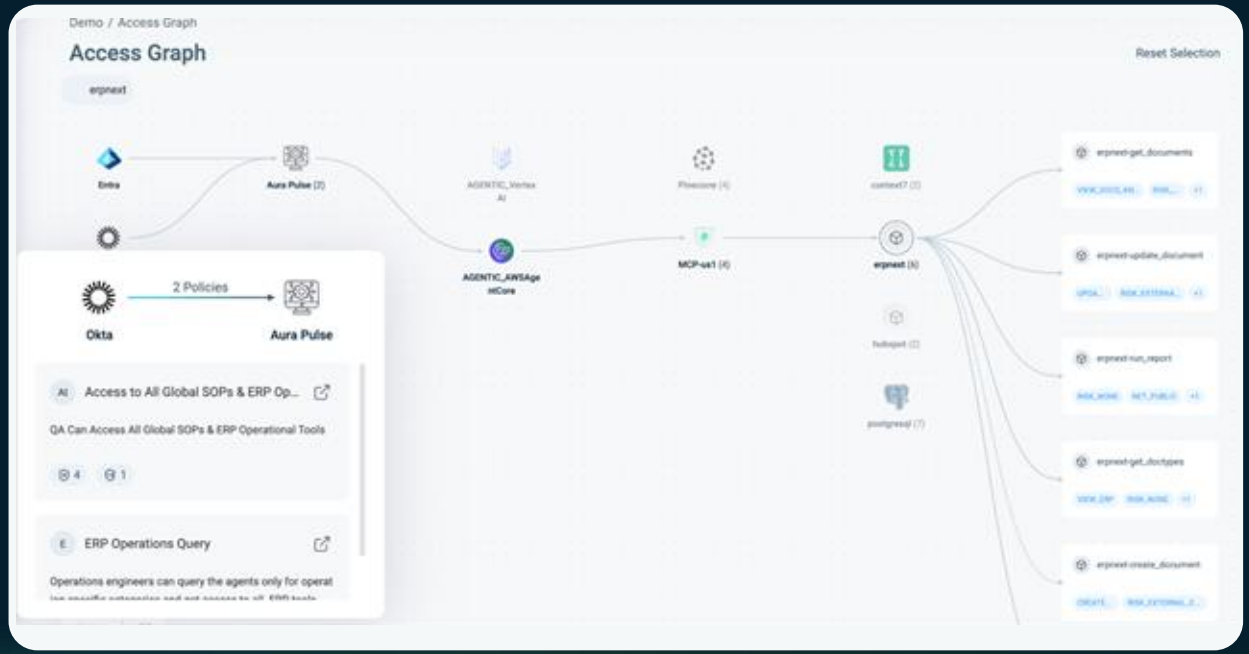

# Business Outcomes that Matter

Shifting to runtime decisions unlocks critical security and operational capabilities that static permissions cannot achieve

## Zero Standing Privileges.

100% alignment with core security principles

## 100% visibility

Who, What and Why

## Time-to-Value

<1 min enforcement for new and updated policies

## Cost Savings

30-50% reduction in cost of authorization solutions and processes

# Market Leadership

Gartner

## Market Guide for Guardian Agents

25 February 2026 - ID G00836388

By: Avivah Litan, Daryl Plummer, Carlton Sapp, Dionisio Zumerle, Tom Coshow, Max Goss, Lauren Kornutick

Initiatives: Delivery of Functional Responsibilities; Govern Agentic AI for Greater Business Autonomy

## Reference Architecture Brief: IAM for AI Agents and Other Workloads

20 April 2026 - ID G00850819

By: Erik Wahlstrom

Initiatives: Delivery of Functional Responsibilities; Build and Evolve a Resilient and Agile Cybersecurity Program

## Innovation Insight: Authorization Management Platforms

8 August 2025 - ID G00828568

By: Paul Mezzera, Nathan Harris

Initiatives: Identity and Access Management

kuppingercole ANALYSTS

## Policy Based Access Management

Graham Williamson & Martin Kuppinger

February 19, 2024

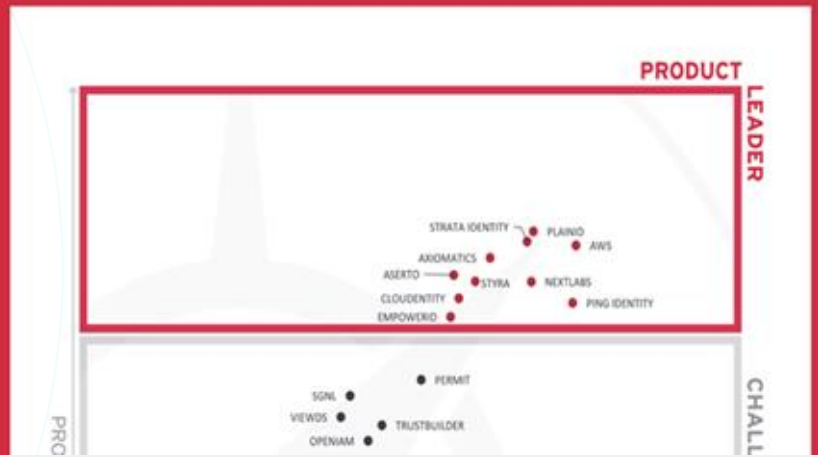

# Key Differentiators

- Identity-first Security
- Centralized management, distributed enforcement
- Runtime Authorization at Enterprise Scale
- Pre-Built Authorization Enforcement Across the Enterprise Stack

# AI agents need boundaries. Most don't have them.

Agents are designed to automate business processes. But without proper controls, they can quickly go off track and expose your organization to critical risks.

Gartner

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

https://www.gartner.com/document-reader/document/6269683?ref=solrAll&refval=464350290

# The market mandates control.

Leading analysts and standards bodies agree: **static access controls are insufficient for the dynamic nature of Agentic AI.**

ANALYST FRAMEWORK

## Gartner TRiSM

Implement end-to-end governance to control AI Trust, Risk, and Security Management (TRiSM).

SECURITY STANDARD

## OWASP GenAI Threat Defense

Specifically address risks emerging from autonomous agents and external tool access, majority are access controls related.

CLOUD SECURITY

## CSA Agentic AI IAM

Adopt identity-first controls that explicitly govern non-human (agent) identities.

CLOUD SECURITY

## OIDC AI Agent Identity

Standardize how agent identities are issued, verified, and trusted across systems.

# The ONLY End-to-End Solution

- **Intent-based, Zero-Standing-Privileges**
  Mapping human context to every autonomous agent action
- **Identity First Security**
  Human, NHI and Agent enforced.
- **Discover ▪ Manage ▪ Authorize**
  Enterprise grade operational framework
- **Full Flow Coverage**
  From prompt to data retrieval, tools, and output

**Centralized control** over what AI agents can **access**, **do**, and **expose** across the full AI flow, through dynamic policy-driven guardrails.

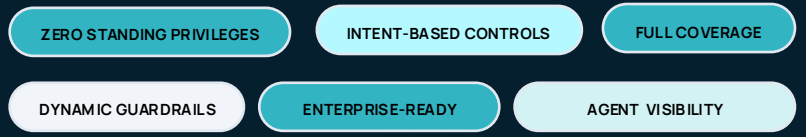

# Enforce control across the full agentic AI flow

Comprehensive guardrails bind every decision to identity and intent, from the first prompt to the final output.

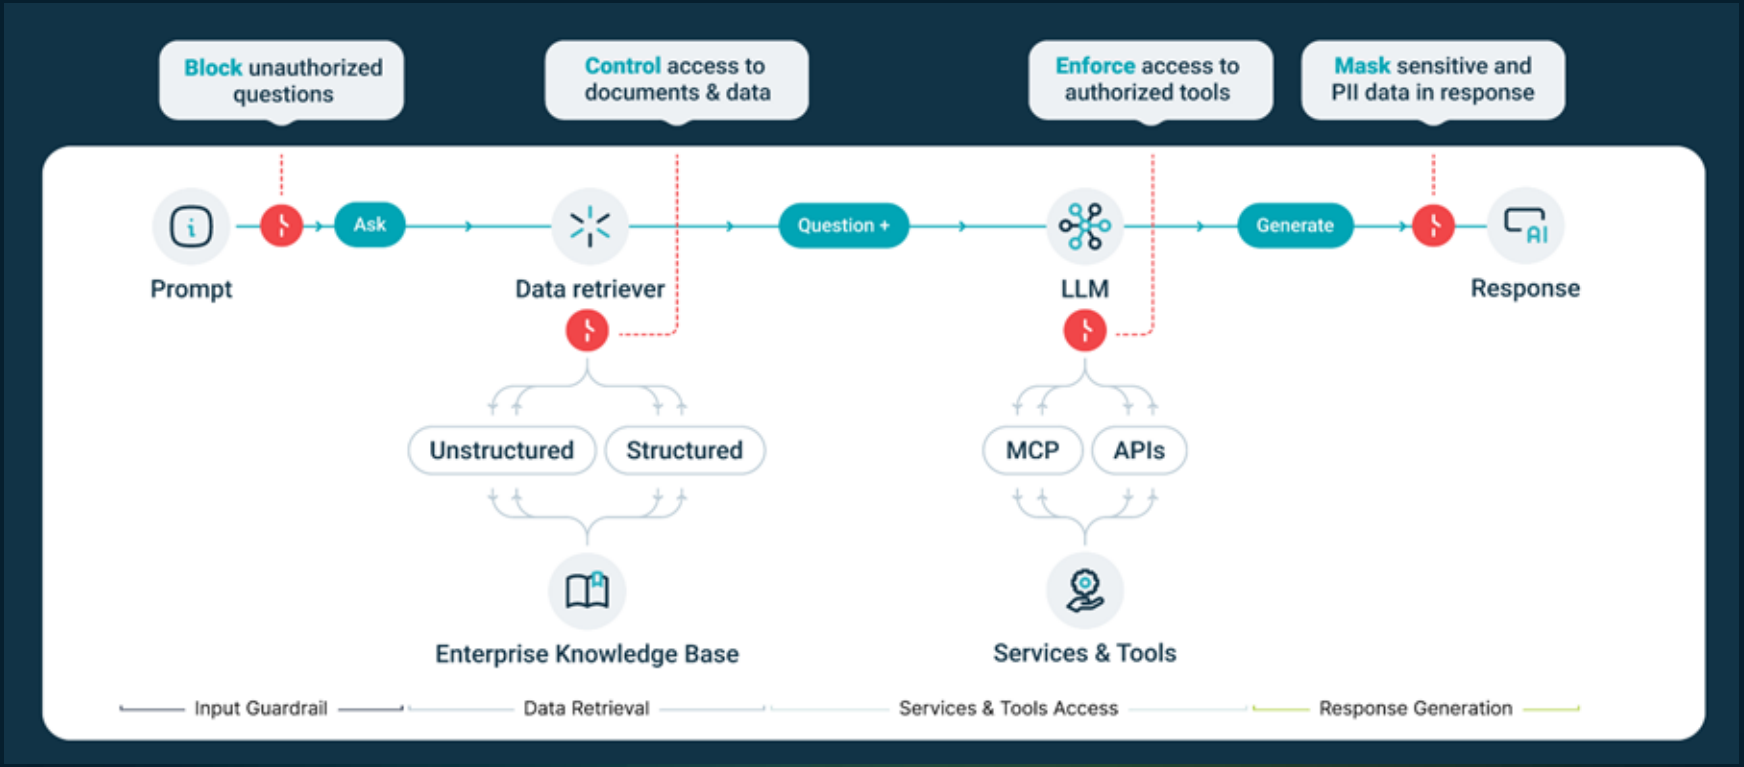

# 01

Guardrail 01

## Input Guardrail

An intelligent front door that validates every prompt before it enters your AI ecosystem.

## Intelligent Front Door

### Interpret Intent & Bind to Identity

Interprets user intent and reads identity signals immediately. Binds the prompt to the specific human and agent identity, ensuring context is established before processing begins.

### Check Policies Before Data Calls

Evaluates permissions against defined policies in real time. Automatically approves, transforms, or blocks prompts before any expensive or sensitive data retrieval occurs.

### Constrain Prompts to Approved Scope

Stops prompt unauthorized attempts at the source. It ensures the agent operates strictly within its approved business scope, preventing scope creep.

# 02

Guardrail 02

## Data Guardrail

A comprehensive filter layer that secures data access at the source, determining who can access what in real time.

## Pre-Retrieval Control

### Filter Structured & Unstructured Data

Apply precise filters to data retrieval calls. Whether it's SQL rows or vector embeddings, ensure only relevant data is fetched based on policy.

### Govern Topics & Metadata Access

Control access based on data classification, authorized topics, and metadata. Enforce granular permissions before the model ever sees the data.

### Prevent Unauthorized Retrieval

Stop the retrieval of unauthorized documents at the source. Focus agents on authorized, relevant content to minimize exposure risk.

# 03

Guardrail 03

## MCP Guardrail

A dedicated control layer for Model Context Protocol (MCP) tools, ensuring agents access only approved capabilities and parameters.

## Tool Use Under Control

### Discover & Classify Tools

Automatically discover and classify MCP servers and tools according to business purpose. This creates a managed inventory for agent usage, ensuring visibility into what tools are available.

### Validate Identity & Context

Enforce real-time identity and context validation before any tool is accessed. Access decisions are bound to the specific user and agent identity, preventing unauthorized tool usage.

### Constrain Actions to Scope

Apply granular constraints to tool actions and parameters. This ensures agents operate strictly within their approved scope, preventing unintended side effects or parameter manipulation.

# 04

Guardrail 04

## Output Guardrail

A smart safety layer that secures what leaves the system, ensuring data privacy and compliance.

## Secure what leaves the system

### Validate Responses Against Policy

Validates every generated response against active policies before delivery. Ensures that only permitted information is returned to the user, masking unsafe content.

### Stop Unauthorized Disclosures

Dynamically masks sensitive fields and PII in real time. Prevents data leakage by redacting or transforming confidential data based on the user's authorization level.

### Bind Output to Identity & Explainability

Records complete evidence for compliance. Ties the final output back to the original identity, intent, and policy decision, providing a full audit trail for explainability.

Platform Capabilities

# End-to-End Agentic Identity Platform

## Discover

### Tools Discovery & Classifier

Automatically detect MCP tools and classify them by usage category.

### RAG Discovery & Classifier

Connect to vector DBs to enrich and learn existing classifications and metadata.

### Data Source Visibility

Gain insight into what data and APIs your agents are attempting to access.

## Manage

### Identity-Aware Controls

Manage access policies for Humans, Agents, and how they interact.

### Investigation View

Graph-based visibility into agent flows, datasets, and related identities.

### No Code AI-Guided Policy Builder

Describe intent in natural language to automate policy creation.

### Full Audit Trail , For Auditors

Access decision logs with plain-language reasoning for compliance.

## Authorize

### End-to-End Guardrails

Enforce policies across Input, Data Retrieval, MCP Tools, and Output.

### Granular Control

Apply filters and constraints before data is retrieved or exposed.

### Built-in Output Masking

Dynamically redact sensitive data based on user entitlements.

### Flexible Integration

Compatible with all major AI frameworks and enterprise stacks.

# Agentic AI changes everything, Runtime Authorization makes it secure

**End-to-end solution**
Discover, Manage, Authorize

**Enforce all identities**
Human and Non-human

**Scale AI responsibly**
With ZSP

Learn more at plainID.com

plainID

Discover Manage Authorize

# Thank You!