

PlainID + Microsoft 365 E5/E7

Agentic IAM Better Together Capability Comparison

Based on PlainID Agentic IAM capability model and public Microsoft E5/E7 capability descriptions.
date: 9 June 2026

Microsoft 365 E5/E7 gives customers a powerful Microsoft security, identity, compliance, and agent-management foundation; PlainID complements it with externalized, fine-grained, cross-platform runtime authorization and policy governance for agent, tool, API, and data interactions.

Executive summary

- **Microsoft 365 E5 is the security and compliance baseline.** E5 brings Microsoft Entra ID P2, Defender, Purview, Intune and advanced analytics/security capabilities that protect users, devices, identity, data and collaboration surfaces. It is strong foundation for human workforce security and AI readiness, but it is not positioned by Microsoft as the full agent-control plane.
- **Microsoft 365 E7 is the Microsoft agent-governance step-up.** E7 includes E5, Copilot, Agent 365 and Entra Suite, and is positioned by Microsoft for agent observability, security, governance, agent compliance, secure access to AI/apps/resources, and agentic operations.
- **PlainID is the externalized authorization and policy control layer.** PlainID focuses on agent discovery and visibility, centralized policy management, runtime authorization enforcement across prompts, tools, APIs and data, and audit/observability. Identity foundations are intentionally integrated with enterprise identity and secrets providers.
- **The combined value is stronger than either layer alone.** Microsoft can identify, protect and govern Microsoft estate users, data and agents; PlainID turns that context into consistent, fine-grained authorization decisions at runtime across heterogeneous agent frameworks, MCP gateways, APIs, SaaS, data sources and RAG/vectorized environments. Additionally PlainID offers a cross-platform support

PlainID makes Microsoft identity, security and data signals actionable as fine-grained authorization decisions for agents, tools, APIs and data across both Microsoft and non-Microsoft environments.

A new architectural layer is required to manage autonomous systems

Gartner.

Market Guide for Guardian Agents

25 February 2026 - ID G00836388 - 35 min read

By: Avivah Litan, Daryl Plummer, Carlton Sapp, Dionisio Zumerle, Tom Coshaw, Max Goss, Lauren Kornutick
Initiatives: Delivery of Functional Responsibilities; Govern Agentic AI for Greater Business Autonomy

Guardian agents supervise AI agents, helping ensure agent actions align with goals and boundaries. They monitor and block risky actions and are evolving from a collection of services to autonomous agents that enforce policies across platforms. AI leaders can use this Guide to understand the market and vendors.

01 Incumbent Vendors Lack the Agility to Secure Fast-Evolving AI Agents

"The inability of incumbent vendors to update as quickly with modern controls could leave critical gaps... Clients should consider picking a best-of-breed provider... to counter the sophisticated emerging threats and unpredictable behavior associated with AI agents."

02 AI Demands the Convergence of Identity and Data Governance

While legacy IAM platforms stop at validating an agent's identity, we natively converge identity and data. We dynamically evaluate the sensitivity of the data the agent is interacting with in real-time to ensure even valid identities cannot expose unauthorized data.

03 An Independent, Universal Enforcement Layer is Mandatory

"No cloud provider can unilaterally enforce runtime control over AI agents once they operate across another provider's cloud... Only a neutral, trusted guardian agent layer... acts as the missing universal enforcement mechanism."

Capability model used for the comparison

The Agentic IAM model is organized around five architectural categories: Agents Identity, Agents Discovery & Visibility, Agents Policy Management, Agent Authorization Enforcement, and Audit & Observability. The model is designed to let organizations discover agents and resources, govern policies centrally, enforce runtime authorization across the agentic flow, and maintain auditability.

Layer	PlainID	Microsoft E5/E7	Together
1. Agent identities	Uses enterprise identity, machine identity and secrets platforms as trusted sources; evaluates human, agent, service and delegated identity context together.	E5 supplies Entra ID P2 identity and access management; E7 adds Entra Suite and Agent 365 agent identity and access package capabilities.	PlainID does not replace Microsoft identity; it consumes and enriches identity context for authorization.
2. Discovery and visibility	Agent registry, profile enrichment, access graph, MCP/tool discovery, resource discovery and categorization.	E7 Agent 365 provides centralized agent registry, usage insights and visual mapping of agent activity/connections. E5 alone is not the dedicated agent-management tier.	PlainID expands visibility into authorization-relevant relationships across agents, MCP/tools, APIs, data and resources. And it adds agent and resource information from other non-Microsoft agentic technologies.
3. Policy management	Policy360, no-code and policy-as-code authoring, policy governance, simulation, certification, delegated admin and cross-identity policies.	Microsoft policies span Entra Conditional Access, Purview data/compliance controls, Defender/Intune security controls, and E7 Agent 365 guardrails/templates.	PlainID provides cross vendor policy management layer. PlainID extends Microsoft identity controls to full business authorization logic including the context of the end use identity and the resources that are accessed.
4. Runtime enforcement	Runtime authorization, enriched context, externalized decisions, SDKs, gateway plugins, MCP controls, tool-level and parameter-level controls, data access controls, input/output guardrails.	E7 Agent 365 adds agent security controls, risk-based access, real-time protection and ability to block unsafe behaviors/tool invocations.	PlainID provides cross vendor runtime enforcement controls, as well as expands the microsoft control to non-microsoft resources. Together: Microsoft provides security signals and platform controls for microsoft manages resources; PlainID adds fine-grained decisions near protected tools, APIs and data.
5. Audit and observability	Authorization decision audit trail, policy-to-decision correlation, cross-flow traceability, access/authorization insights, admin audit.	E7 includes agent observability logs, agent analytics, logging/reporting/audit for agent actions and interactions.	PlainID provides cross vendor observability layer. Additionally, PlainID can provide explainable authorization evidence: what policy decided, why, and what context was evaluated.

Better Together Summary

Area	Microsoft E7 / Agent 365	With PlainID
Agent identity	Entra Agent ID gives agents first-class identity and lifecycle controls.	PlainID consumes agent identity as an authorization attribute and combines it with user, resource, action, risk, and business context.
Authentication / session access	Conditional Access evaluates agent/user/session risk, device, location, and network.	PlainID adds fine-grained authorization after authentication: what the agent can do, on which resource, for which purpose, under which business conditions.
SharePoint / OneDrive access	Enforces Microsoft 365 permissions, restricted access control, restricted content discovery, and sharing policies.	PlainID extends the same policy model to non-Microsoft repositories, databases, APIs, MCP servers, RAG/vector stores, and application data.
Microsoft 365 data governance	Purview provides sensitivity, DLP, audit, DSPM, and compliance signals.	PlainID can use sensitivity/classification context as input to runtime authorization decisions and masking instructions.
Agent risk and detection	Defender and Agent Registry surface risks such as prompt injection, sensitive data access, excessive permissions, and shadow agents.	PlainID can use risk as an authorization condition and enforce allow/deny/mask/step-up decisions inline.
Runtime enforcement	Strongest for Microsoft-native access paths, SharePoint/OneDrive, Conditional Access, DLP-supported flows, and Defender-supported protections.	PlainID provides externalized runtime authorization for Microsoft and non-Microsoft agents, MCP tools, APIs, data, parameters, prompts, and outputs.
Contextual authorization	Microsoft considers user, agent, session, risk, device, location, permissions, and content governance signals.	PlainID adds policy decisions that combine end user + agent + delegated identity + action + resource + data sensitivity + business context + environment + risk.

Detailed capability comparison

Legend: Strong = broad native coverage; Partial = coverage depends on Microsoft workload, configuration, add-ons, or scope; Complementary = PlainID extends or deepens the Microsoft capability; Gap/Not primary = not the main job of that Microsoft SKU based on public descriptions.

Capability area	PlainID	Microsoft 365 E5	Microsoft 365 E7	Better together message
Agent/workload identity	Cross-identity model for human, non-human, machine and agent context; delegated/on-behalf-of accountability; JIT/zero-standing-privilege patterns; identity enrichment.	Strong human/workforce identity foundation with Entra ID P2, risk-based Conditional Access and privileged access for critical resources.	Strong plus Agent 365 agent identity, Conditional Access, identity protection and access packages for agents, plus Entra Suite.	Use Microsoft as identity source of truth and risk signal provider; use PlainID to translate identities and delegated context into fine-grained authorization decisions at runtime.
Agent registry and inventory	Authoritative agent registry, profile enrichment, ownership, environment, frameworks, connected systems and business purpose.	Gap: E5 positioning does not describe a dedicated agent registry as included by default.	Partial to strong depending on integration: Agent 365 provides centralized agent registry and views into adoption, activity and health.	E7 / Agent 365 gives customers a Microsoft-native agent inventory and governance plane. PlainID complements this by extending discovery, relationship mapping, and runtime authorization across heterogeneous agent ecosystems, including MCP servers, tools, APIs, data sources, RAG/vector repositories, and non-Microsoft runtime enforcement points. PlainID can give customers a deeper MCP/tool authorization inventory and make that inventory actionable in policy.
MCP and tool inventory	MCP discovery, tool categorization/enrichment, resource discovery, resource classification and tool change detection.	Gap: public E5 docs do not present granular MCP/tool inventory as a core E5 capability.	Partial to strong depending on integration: E7/Agent 365 provides activity mapping, public docs emphasize agents more than detailed MCP parameter inventory.	
Policy authoring and governance	Policy360, native policy view, no-code authoring, policy as code, relationship-based rules, dry-run, simulation, certification, explainability and delegated administration.	Partial : Policy foundation across Entra, Purview, Defender and Intune.	Partial: Microsoft governance layer is mostly based on extended conditional access, and existing security layers on data sources; It is centered on Microsoft control surfaces.	PlainID provides a centralized authorization policy plane that can span Microsoft and non-Microsoft tools, APIs and data.
Runtime authorization	Externalized, context-rich authorization decisions at execution time for agents and humans; runtime decisions based on identity, action, resource, environment, sensitivity and risk.	Partial: E5 supports risk-based access and security controls, but not a purpose-built externalized authorization layer for every agent/tool/API/data action.	Partial: Agent 365 extends risk-based access and security to agents, however the focus is mostly on the identity (extended conditional access).	Microsoft evaluates access and risk in the Microsoft ecosystem, with strong focus on the identity side; PlainID adds application- and resource-level authorization decisioning across heterogeneous runtime points.
MCP, API and tool enforcement	MCP tool control, tool-level authorization, tool parameter controls, dynamic scope adjustment, allow/deny lists, gateway plugins, SDKs and MCP Gateway.	Gap: E5 security controls are not positioned as MCP/API parameter enforcement.	Partial: Microsoft states E7/Agent 365 can detect suspicious agent activity and block tool invocations within the Microsoft ecosystem.	PlainID deepens enforcement granularity: who/which agent can call which tool/API, with which parameters, under which context.
Data access controls	Structured/unstructured data access control, row-level filtering, column masking, vector/RAG discovery and metadata enrichment.	Partial : Mostly classification in Microsoft Purview data security/compliance.	Partial: Mostly based on pre-defined group assignments.	Purview classifies and protects data; PlainID makes real-time authorization decisions before retrieval/exposure and can apply masking/filtering instructions.
Input and output guardrails	Classifies/restricts prompts, detects sensitive intent/disallowed actions, categorizes outputs, masks/encrypts sensitive output elements.	Gap: E5 is not a generalized agent prompt/output policy engine.	Partial: E7 adds agent prompt injection controls.	PlainID provides input and output guardrails, based on business context, at runtime. That includes allow/block prompt, and masking of sensitive and PII data in the response.
Audit and explainability	Authorization decision audit trail, policy-to-decision correlation, business-readable reasoning, admin audit, simulation audit and cross-flow traceability.	Strong Microsoft auditing and compliance reporting in Purview/Defender/Sentinel contexts, but authorization explanation is distributed by workload.	Strong agent observability/logging/reporting/audit for agent actions and interactions in the Microsoft ecosystem.	PlainID provides full audit and observability for all runtime events, across technologies.