Microsoft Wakes Up to the Shadow AI Problem:
Agent 365 Goes Live for the Enterprise

With autonomous agents quietly multiplying inside corporate networks, Microsoft is shipping the governance layer it argues every enterprise will soon need.

Microsoft has officially moved Agent 365, its enterprise control plane for AI agents, out of preview and into general availability. The release is more than a routine product graduation. It is the company's most direct response yet to a phenomenon spreading inside organizations of every size: employees installing AI coding assistants, productivity bots, and autonomous workflow tools on their work devices, often well outside the line of sight of IT and security teams.

Microsoft has a name for this trend, and it is one that security leaders are about to hear a lot more often: shadow AI.

The new shadow problem inside corporate walls

Shadow IT was the headache of the last cloud era, when business units quietly spun up SaaS subscriptions without telling the security team. Shadow AI is its faster, stranger sequel. Instead of unauthorized apps, organizations are now dealing with autonomous software that can read documents, write code, send emails, and trigger workflows on a user's behalf, sometimes with surprisingly broad access.

01

In an interview tied to the launch, David Weston, Corporate Vice President of AI Security at Microsoft, framed the dilemma neatly. Most enterprises, he said, are stuck oscillating between two extremes: a permissive stance where any agent is allowed to run, and a restrictive one where nothing meaningful is allowed at all. Agent 365 is being pitched as the middle path, a way to permit useful agentic work while keeping a clear ledger of what is running, who owns it, and what it can touch.

02

The timing is not accidental. Industry analysts have been warning that agent counts inside large organizations are about to explode. IDC has projected more than a billion AI agents in active use within the next few years, and Microsoft itself has cited research showing the vast majority of Fortune 500 companies already have some form of agent activity in production. The governance scaffolding to manage that wave has, until now, lagged badly.

What Agent 365 actually does

At its core, Agent 365 is a control plane. It plugs into the security and identity stack most large Microsoft customers already run, and gives them a single place to inventory, govern, and secure AI agents across the organization. Agents get assigned their own identities through Microsoft Entra, governance policies through Microsoft Purview, and runtime protection through Microsoft Defender. Device-level enforcement runs through Microsoft Intune.

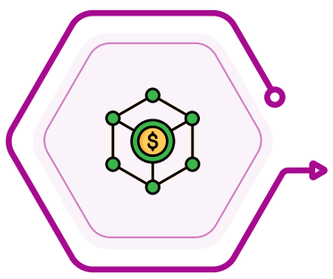

The pricing model reflects a deliberate design choice. Rather than charging per agent, which would punish organizations for the very sprawl Microsoft is trying to help them tame, Agent 365 charges per person who interacts with the agent ecosystem. The standalone license is priced at fifteen dollars per user per month, with bundled access available through the new Microsoft 365 E7 Frontier Suite that also rolls in Copilot, the Entra Suite, and E5 security tooling.

The headline feature: hunting down local agents

The most consequential addition in this release is the platform's ability to find and manage agents that live directly on a user's Windows device. These are the tools developers and knowledge workers tend to install on their own initiative, sometimes downloading them from open repositories without ever consulting IT.

From day one of general availability, organizations enrolled in Microsoft's Frontier program can use Agent 365, working through Defender and Intune, to detect OpenClaw agents running on managed Windows endpoints. Administrators get a new Shadow AI page inside the Microsoft 365 admin center that surfaces which devices are running these tools, and they can apply Intune policies to block common ways of executing them. Microsoft has also signaled that detection coverage will expand soon to include other widely used local agents such as GitHub Copilot CLI and Claude Code.

Reaching beyond the Microsoft cloud

Agent sprawl does not respect cloud boundaries, and Microsoft has acknowledged as much. Alongside the GA release, the company introduced public previews of registry sync connections with AWS Bedrock and Google Cloud's Gemini Enterprise Agent Platform. The integrations are designed to let IT teams automatically discover and inventory agents running on those platforms, with basic lifecycle controls such as start, stop, and delete actions on the roadmap.

01

Microsoft has also lined up an unusually broad set of partners to extend Agent 365's reach. Adobe, SAP, Zendesk, Manus, and others have agreed to make their agents fully manageable through the platform. On the services side, integrators including Accenture, KPMG, Capgemini, Protiviti, and Slalom, joined by close to two dozen others, are positioning themselves as launch partners offering inventory assessments, least-privilege rollouts, and ongoing lifecycle management.

02

Adding to the announcement, Windows 365 for Agents entered public preview in the United States. The new Cloud PC variant is purpose-built for running agents in a policy-controlled environment, isolated from human user sessions but still subject to the same security and management controls IT teams already use for employees.

A phased path: crawl, walk, run

For organizations weighing how to roll Agent 365 out, Weston offered a three-stage progression that Microsoft believes is realistic inside a 90-day window. The first stage, the crawl phase, is purely about visibility. Teams cannot secure what they do not know exists, so the immediate priority is building an inventory of agents already operating across the environment.

01

The walk phase shifts attention to identity and access. Each discovered agent gets a managed identity and a clearly bounded set of permissions, treating the agent more like a constrained employee account than an anonymous process. The run phase is where the deeper security controls come into play: isolation through Windows 365 for Agents, runtime blocking of suspicious behavior, and blast radius mapping to understand what damage a compromised agent could realistically cause.

02

Whether enterprises actually move through these stages on Microsoft's suggested timeline is another question. The pace will depend on how mature an organization's existing endpoint management and identity infrastructure already is, and how quickly shadow AI is multiplying within its walls.

Why Microsoft thinks it can win this market

Microsoft is far from alone in chasing the agent governance opportunity. Google, Amazon, and Salesforce are all developing their own orchestration and oversight tools, and a long list of cybersecurity vendors are racing to attach AI-specific capabilities to their existing platforms. What sets Microsoft's pitch apart is integration depth. The company already sits at the center of how most large enterprises handle endpoints with Intune, threat detection with Defender, identity with Entra, and productivity with Microsoft 365.

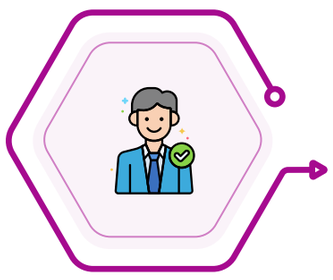

Agent 365 leans on those existing rails rather than asking customers to learn yet another security console. For a CISO already drowning in tooling, the appeal is obvious. The risk, also obvious, is that any product depending on so many surfaces can take time to mature evenly across all of them, a pattern Microsoft watchers have seen before.

What still sits behind the preview gate

Not everything announced is generally available. Several capabilities remain in Frontier preview, including Microsoft's developer SDK for building Agent 365-aware agents, agent identity authentication features, and what the company calls agentic users, which are agents that operate as their own first-class identities rather than acting on behalf of a specific human user. Context mapping, policy-based controls, and runtime blocking and alerting through Intune and Defender are scheduled for public preview in June 2026.

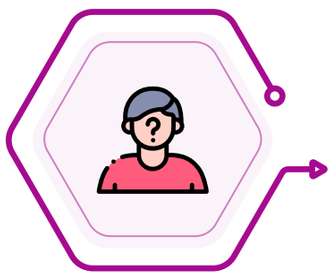

That phasing is intentional. The on-behalf-of model maps cleanly onto the identity governance patterns enterprises already understand. Fully autonomous agentic identities raise harder questions about audit trails, accountability, and ongoing oversight, and Microsoft has chosen to keep them behind a beta gate while collecting more deployment data.

The real takeaway

The bigger story behind this release is not a feature list. It is the implicit acknowledgment from one of the largest enterprise software vendors in the world that AI agents have already moved faster than the controls built to manage them. The shadow AI problem is not a future concern that organizations have time to plan for at leisure. It is a present-day reality showing up on devices, in code repositories, and inside SaaS workflows right now.

Agent 365's arrival in general availability marks the moment when agent governance crosses from optional infrastructure into something Microsoft expects every serious enterprise to be running, in some form, within the next year or two. Whether that turns out to be the right call depends on how aggressively shadow AI keeps spreading, and how quickly Microsoft can make this control plane feel less like a checkbox and more like the operating layer for a workforce in which a growing share of the work is no longer done by humans at all.