ME7/A365 – Launch FAQ for Partners

LAST UPDATED: March 9, 2026

Contents

Summary

This document provides a consolidated set of frequently asked questions (FAQs) for partners regarding the launch of Microsoft 365 E7 and Microsoft Agent 365. It covers product details, licensing, billing, promotions, incentives, compliance, and adoption guidance. Use this as your go-to resource for customer conversations and internal readiness.

Resources

Partner Center

- **Link**: https://partner.microsoft.com/

Pricelist

- **Link**: https://partner.microsoft.com/dashboard/pricing/pricelist

Microsoft 365 Roadmap

- **Link**: https://www.microsoft.com/microsoft-365/roadmap

Microsoft 365 E7 Product & Value

What is Microsoft 365 E7?

Microsoft 365 E7 is the productivity infrastructure for a human-led, agent-operated enterprise, turning human intent into action with AI that runs securely at scale. ME7 brings together Copilot, agents, and built-in governance via Agent 365 so organizations can operationalize AI safely at scale.

ME7 combines the full E5 security and compliance foundation with a new, integrated AI stack delivered as a single, governed platform. Unlike past suites that primarily added features or workloads, ME7 introduces a human-led, agent-operated system of work where AI can take action, not just generate content. It also adds a unified governance layer for AI agents and cross-app intelligence that earlier suites did not include, enabling organizations to securely operationalize AI at scale.

Microsoft 365 E7 combines the following components:

- **Microsoft 365 E5:** Provides the core productivity, security, identity, and compliance foundation required to run work securely at enterprise scale.
- **Microsoft 365 Copilot:** Embeds AI directly into everyday work experiences to help people reason, create, and act with context from their organization's data.
- **Microsoft Entra Suite:** Extends identity and access controls to users, apps, and agents, enabling secure access, conditional controls, and least-privilege governance.
- **Agent 365:** Delivers centralized governance, visibility, and control for AI agents—so organizations can deploy and scale agents safely across the enterprise.

What did we announce on 3/9 about Microsoft 365 E7?

Microsoft publicly **announced the launch of Microsoft 365 E7, the suite for the Frontier Worker**. The announcement positioned ME7 as the solution to help organizations move from AI experimentation to enterprise-wide execution.

Additional blogs and announcements can be found at the below links:

Microsoft 365 Blog | Latest Product Updates and Insights

AI Business Solution for Partners Blog

Who is Microsoft 365 E7 designed for?

Microsoft 365 E7 is designed for organizations that want AI to work across the entire business, not just in pilots or isolated teams. It equips every employee to work with Copilot

and AI agents in their daily tools, while giving the organization a single, trusted foundation to run AI at scale.

Employees interact with ME7 through Copilot and agents embedded in the flow of work. IT and security teams set the guardrails—identity, data protection, governance, and oversight—so AI can take action safely and consistently across the organization.

ME7 is ideal for enterprises moving from AI experimentation to enterprise-wide operationalization, where AI needs to be productive for employees and governed for the business.

What business problems does Microsoft 365 E7 solve?

As organizations move beyond isolated AI pilots, they face three core challenges that prevent AI from delivering enterprise-wide impact:

1. AI is fragmented and disconnected from real work

Most AI deployments live in silos—separate tools, point solutions, or experimental workflows. Employees have to leave their everyday apps to access AI, and agents often operate without full context of how work actually gets done across teams and systems.

ME7 solves this by embedding AI directly into the flow of work.

With Copilot and agents grounded in Work IQ, AI understands organizational context—people, roles, content, and workflows—and can take action across Microsoft 365 apps and business systems. Employees interact with AI where they already work (Teams, Outlook, Word, Excel), turning intent into execution without switching tools or learning new interfaces.

2. AI experimentation doesn't scale to the whole organization

Many organizations can pilot AI in pockets, but struggle to roll it out broadly. Without a shared foundation, AI adoption varies by team, value is inconsistent, and benefits never compound across the business.

ME7 solves this by providing a single, integrated AI platform for the entire workforce.

ME7 is designed to be deployed broadly so every employee can use Copilot and agents (1P, custom-build and 3P) in their daily work—while drawing from the same intelligence layer and governed environment. This enables consistent adoption, shared learning, and scalable productivity gains across functions, not just isolated teams.

3. Organizations can't trust AI to operate safely at scale

As AI usage grows, so do concerns about security, data leakage, compliance, and uncontrolled proliferation of agents. Without built-in governance, AI becomes a risk multiplier instead of a business advantage.

ME7 solves this by making trust a built-in property of AI.

Through Agent 365 and the E5 security foundation, ME7 extends enterprise-grade identity, access, data protection, and audit controls to Copilot and agents (including 1P and 3P)—just as organizations already do for users. IT and security teams define the guardrails once, and those controls apply consistently across users, data, and agents as AI scales.

In short, Microsoft 365 E7 helps organizations move from AI experimentation to enterprise execution—embedding AI into everyday work, scaling it across the workforce, and governing it with the same rigor as the rest of the enterprise.

When can customers purchase Microsoft 365 E7?

Microsoft 365 E7 will be generally available for purchase starting May 1 (5/1). Customers will be able to buy Microsoft 365 E7 through standard commercial purchasing channels at GA.

Does every employee type need Microsoft 365 E7 or is this only for IT Admins who manage agents?

ME7 is designed to be deployed broadly because governance and control need to apply wherever AI is used. If employees who aren't covered can still interact with assistive agents (e.g., in M365 app surfaces), then IT won't have consistent policy enforcement, visibility, and controls across the agent estate—creating gaps and "shadow" usage. IT admins manage the guardrails, but the guardrails are meant to protect the whole organization—so ME7 is fundamentally about enterprise-wide operationalization, not an IT-only SKU.

What are the core benefits for customers adopting Microsoft 365 E7?

Microsoft 365 E7 (ME7) delivers a complete, integrated foundation for enterprise AI by unifying intelligence, execution, innovation, and control. It enables organizations to move from isolated AI experiments to durable, enterprise-wide impact—without sacrificing security, governance, or visibility.

1. Intelligence that understands you, your work, and your organization

- ME7 provides an intelligence layer grounded in real work signals so AI can act with context, continuity, and relevance across the organization.

- Work IQ builds on signals across email, meetings, documents, chat, and collaboration to give AI full situational understanding.
- Context continuity preserves intent across tasks, apps, and moments of work, enabling consistent outcomes aligned to real priorities.
- Inference at work scale brings together models, skills, and tools so Copilot and agents don't just understand work—they move it forward.

2. AI embedded in the flow of work

- ME7 embeds AI directly into the tools employees already use, reducing friction and accelerating adoption.
- Copilot in Microsoft 365 apps (Word, Excel, PowerPoint, Outlook, Teams) delivers assistance and execution without switching contexts.
- AI-powered collaboration enables shared, multiplayer workflows where Copilot maintains context and momentum.
- Work partner model uses AI to complete end-to-end workflows, freeing employees to apply judgment and direction.
- Model choice allows the best model for the job, improving quality and consistency of results.

3. Agents at your business edge

- ME7 enables innovation to happen wherever work happens—across roles, teams, and functions.
- Agents at the business edge accelerate time to value by turning intelligence into action across real workflows.
- Agent flexibility lets customers choose from prebuilt Microsoft agents, third-party agents, or build custom agents when needed.
- Low and no-code tools empower domain experts—not just developers—to create and use agents.
- Broad ecosystem support ensures agents built with Copilot Studio, Azure AI Foundry, and partner tools integrate seamlessly into Microsoft 365.

4. Control plane for AI adoption

- ME7 provides enterprise-grade security and compliance that protects users, data, and agents together—so innovation can move fast without compromising trust.

- Centralized governance ensures agents, Copilot experiences, users, and data operate within clearly defined security and compliance boundaries.
- ME7 gives organizations confidence to scale AI without losing visibility or control:
  - Built-in visibility makes AI activity observable across the organization, enabling leaders to understand what's running, where, and why.
  - Policy-driven controls help organizations manage risk proactively as AI usage expands across roles, teams, and tools.

What exactly is Work IQ?

Work IQ is a semantic + signals + memory layer over Microsoft substrate data (emails, files, meetings, chats). It existed conceptually before (semantic index in M365) but is now formalized and branded as "IQ". Work IQ is the intelligence layer in Microsoft 365 E7 that understands how work actually happens across an organization and uses that shared understanding to power Copilot and agents with relevance, continuity, and precision. By maintaining a living view of work in motion, Work IQ enables AI to move beyond generic assistance and instead help employees and agents actively move work forward in ways that reflect real priorities and how the organization operates.

Work IQ is built on three core capabilities:

- **Data:** Work IQ is built on rich signals across email, meetings, documents, chat, and collaboration, giving Copilot and agents the inputs they need to act with full contextual understanding of work as it happens.
- **Context:** Work IQ maintains continuity of context across tasks, apps, and moments of work, enabling Copilot to deliver consistent, relevant outcomes aligned to real business priorities — not isolated prompts or one-off interactions.
- **Inference:** Work IQ brings together models, skills, and tools in connected capabilities so Copilot and agents don't just understand work, but can reason over it and help drive progress end-to-end.

Without Work IQ, Copilot and agents operate without continuity of context, limiting their ability to deliver relevant, consistent outcomes that reflect how the organization truly works.

For example, you can now ask your analytics agent to build an excel model to calculate KPIs for the quarter, referencing a transcript from Microsoft Teams calls, email chains, and previous presentation documents.

How is Microsoft 365 E7 positioned in the marketplace?

Microsoft 365 E7 is positioned as the productivity infrastructure for a human-led, agent-operated enterprise. It is designed for organizations moving beyond AI experimentation and pilots to make AI a durable, business-critical capability. Unlike standalone AI tools or point solutions, E7 brings together productivity, AI, security, and governance into a single, integrated system—so AI can take real action in the flow of work while remaining visible, secure, and accountable at scale. This positions E7 not as "more AI," but as the operating model enterprises need to run AI safely across their workforce and business processes.

In the market, E7 differentiates by extending the Microsoft 365 foundation customers already trust to manage users, data, and devices, and applying that same model to AI and agents. As AI shifts from generating content to executing tasks, governance and control become the key differentiators. Microsoft 365 E7 uniquely treats agents as first-class entities with identity, ownership, lifecycle management, security, and auditability—deeply integrated across productivity apps, identity, data protection, and security. Competing offerings may provide AI capabilities or monitoring, but lack this level of native, end-to-end integration. E7 is positioned as the premium enterprise suite for organizations that want AI to operate broadly, safely, and confidently as part of how work gets done.

When and where will Microsoft 365 E7 be generally available?

General Availability and full transactability will be on 5/1. Pricelist preview will be available on 4/1 for EA and partner channels.

What countries will Microsoft 365 E7 be available in at launch, and are there any language limitations?

For General Availability (launch), Microsoft intends for Microsoft 365 E7 to be a worldwide service, available in all the same regions and datacenters as Microsoft 365.

What capabilities are available today?

As of General Availability, ME7 provides the following **core capabilities** to enable organizations to implement AI securely and at scale:

- **Microsoft 365 Productivity**. Core productivity apps for email, meetings, documents, collaboration, and file sharing.
- **Microsoft 365 Copilot.** AI assistance embedded directly into Microsoft 365 apps. Draft, analyze, summarize, plan, and take action using work context.
- **Copilot Studio.** Tools to build, customize, and extend Copilot and AI agents for business processes. Supports low-code and pro-code creation path.

- **Work IQ.** Intelligence layer that gives Copilot and agents understanding of work. Connects signals across content, meetings, tasks, workflows, and relationships.
- **A365.** Enterprise capabilities to discover, govern, secure, and manage AI agents at scale. (See below for more detailed information).
- **E5 security and compliance.** Enterprise-grade security and compliance foundations extended to AI and agents. Including Microsoft Purview, InTune, Defender and the Entra Suite.

How is Microsoft 365 E7 positioned in the marketplace?

Microsoft 365 E7 is supported through the same enterprise-grade support channels customers already use for Microsoft 365 services. Once generally available, issues related to Microsoft 365 E7—including Agent 365—can be logged through standard Microsoft 365 support routes in the Microsoft 365 Admin Center, and will be handled by trained Microsoft support engineers or routed to the appropriate product experts.

During early access programs, support is provided through dedicated channels. This includes internal Microsoft Teams support forums for participating customers and field teams, as well as direct engagement with the product group for nominated design-partner customers. Early access participants are also encouraged to share feedback directly with the product team through meetings, surveys, or private collaboration spaces.

Customers with Premier or Unified Support can continue to escalate Microsoft 365 E7 issues through their existing support agreements and Technical Account Managers, just as they would for other Microsoft 365 services. In addition, Microsoft will provide documentation and community forums (such as Microsoft Tech Community) where customers and IT professionals can discuss issues, share guidance, and surface non-urgent questions—though formal issue resolution should always go through official support channels.

How do customers report bugs or feature requests for Microsoft 365 E7?

Customers can use standard Microsoft channels:

Microsoft Support tickets – Customers should open a support case through the Microsoft 365 Admin Center. If an issue is confirmed as a product bug, it is logged and escalated through Microsoft's internal engineering pipeline, with a case ID for tracking.

Microsoft Feedback portal – Customers can submit feature requests or upvote ideas at feedback.microsoft.com (Microsoft 365 Feedback). These requests are reviewed by product teams as part of roadmap planning.

Microsoft account team – For high-priority or strategic feature requests, customers can work through their Microsoft account team, who can escalate feedback through internal product and engineering channels.

Community forums and Tech Community discussions can also be used for knowledge sharing and minor issues, but formal bugs and feature requests should go through the channels above to ensure proper tracking and follow up.

Microsoft Agent 365 Product & Value

What is Microsoft Agent 365?

Agent 365 is the control plane for agents. It extends the infrastructure that customers use for managing users and leverages familiar capabilities that have been adapted to agent needs. Agent 365 includes leading Microsoft security and compliance solutions — Defender, Entra, and Purview— to protect and govern agents.

With Microsoft Agent 365, organizations don't need to reinvent the wheel because the fastest path to confidently deploying agents is by managing them in a similar way to how they manage and secure their users. It delivers observability, governance, and security for all agents across their organization.

What did we announce on March 9th?

- **New Agent 365 pillars:** Agent 365's value proposition has been updated to three core pillars:
  - **Observe –** Monitor and manage agents in real time
  - **Govern** – Establish guardrails for agents and users
  - **Secure** – Protect all agents end-to-end
- **Pricing availability:** The price list for Agent 365 is scheduled to become available on **April 1, 2026.**
- **SKU readiness:** The SKU will be live and ready for transaction on **May 1, 2026,** marking the beginning of general commercial availability.

Is Agent 365 an AI agent, a platform, or a development tool?

Agent 365 itself is ***not*** an AI agent or a tool to build agents – rather, it is the **infrastructure layer** that supports and governs agents. It's not a new Copilot or some kind of "super agent." It's also not a replacement for Copilot Studio or Microsoft Foundry (which are agent development platforms). Instead, **Agent 365 is analogous to how IT and Security teams are currently securing, managing and controlling users in their M365 environment**. Once you have an AI agent, **Agent 365** helps you wrap the agent with an

identity, security policies, compliance controls, and administrative oversight. This lets the agent operate within your organization just as a user would, subject to the same rules and integrations.

Who is Agent 365 designed for?

Agent 365 is designed for organizations that are adopting AI agents and need to **manage them at scale securely**. The solution primarily speaks to **IT and Security leaders (CIOs, CTOs, CISOs)** responsible for enabling new technology while protecting the enterprise. In terms of organization size, initial focus is on **medium and large enterprises** – those likely to be actively investing in agentic AI.

These customers want to **harness AI agents** for productivity but are concerned about oversight. Agent 365 is built for *any enterprise that plans to deploy AI agents (from any source) and wants to do so in a controlled, compliant manner*. More specifically, it's for customers who have or anticipate scenarios like AI assistants in departments, autonomous workflow bots, or third-party AI services integrated into their environment.

Microsoft expects interest across industries: from **financial services** (where compliance is paramount) to **manufacturing** (using agents for automation) to **IT services** (i.e. managing fleets of AI support agents).

In short, **Agent 365 is aimed at IT and Security decision-makers** in organizations embracing AI – it gives them a complete solution to accelerate adoption *safely*.

What business problems does Agent 365 solve?

- **Observability**
    - **Business problem**: Organizations lose visibility as agents rapidly proliferate across teams, tools, and platforms, leading to unmanaged shadow agents, unclear ownership, and reactive fire drills when issues arise.
    - **How Agent 365 helps**: Agent 365 provides a unified control plane that gives IT and security leaders full visibility into all agents in the environment, how they're used, who's using them, and how they behave. With a centralized agent registry, analytics, and relationship mapping, organizations can proactively identify agent sprawl, usage trends, and emerging risks before they impact the business, shifting from reactive cleanup to proactive management.
- **Governance**
    - **Business problem**: Without consistent governance, agent adoption becomes chaotic. Some agents are reviewed and compliant, others are

not. Ownership and lifecycle management break down and IT loses confidence in scaling agents safely.
    - **How Agent 365 helps**: Agent 365 applies the same policy-driven governance model used for users and apps to agents. It enables IT-led onboarding, least-privilege access, lifecycle controls, audit logging, and built-in compliance. This ensures every agent starts governed and compliant from day one, allowing IT and Security teams to confidently innovate while maintaining accountability, audit readiness, and regulatory alignment.
- **Security**
    - **Business problem**: Agents behave like users but scale like applications, dramatically expanding the attack surface and introducing new risks such as excessive access, data leakage, prompt injection, and AI-specific threats, often without obvious warning signs.
    - **How Agent 365 helps**: Agent 365 extends enterprise-grade security controls to agents, including identity protection, conditional access, data loss prevention, and AI-specific threat protection. Security teams gain visibility into agent threats, vulnerabilities, and attack paths, with the ability to detect, investigate, and block risky behavior in real time.

Which Agent 365 capabilities are included with Microsoft 365, and which require an Agent 365 license?

Some Agent 365 capabilities are included with a Microsoft 365 plan and available at the tenant level, including:

- The Agent 365 Registry for visibility into agents
- Limited usage insights
- Content search

How do we help security teams understand how Entra, Purview, and Defender integrate to provide a unified control plane for agents?

Agent 365 serves as a unified control plane that extends Microsoft's core security and identity stack to the agentic workforce. Rather than introducing a siloed security model, it embeds agent management directly into the existing workflows of IT and Security teams, ensuring that digital workers are governed with the same rigor as human employees.

Microsoft Entra: Every agent is assigned a unique Entra Agent ID. This allows security teams to treat agents as first class identities, applying familiar controls like Conditional Access and least-privilege permissions to manage what resources an agent can access on a user's behalf.

Microsoft Defender: Defender provides a dedicated AI Agent Inventory to detect shadow agents and monitor for behavioral anomalies. It offers real-time protection against agent-specific risks, such as prompt injection or data exfiltration, by correlating agent activity with threat signals.

Microsoft Purview: Purview extends data protection policies to agent interactions. It enables Data Loss Prevention to block sensitive data in real-time and maintains a comprehensive Audit Trail of all agent-to-tool and agent-to-human communications for eDiscovery and regulatory compliance.

Agent 365 leverages the Registry as a single source of truth, feeding metadata into Entra for identity, Defender for threat protection, and Purview for data security and compliance, creating a unified governance experience.

What is the relationship across Agent 365, Microsoft 365, and M365 Copilot?

- Agent 365, M365 and M365 Copilot are separate offers that are complementary.
    - Microsoft 365: productivity and security for users
    - Agent 365: observability, governance and security for agents
- M365 and Agent 365 can work in tandem as organizations need to manage across users and agents.
- M365 and M365 Copilot are not a prerequisite to Agent 365.
- Agent 365 is not part of M365 or M365 Copilot.

How does Agent 365 compare with Foundry Control Plane and the Microsoft Security products?

As agents become integrated into the workforce, IT, developers, and security teams all play critical roles. At the core of these three teams is a shared security foundation that enables consistent identity, monitoring, and governance across the agent lifecycle. Security is the core primitive that powers every persona:

- **IT teams** must enable no and low code agent builders to innovate while being able to discover and manage agents across their environment. To ensure the responsible and secure deployment of agents into an organization, IT needs a unified agent registry and the ability to assign an identity to every agent. In addition, IT needs to be able to assign access to email and file storage. While monitoring their entire agent estate for risks such as over permissioned agents.
- **Development teams** need to build and test agents, apply security and compliance controls by default, and ensure models are evaluated for content safety and vulnerabilities. Post deployment, development teams must observe that agents are staying on task, accessing tools appropriately, and running with the appropriate cost efficiency and performance.
- **Security & compliance teams** must ensure overall security and compliance across their entire estate, including their AI platform, apps, and agents. Disparate tools and policies for AI create additional risk by reinforcing silos within the environment and across teams within the organization. Security teams need solutions that extend their existing security investments with purpose-built capabilities for AI platforms, apps, and agents. These capabilities should be built into the AI platforms and the tools that use them to enable secure and trustworthy AI innovation.

To deliver trustworthy AI, these teams need tools that are seamlessly integrated into their everyday platforms and that rely on the same consistent security data, alerts, policies, and controls. We call these 'shared security primitives and they are at the heart of what enables observability (security, management, and monitoring) at every layer of the stack.

Microsoft uniquely enables organizations to drive observability into every layer of the stack, backed by shared security primitives.

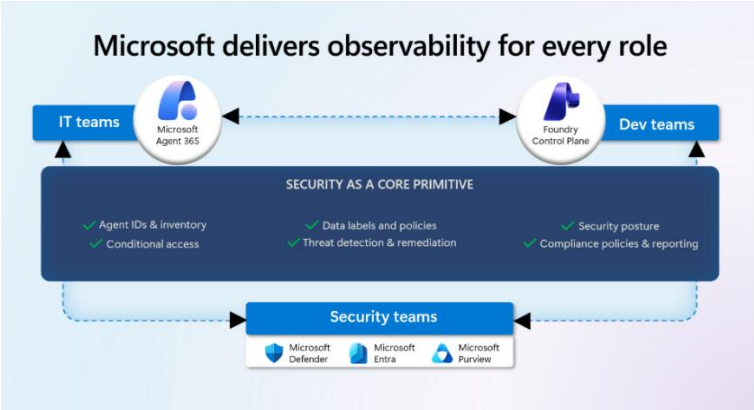

In simple terms, **Foundry Control Plane** is aimed at developers (it helps observe and govern agents during development and testing), whereas **Agent 365 is the**

IT admin's control plane for agents in production. As agent innovation accelerates, organizations need both capabilities – developer-side control *and* enterprise IT oversight. Microsoft is designing these to be **"better together"**. Agent 365 extends the familiar **Microsoft 365 Admin Center** to let IT admins discover, manage, and secure all agents alongside user accounts.

Meanwhile, Foundry Control Plane focuses on the developer observability required for AgenticOps. It brings controls, observability, security, and governance into one connected experience so developers can build, operate, and manage every agent in one place. Foundry Control Plane connects the entire lifecycle from development to deployment to oversight, allowing teams to move from alert to fix without losing context or flow.

They both share core security primitives – Agent ID & inventory, conditional access, data labels & policies, threat detection & mitigation, security posture, and compliance policies & reporting. For example, Agent 365 and Foundry Control Plane both enable Entra Agent ID. That means an agent created via Foundry is automatically registered with an Entra Agent ID and will show up in Agent 365's inventory automatically. Security signals and policies are also unified – e.g. a threat related to tool usage detected in a Foundry-built app managed by Agent 365 can raise an alert in Microsoft Defender.

- Agent 365 = governance, registry, security, oversight for all enterprise agents
- Foundry Control Plane = optimized for developer persona building AI-native agents
- Both offer complementary, not competing, controls.

What are the core benefits for customers adopting Agent 365?

Agent 365 helps organizations confidently scale AI agents across the enterprise by providing a unified control plane to observe, govern, and secure agents—treating them with the same rigor as users, apps, and devices. Together, these capabilities reduce risk, prevent sprawl, and enable faster, more sustainable agent adoption.

Observe: Monitor and manage agents in real time

- Gain visibility into every agent in your environment, understand how they're used, and act quickly on performance, behavior, and risk signals before they impact the business.

Govern: Establish guardrails for agents and users

- Establish guardrails for agents and people, onboard agents with IT oversight, and govern agent access to resources and data. Be audit ready with built-in compliance and data retention.

Secure: Protect agents comprehensively

- Secure agent identities, control access to resources, prevent data oversharing and leaks, and defend against threats and vulnerabilities with enterprise-grade security solutions.

How is Agent 365 positioned in the marketplace?

Agent 365 is positioned as the control plane for AI agents, a category that is just emerging. In the market, there are a few tangential offerings (some startups focusing on AI governance, a few cloud providers' AI management tools), but Agent 365 is somewhat unique. The key positioning points are:

- **Extension of Microsoft 365:** Rather than being a standalone AI tool, Agent 365 is an **extension of the existing Microsoft 365 platform**. This means customers don't need to adopt a new silo for agent management – it's built into the environment they already use for users and devices. This gives customers an advantage: if you trust M365 to manage your people and data, you can trust it to manage your AI agents too.
- **Extension of Microsoft Security:** Trustworthy agentic AI requires security and governance built into the fabric of IT and development platforms. Agent 365 is part of that story: it uses **Entra (identity)**, **Purview (data security and compliance)**, and **Defender (security)** all together. Competing offerings often cover one dimension. **Only Microsoft** can offer an agent solution that is so deeply integrated across identity, security, data security, and compliance,
- **Bridging AI and Enterprise IT:** Agent 365 is positioned to **bridge the gap between AI development and enterprise IT governance**. It's the answer to CIO/CISO concerns around AI: "We have these Copilot and GPT-based agents coming – how do we control them?" Microsoft's answer is Agent 365. In essence, we position it as **the missing puzzle piece** for organizations to safely embrace AI agents at scale. Without it, there's chaos; with it, you bring order.
- **Part of the broader Copilot/AI strategy:** Agent 365 is complementary to Copilot. Where Copilot provides AI to end-users, Agent 365 provides management of AI agents built by the organization or brought into the organization. As Microsoft talks about **Copilot (for users)** and **Foundry (for developers)**, it will talk about **Agent 365 (for IT and security)** – framing a complete story that spans all

personas. This positioning shows Microsoft has a holistic approach to Enterprise AI (not just throwing AI at end-users but also giving IT and Security teams the tools to manage it).

When talking to customers, we'd say: *"Just as Microsoft 365 E5 is the leader in securing your users and data, Agent 365 will be the leader in securing and managing your AI agents. It's a natural extension of your Microsoft ecosystem, not a bolt-on from another vendor."*

When and where will Agent 365 be generally available?

Agent 365 will be generally available on May 1. At GA, Agent 365 will be available in all markets supported by Microsoft 365 today (246 countries and regions) and will support all applicable Microsoft 365 languages.

What types of agents are in scope of GA?

When Agent 365 becomes generally available (GA) on May 1, it will cover the On Behalf Of (OBO) agents. For other types of agents, observability, security, and governance capabilities will remain in preview.

What capabilities are available today in Agent 365?

- **Observability**
  - **Agent Registry** (Register and track every agent) - Get a complete view of all agents in your organization, including agents built with Microsoft AI platforms, agents from our ecosystem partners, and any agents you register yourself.
    - **Product**: Agent 365 portal (MAC)
  - **Agent Analytics** (Measure agent effectiveness) **-** Track agent performance, speed, quality, business impact and ROI to make informed decisions.
    - **Product:** Agent 365 portal (MAC)
  - **Agent Map** (Visualize agent usage and behavior) - Visualize how agents fit into the broader ecosystem, connect with other agents, and perform over time to simplify monitoring and accelerate issue identification.
    - **Product:** Agent 365 portal (MAC).
  - **Role-based Oversight** (Extend role-based reporting) **-** Extend agent oversight to Security Leaders to manage agent risks and Business Leaders to monitor business metrics and ROI.

        - **Product:** Agent 365 portal (MAC), Defender, Entra and Purview portal for Security actions, Teams and M365 Copilot for end-user experiences.
- **Governance**
    - **Agent onboarding** (Bring agents under control from day one) - Onboard agents through IT-controlled workflow, applying security policy templates so every agent starts secure, governed, and compliant.
        - **Product**: Agent 365 portal (MAC) and Entra Governance
    - **Access and integration governance** (Control what agents can access and do) - Enforce least-privilege access by controlling which users, data, and tools or MCP servers agents can use. Limit access to only the resources and other agents they need.
        - **Product**: Agent 365 portal (MAC), Entra Governance, Agents tools gateway
    - **Lifecycle management** (Automate on-going agent governance) - Leverage rules-based agent management to automatically enforce lifecycle policies, such as sharing agents with specific users and groups, expiring inactive agents, flagging and re-assigning ownerless agents.
        - **Product**: Agent 365 portal (MAC) and Entra Governance
    - **Audit and logging** (Capture and trace agent activity) - Strengthen visibility and traceability of agent actions and interactions with logging, reporting, and audit capabilities.
        - **Product**: Agent 365 portal (MAC) and Purview
    - **Data compliance** (Meet AI regulations and policies) - Establish content safety controls to detect, retain, and investigate unethical agent interactions.
        - **Product**: Purview
- **Security**
    - **Access control** (Protect agent identities and access) - Protect agent identities and prevent breaches by extending conditional access and internet traffic filtering policies from users to agents.
        - **Product**: Agent Tools Gateway (ATG), Entra, Purview
    - **Data security** (Prevent oversharing and data leaks) - Gain visibility into AI-related data exposure, protect the data agents create and access from oversharing, leaks, and risky behavior.
        - **Product**: Purview

- **Threat protection** (Defend against threats and vulnerabilities) - Protect agents from threats and vulnerabilities, and adversarial attacks. Detect, investigate and remediate incidents quickly, with visibility into attack paths.
  - **Product**: Defender

What is the Agent Registry?

The Agent Registry is the central inventory of all AI agents operating in your tenant, providing a single authoritative view of every agent—Microsoft-built or third-party, approved or shadow. It stores rich metadata that makes agents discoverable, governable, and secure across your environment.

The registry provides visibility into:

- Agent identity
- Permissions & access rights
- Data access patterns
- Tools, plugins, and protocols the agent can use
- Usage patterns, activity trails, and interactions
- Risk posture and anomalous behaviors

Because it centralizes this information, the registry acts as the source of truth for Defender, Entra, Purview, Foundry, and Agent 365—enabling consistent governance, security assessments, and oversight across the entire agent estate.

What is the practical difference between registry-only agents and agents with Entra ID?

Registry-only agents provide visibility and inventory, such as discovery, naming, and basic metadata. Agents with Entra Agent ID enable identity-based governance, including Conditional Access, Identity Governance (access reviews, expiration, sponsor enforcement), and stronger lifecycle and policy controls. Bottom line: Registry provides visibility. Agent ID provides control.

How does Agent 365 discover shadow agents not built with Microsoft tools?

Discovery today is indirect and not identity based. It relies on Microsoft security signals from Defender for Cloud Apps, Defender for Endpoint, Entra Secure Web / AI Gateway, and network and data access telemetry. These agents appear as shadow activity, not full Agent IDs, and can be restricted at the data or network level.

Are discovered shadow agents automatically governed by Agent 365?

No. Discovery does not equate to governance. While Microsoft security signals can identify shadow agents active in your environment, these agents are not automatically enrolled in the Agent 365 management framework. To move a shadow agent from discovered to governed, it must be explicitly onboarded. There is currently no automatic conversion process; instead, an admin must:

- **Register the Identity**: Use a supported Microsoft agent platform or the Agent 365 SDK to provision a dedicated Entra Agent ID.
- **Formalize Oversight**: Once registered, the agent is recognized as a first-class citizen, allowing for full lifecycle management, Purview compliance auditing, and granular access controls.

What types of agents will Agent 365 work with?

The intent is that A365 work with all agents.

- **Supported agents**: Agent 365 value is provided for Copilot Studio, Foundry, and third-party agent via Agent 365 SDK within the Microsoft Agent Framework. Agent 365 will support Copilot Studio Lite, M365 Copilot first-party agents (Researcher, Sales, etc.) within the inventory, with more capability support to follow
- **Pre-existing agents**: Agents created prior may need retrofitting to support Agent 365. Identity engineering work is underway to minimize customer effort.
- **Partner ecosystem agents**: A curated list of first-party (1P) and third-party (3P) agents will be available for use with Agent 365.
- **Customize agents for Agent 365**: The Product Group will provide guidance on how to make agents compatible with A365.

Why does Entra Agent ID appear only for new agents?

Entra Agent ID is provisioned at the time an agent is created on supported Microsoft platforms. As a result, newly created agents automatically receive an Agent ID. Agents

created prior to Entra Agent ID integration cannot yet be retroactively backfilled. Support for backfilling existing agents is on the engineering roadmap.

Do all agents automatically receive an Entra Agent ID? Is Agent ID tied to licensing?

Entra Agent ID assignment depends on how the agent is created and onboarded. Agents created on supported Microsoft platforms receive Agent IDs as integrations roll out. Third-party or custom agents must be onboarded through the Agent 365 SDK. An Agent ID is an identity construct, not a license. Licensing applies when an autonomous agent-instance executes work.

What are the technical requirements for an agent to be 365-enabled?

To meet the minimum bar of A365 enablement, an agent needs to be part of the Agent 365 registry and needs to be observable. Not all agents will require advanced data and tools. As a result, integration with data and tools is optional.

- [Required for A365] Agent identity: For agents built on 3P creation platforms, this means integrating with Agent 365 SDKs. For agents built on Microsoft's agent creation platforms, this integration will be covered by default. Most of these agents will part of the A365 registry, with the remaining few joining in some months.
- [Required for A365] Agent traceability, governance, & security: for agents built on 3P creation platforms, this means adopting the A365 Observability SDK. For agents built on Microsoft agent creation platforms, this integration will be covered by default and will roll out over the next weeks.

Can customers bulk assign Entra Agent IDs to existing Copilot Studio agents?

No. There is no bulk or automated way today to retroactively assign Entra Agent IDs to existing Copilot Studio agents. Entra Agent IDs are created at agent creation time for supported platforms. Existing agents must be rebuilt or republished to receive an Agent ID.

How do I retroactively enable my existing agents for Entra Agent ID?

Currently, there is no retroactive process or automated migration to enable Entra Agent ID for agents created prior to GA. Agents built prior to GA do not automatically inherit the new identity primitives.

Can customers register or govern custom MCP servers in Agent 365 today?

No. Agent 365 supports Microsoft-provided MCP servers with limited enable/disable controls. Custom MCP registration and governance are not available today. For custom MCP endpoints, AI Gateway provides network-level enforcement, while Agent 365 focuses on agent-centric governance (identity, lifecycle, policy).

Are MCP servers free?

Today, MCP usage is gated by Copilot and Agent 365 licensing (USL/ASL). There are discussions about future pay as you go access for non-licensed users.

Are these MCP servers built on Power Platform connectors?

Not the ones built for Agent 365. They rely directly on Microsoft Graph, Copilot APIs, and the Model Context Protocol. Some future MCPs may reuse PP connectors.

Does Agent 365 support the Copilot APIs (Chat, Retrieval, Search)?

Yes. MCP servers in A365 use Copilot Chat APIs and will additionally support Retrieval and Search APIs as they're made available.

How is MCP server governance handled - Agent 365 or AI Gateway?

Agent 365 provides agent-centric governance including identity, lifecycle, and policy controls. AI Gateway-related controls provide network-level enforcement for MCP traffic. Agent 365 supports Microsoft-provided MCP servers with limited enable or disable controls; a dedicated MCP registry is not available today.

What is the future path of Agent 365 Registry and Entra Agent Registry?

Agent 365 Registry and Entra Agent Registry have now converged into a single, unified inventory. Agent 365 Registry serves as the central system of record for agent metadata and management. While the two registries operated separately during early development, they now function as one shared inventory, providing a consolidated view of all agents across the environment

How do agents appear in the Agent 365 Registry or Microsoft Admin Center?

Agents enrolled in Agent 365 will appear in your directories and admin portals in a manner similar to users or service principals, with clear indicators that they are "Agents." Here's what to expect:

- **In Microsoft Entra admin center**: When an agent is given an **Entra Agent ID**, it essentially creates an identity record in Microsoft Entra. Depending on implementation, this might show up as a special type of **Service Principal or Application** or a new object class for Agents. For example, if your agent is named "Contoso Sales Bot", Microsoft Entra might list "Contoso Sales Bot – Agent" with an icon or description indicating it's an agent. Under the covers, the **Object Type** could be something like "Application (Agent)" or an enterprise application representing the agent's logic.
- Additionally, Microsoft Entra will **store attributes** for the agent: e.g., "**Agent sponsor**" might be an attribute pointing to a user; "**Agent expiration date**" if set; etc., to facilitate governance. From Microsoft Entra's perspective, you manage this agent identity as a special identity type – you can reset credentials (if any), disable sign-in, view its sign-in logs, assign it to groups, etc. This will be available to identity admins in dedicated Agent ID overview blade in Microsoft Entra admin center.
- **In Microsoft 365 Admin Center**: There will be a dedicated area for agents. In early access, agents appear in a new **"Agents" section under Users** (or parallel to Users/Groups). The UI lists all agents by name with some properties (owner, status, last activity, etc.). Admin Center might also show agents in certain relevant lists: for instance, in Exchange Admin Center, you might see an agent if it has a mailbox (some agents might need a mailbox to send/receive emails). But primarily, expect a distinct listing separate from the user list.

When you click on a particular agent in Admin Center, you get a detail pane similar to a user detail pane. This includes information like:

- **Agent Name** (and perhaps a generated email or ID).
- **Agent ID (Client ID):** a unique identifier.
- **Sponsor/Owner**: which user or team is responsible for it.
- **Status**: Active, Suspended, etc.
- **License Assignment** (if Agent 365 requires a license per agent, it might show if a license is assigned).
- **Activity summary**: e.g., "X actions in last 24h", "Last action: 2 hours ago".

- **Policies applied**: e.g., it might list any conditional access or compliance policies specifically targeting this agent.

Essentially, this detail page in Admin Center becomes an **Agent Profile**. Admins can take actions from here like "Disable agent" (which would block its sign-in and activity), "Reassign Owner", "View Audit Log" (a button that deep links to the audit logs filtered for this agent).

- **In Microsoft Entra Org Chart / People experiences**: Agent 365 will also integrate with the Org chart in tools like Delve or Viva or Teams @ mention directory. This implies that if an organization chooses, they could represent agents in their internal org chart (perhaps under an owner or in a separate "AI Agents" department). Initially, this might not be auto-populated, but an admin could place agents in an org structure for visibility. For example, an agent could have a defined "Manager" field pointing to the sponsor. Then if an employee looks at the sponsor's org chart in Delve, they might see the agent as a subordinate (this is speculative but aligns with the concept of "digital workers" in org structure).
- **Identification**: Anywhere an agent appears, it's clearly identified as such so as not to be confused with a user. This could be via an icon or a suffix in the display name, or a special context label. For instance, in Teams, if an agent posts a message, its name is shown with a small badge indicating "Agent 365 agent" or such, to inform users. Similarly, in audit logs, the actor's user principal might include something that denotes it's an agent.

One can think of it like how we manage service accounts or shared mailboxes today – they appear in our consoles with slight differences – similarly, agents are a new type of identity/object that appears throughout the admin ecosystem. Also, Microsoft Entra Agent ID new features will include an **"Agent registry"** where all Agent IDs are registered and can be searched, plus all agents without an identity, including shadow agents, can be visible to identity administrators.

Operationally: to find an agent in Microsoft Entra, an admin will get a full overview of their agent identities in a dedicated Agent ID overview blade. In M365 Admin, they go to the Agents section and see it directly. Agents do not consume some UI places (like they won't show up in Exchange GAL unless given a mailbox, etc., and even if so, you might hide them from GAL as they aren't people). But they are **present in all the IT management surfaces** for full visibility and control.

How do inline or child agents created in Copilot Studio appear in the Agent 365 registry?

Agent 365 registry currently prioritizes Top-Level visibility. This means only primary or connected agents are surfaced, while nested logic remains under the parent's umbrella. If an agent is built to function solely within a parent agent, it is not visible as a standalone entry in the registry today.

Why do customers see agents or app in the Admin Center that do not appear in the Agent 365 Registry?

The Agent 365 Registry shows agents recognized for agent-specific governance and lifecycle management. Broader Admin Center views may include apps or services that are not Agent 365-managed agents. Only agents explicitly onboarded into Agent 365 participate in agent-level governance workflows.

Is there a dedicated "agent owner" or "agent boss" experience for business users?

No single new portal exists. Ownership and oversight are distributed by role:

- **IT Admins**: Microsoft 365 Admin Center (Agent 365)
- **Security Teams**: Entra, Defender, Purview
- **Agent Owners / Managers**: Teams activity and interaction history
  This approach meets users where they already work rather than introducing a new standalone console.

How does an IT Admin determine when their agents need Agent 365?

Not every simple agent may need the full Agent 365 capabilities, so IT admins will assess agents based on certain attributes and risk factors to decide if they should be onboarded into Agent 365. Essentially, **any agent that is empowered for actions and/or interacts with important corporate resources or multiple users should be under Agent 365**. Here are guidelines an admin can use:

- **Does the agent access corporate data beyond a single user's scope?** If yes, it needs Agent 365. For example, an agent that only helps a single user summarize their personal emails might not need its own identity (it works under that user's context). But as soon as an agent accesses shared/org data (SharePoint files, Teams channels, etc.), you'd want it in Agent 365 so you can control and audit those accesses.

- **Does the agent perform autonomous actions on behalf of the organization?** If an agent can take actions that affect systems (like sending emails, posting messages, modifying records), that's a candidate for Agent 365. The more autonomous an agent is (acting without direct user initiation each time), the more it should be governed.
- **Is the agent used by multiple people or a whole department?** Multi-user or organizational agents belong in Agent 365. E.g., a "Finance Report Agent" used by the finance team or a HR Q&A bot accessible to all employees – those should be under central management.
- **Does the agent handle sensitive information or critical processes?** Agents involved in regulated data (PII, financial data, IP) or critical business processes (like approving expenses or running build pipelines) should be under Agent 365's security blanket. If an agent touches anything you'd classify as sensitive if a user did it, treat it the same and put it in Agent 365 so DLP, audit, etc. apply
- **Was the agent built outside of IT's direct supervision?** If business units or citizen developers are creating bots via Copilot Studio or other platforms, those agents should be pulled into Agent 365 when they move from "experiment" to "production use." Essentially, Agent 365 is the safety net once something goes beyond personal experimentation.
- **Does the agent have persistent logic or "life" of its own?** For example, a ChatGPT plugin that only responds in a user's session vs. an agent service that runs continuously. The latter (persistent agent) needs an Agent 365 identity to track its lifetime. The former might end with the session (though even then, if it accesses corporate data, caution).
- Admins can ask: *If this agent malfunctioned or was compromised, could it do harm?* If the answer is yes (it could leak data, make unwanted changes, confuse users with bad info), then it should be managed under Agent 365 to mitigate those risks with monitoring and controls.

As a baseline, **any agent integrated into official business workflows or customer-facing functions should be in Agent 365**. Whereas a trivial bot that just formats data locally might not warrant the overhead.

In practice, an IT admin might maintain a checklist or rubric:

- **Scope of data** (personal vs org-wide),
- **Criticality of tasks** (advisory vs decision-making),
- **Degree of autonomy** (fully autonomous vs always user-triggered),
- **Volume of use** (one-off bot vs frequently used service).

If an agent scores high on those dimensions, it needs Agent 365.

How do Admins monitor and manage active agents in Agent 365?

Admins will use Agent 365's unified control panel (integrated into Microsoft 365 admin experiences) to continuously **monitor and manage** all active AI agents. Key tasks and tools include:

- **Agent Inventory & Registry**: Admins have a dashboard listing every registered agent, with key metadata (name, owner, last active time, etc.). From here, they can select an agent to see details. This inventory view is the starting point for monitoring – it shows which agents are currently active in the organization.
- **Real-Time Status and Health**: In the admin portal, agents have status indicators (Active, Suspended, Error) and health metrics. For instance, if an agent hasn't performed any action in a long time, it could be flagged as dormant. Or if an agent triggered a security alert, it might be flagged. The admin can quickly identify any agents requiring attention.
- **Activity Logs and Audit Trails**: For supported agents, admins can understand agent interactions with resources and other agents. Additionally, audit logs are sent to Purview for visibility into Data Security Posture Management, Audit, and Compliance Manager.
- **Security Dashboards & Alerts**: Across Microsoft Defender, Entra and Purview, security admins get alerts if an agent does something suspicious or violates a policy. These alerts surface in security dashboards. Admins can monitor risk levels of agents (like a list of "High risk agents" if any) and respond accordingly. For example, if Agent A has unusually high data downloads, an alert might appear, prompting the admin to investigate that agent.
- **Policy Controls**: Admins manage active agents by setting and adjusting policies. In the Agent 365 interface, an admin can apply organizational policies that automatically govern agents (like requiring sponsor re-attestation every N days, or limiting which data certain agent roles can access). They can also set agent-specific policies (for example, limit Agent X to read-only mode if needed for a period). Active agents are continuously evaluated against these policies and administrators can tweak them as necessary from the portal.
- **Admin Actions (CRUD for Agents)**: For management, admins can perform actions on agents:
    - **Create/Onboard New Agent**: When a new agent needs to be introduced, the admin can onboard it (assign identity, license, etc.) via Agent 365.

    - **Update Settings**: For an active agent, admins can update its attributes – e.g., change its sponsor/owner if the responsibility shifts, adjust its permission levels (i.e. put it into an AD group that gives more or less access).
    - **Suspend/Pause Agent**: If an agent is misbehaving or is not needed temporarily, an admin can disable its account or pause its operation. This might be a one-click "Disable" button that stops the agent's credentials from working (so it can't log in or take actions until re-enabled).
    - **Retire/Remove Agent**: If an agent is decommissioned, the admin can delete its account (and optionally archive related data). This is analogous to offboarding a user – typically removing licenses, access, and then deletion with content preservation.
    - **License Allocation**: If Agent 365 uses licensing, admins manage license assignment to agents just like for users (though current info suggests they haven't finalized pricing model – but if an agent consumes a license or meter, admin would allocate those and monitor use).
    - **Monitor Usage Metrics**: The admin interface or a PowerBI report might show metrics like number of tasks completed by each agent, response times, etc. Admins (or business owners) can monitor these to gauge performance and ROI of agents.
- **Integration with ITSM**: Admins can incorporate agent monitoring into their IT Service Management processes. For instance, if an agent fails (the underlying service goes down or it encounters an error it can't handle), that could trigger an alert in the monitoring system, and an admin might get a ticket or notification. Admins then use Agent 365 tools to troubleshoot – checking logs, re-running tasks, or contacting the agent's sponsor to adjust something.
- **Compliance Oversight**: Admins (and compliance officers) manage active agents by periodically auditing them. E.g., running an eDiscovery search to see what content an agent has generated or accessed in the last quarter, ensuring it aligns with compliance expectations. Agents may be subject to compliance reviews just like how you'd review user access or communications.
- **Separation of Responsibilities**: Some organizations might designate an "Agent Administrator" role who specifically focuses on Agent 365. That admin (or team) would have the job of daily/weekly checking on agent statuses, reviewing any anomalies, and optimizing the agent settings. The Agent 365 portal will support role-based admin so someone can be an Agent admin without being a full tenant admin, etc. That helps delegate the operational monitoring to the appropriate team (i.e. the automation CoE or IT ops team).

In everyday practice: an admin logs into the Microsoft 365 admin center each morning, glances at the **Agents section to see all green** (meaning all agents are functioning within norms). Perhaps one agent shows a warning icon (it's approaching an expiration date or triggered a DLP block recently) – the admin drills into that agent's details, sees the issue (like "tried to email sensitive file, blocked by DLP at 5pm yesterday"), and can follow up (remind the agent's sponsor to update its training or restrictions). If a new request comes like "Team X wants to use a new agent," the admin onboards it via Agent 365, so it appears in this monitoring list from then on.

Where do we manage agents – MAC, Foundry, or Power Platform Admin Center?

- Microsoft Admin Center (MAC) is the centralized control plane for managing agents across the organization. Agents built in Azure AI Foundry and Copilot Studio are surfaced and governed in MAC, providing a unified management experience.
- Azure AI Foundry and Power Platform Admin Center (PPAC) are environments for building and configuring agents within their respective platforms. However, MAC serves as the cross-platform layer for centralized visibility, governance, and lifecycle management across all agent types.

How does Agent 365 handle agent lifecycle (activation, suspension, retirement)?

Agent 365 provides full lifecycle management for AI agents, from their initial activation through suspension and eventual retirement:

- **Activation/Onboarding**: When an agent is first created or brought into the organization, Agent 365 "activates" it by assigning it an identity (Agent ID) and applying initial policies and configurations. This is analogous to onboarding a new employee – setting up their account and permissions. Admins (or automated workflows) can mark the agent as Active once all prerequisites (like approvals, sponsor assignment, license assignment) are done. At activation, the agent might be issued credentials or API tokens that allow it to start operating under supervision. Agent 365 logs the activation event and the agent now appears as an active entity in the environment.
- **Monitoring During Active Use**: Once active, the agent remains in service performing its functions. Agent 365 continuously monitors it as discussed (audit logs etc.). The lifecycle can include periodic attestations – for instance, every 90 days the agent's sponsor might need to confirm that the agent is still needed and working properly. If the sponsor doesn't re-attest by the due date, Agent 365 could

flag the agent for suspension. This is a preventive lifecycle step to ensure no "set and forget" beyond allowed time.

- **Suspension**: If an agent needs to be paused (it's malfunctioning, or a business scenario ends, or security found an issue), an admin can suspend the agent. Suspension in Agent 365 means disabling its Entra ID (so it cannot authenticate or get access tokens) and halting any scheduled tasks the agent performs. The agent's status changes to Suspended. While suspended, the agent essentially goes inert – it won't respond or act. However, all its data (identity, configuration, logs) remain intact in Agent 365. This is similar to disabling a user account (the user can't log in but their mailbox and files remain). Agents can be suspended manually or automatically (via policy triggers). For example, if an agent's expiration date passes with no renewal, Agent 365 could auto-suspend it. Or if the agent's sponsor left the company and no replacement sponsor was assigned in a given time window, policy might suspend the agent until governance is sorted out.
- **Resumption/Reactivation**: If conditions allow, a suspended agent can be reactivated (unsuspended). For instance, after an investigation clears the agent of wrongdoing, or after updating the agent's programming to fix an issue, an admin can click "Resume" and re-enable its credentials. The agent picks up operation again. Agent 365 will note the suspension period in logs for auditing.
- **Retirement/Decommissioning**: When an agent is no longer needed permanently, it should be retired. Retiring an agent in Agent 365 would involve:
    - **Revoking its access** – ensure all credentials/tokens are invalidated so it can no longer function.
    - **Archiving or transferring any outputs/data** – if the agent had a mailbox or storage, decide whether to preserve that data (e.g., export chat logs, keep documents it created).
    - **Removing it from inventory** – an admin marks it as Retired in Agent 365, and then likely deletes the agent's Entra ID and any license assignment. Agent 365 could keep a tombstone record (for audit history, an entry that agent existed from date X to Y).
    - **Notifying relevant parties** – i.e. alert its sponsor or interested users that the agent is being retired and will no longer respond.
    - After retirement, the agent no longer appears in the active agent list. It's essentially gone as an entity that can act in the tenant. If later a similar agent is needed, a new one would be onboarded – you wouldn't typically "unretire" because deletion is final (whereas suspension is the reversible state).
- **Orphan Handling**: Part of lifecycle is dealing with orphaned agents gracefully. If an agent's sponsor (owner) leaves the company or changes roles, Agent 365 can notify

admins and/or the sponsor's manager to reassign the agent's sponsorship. If no sponsor is found, the policy might trigger automatic suspension or eventually retirement for that agent (since no one is accountable for it).

- **Change of Ownership/Transfer**: If an agent changes hands (say a project's ownership moves to another department), Agent 365 supports updating the agent's metadata. The agent continues operating but under new oversight. This is akin to transferring an asset to a new department.
- **Audit Trail in Lifecycle**: Every stage change (Activated -> Suspended -> Resumed -> Retired) is logged by Agent 365 with timestamp and reason. This provides an audit trail for compliance to show that agents are being managed systematically.

Does Agent 365 automatically expire or clean up unused agents?

Agent 365 supports manual lifecycle actions such as assigning sponsors, suspending agents, and retiring agents. Policy-based lifecycle automation such as expiration warnings or inactivity-based cleanup is planned and will roll out incrementally.

Can external users email or contact an agent?

Yes — external users can email or contact an agent, but only if your organization's policies allow it. By default agents behave as internal identities, and external access is controlled or restricted through admin settings.

What happens if an agent is deleted or deactivated?

The associated license is released and becomes available for another agent instance.

Can customers prevent "shadow agents" or unapproved agents?

Yes—customers will be able to require that only registry-approved agents can run, block or restrict unregistered agents, and use policies to enforce approval before deployment, with preview already providing visibility and enforcement expanding over time.

How does Agent 365 support onboarding, interoperability, and governance for third-party and custom agents?

Agent 365 brings third-party and custom agents under the same identity, security, and compliance framework as native agents by:

- **Onboarding**: Assigns an Entra ID and registers the agent in the catalog for identity and compliance.

- **Interoperability**: Uses SDK, Graph APIs, and connectors so external agents can integrate, authenticate, and exchange telemetry.
- **Governance**: Applies the same security, compliance, and audit policies uniformly

Can Agent 365 govern agents built outside Microsoft (e.g. ChatGPT Enterprise, AWS, GCP)?

Yes—if the agent is onboarded using the Agent 365 SDK. SDK integration enables inclusion in the Agent 365 Registry, observability telemetry, and policy enforcement via Entra, Defender, and Purview. Agents not integrated via the SDK are not governed at the agent identity level.

External or SAAS agents that are not onboarded via the SDK cannot be governed at the agent identity level. However, they can still be discovered as shadow agents through Microsoft security signals such as Defender and Entra. Discovery allows organizations to detect and restrict activity, but full governance requires explicit onboarding and registration in Agent 365.

What insights can we show for third-party agents?

When a third-party agent is onboarded through the Agent 365 SDK and assigned an Entra Agent ID, it appears in the Agent 365 dashboard. From there, admins can see:

- Active users and departments
- Frequency and usage trends
- Lifecycle status
- Agent identity and ownership details

This allows security and IT teams to distinguish approved agents from unsanctioned shadow agents and take action directly in the admin center

For deeper content and threat insights, Agent 365 extends into the broader Microsoft security ecosystem. Through the Microsoft Purview SDK, organizations can audit prompts and responses and apply DLP controls to redact sensitive information before it leaves the agent.

Integration with Microsoft Defender enables monitoring and detection of risks such as prompt injection or suspicious data access by triggering automated alerts and access blocks at the runtime layer.

What is the Agent Framework?

Agent 365 SDK is part of Microsoft's Agent Framework and provides developers with tools to build, integrate, and manage intelligent agents across Microsoft 365 services. It enables:

- Multi-channel deployment: Agents can run in Microsoft 365 Copilot, Teams, web apps, and custom channels.
- Agentic patterns: Offers scaffolding for state management, activity handling, and orchestration.
- AI flexibility: Works with Azure OpenAI, Semantic Kernel, or other AI services without locking you into a single stack.
- Security & Identity: Built-in agent ID and integration with Entra registry for unified inventory and observability, and

Core capabilities are:

- Provides a container model for agent state and orchestration.
- Supports authentication via Azure AD app registration and secure endpoints.
- Offers development tools for C#, JavaScript, and Python, plus CLI and VS Code extensions.
- Enables testing in the Agents Playground and deployment through Azure Bot Service.

What is Agent Map?

Agent Map is a visual graph inside Agent 365 that shows how every agent in your organization connects and interacts with other agents. It groups agents into categories and reveals their relationships and dependencies, giving admins a clear picture of how the agent ecosystem actually behaves.

Can employees still use agents without registering them in Agent 365? How do we enforce only registered agents?

Today: Shadow agents can still run. Agent 365 discovers them, surfaces their activity, and flags unmanaged or risky agents — but it does not hard-block them yet.

In the future, on our roadmap: Policy enforcement will allow organizations to require that only agents registered in the Agent 365 Registry may operate, blocking unregistered or shadow agents from running or accessing protected resources.

What value does Agent 365 SDK enable?

- **Enterprise Integration**: Agents can have an identity in Microsoft Entra ID, aligning with corporate security and compliance policies.
- **Unified Development**: Build once and deploy across multiple Microsoft 365 channels without rewriting logic.
- **Extensibility**: Connect to Microsoft Graph, external APIs, and any AI service (Azure OpenAI, OpenAI, or custom models).
- **Governance**: Maintain enterprise-grade security

Can Foundry or Copilot Studio agents be published into Agent 365 without CLI or SDK steps?

Onboarding Foundry or Copilot Studio agents into Agent 365 currently requires CLI or SDK-based steps to enable identity, observability, and governance. Simplifying onboarding is an active engineering focus, but CLI or SDK-based onboarding is required today.

Does Agent 365 require Entra ID for agent identity and can external identity providers (e.g., Okta) be used?

Microsoft Entra ID is a fundamental requirement for the Agent 365 identity framework. Every managed agent must be assigned an Entra Agent ID to be discovered, governed, and secured within the Microsoft 365 ecosystem. However, organizations using external Identity Providers like Okta can still integrate with Agent 365 through standard federation patterns.

Does an external agent (e.g. AWS-built agent) with an Amazon identity federated with Entra ID still need Agent 365 registration?

Yes. While identity federation allows an AWS Bedrock agent to authenticate using Entra ID, it does not automatically "enroll" the agent into the Agent 365 management framework. To achieve governance, observability, and compliance, the agent must be explicitly registered as a managed entity

Federation manages the connection between AWS and Entra, but Agent 365 registration creates the Digital Employee Record. Without this, the agent remains a shadow agent rather than a governed agent.

Registration through the Agent 365 CLI provisions a dedicated **Entra Agent ID**. This ID is the piece that enables Agent 365 to:

- Surface the agent in the Registry for administrative oversight

- Apply Purview data protection and audit logging to the agent's actions
- Trigger Defender threat protections specific to agentic behavior

You should use the Agent 365 CLI to create an Agent Blueprint. This process links your AWS endpoint to a formal Entra Agent Identity, ensuring it is recognized as a first-class citizen within the Microsoft 365 ecosystem.

How are "risky" agents determined? Based on what model?

Risky agents are identified through Defender's observability signals, which look for patterns such as external sharing, unusual behavior, or other anomalous activity. These determinations are driven by Defender's built-in policies and defaults, which customers can configure to match their environment and governance needs. The approach is not based on or mapped directly to NIST or ISO AI-risk frameworks, but rather relies on Microsoft's security telemetry and behavioral indicators to assess agent risk in real time.

What telemetry does Agent 365 provide for agent performance?

Agent 365 provides rich telemetry on agent activities and some performance insights, with more analytics features coming. Out-of-the-box, every action by an agent is logged (who/what/when) in the Microsoft Purview Audit logs. This covers things like files accessed, messages sent, errors encountered, etc., ensuring a complete audit trail.

Beyond raw logs, Agent 365 is introducing usage dashboards and reports in the Microsoft 365 Admin Center. These will show metrics such as: number of tasks completed by each agent, frequency of agent use, active vs inactive agents, and policy violations or alerts triggered.

Some telemetry points likely available:

- **Agent engagement**: e.g. how many user queries or requests an agent handled, response rates, and resolution rates.
- **Agent efficiency**: e.g. average time an agent took to complete a certain request (if measurable), number of hand-offs to humans required.
- **Policy impact**: e.g. count of times an agent was blocked by DLP or had to elevate for approval.
- **Reliability**: e.g. uptime of agent (if it's supposed to run continuously), error rates (like how often it failed to fulfill a request).
- **Security signals**: e.g. agent risk score over time (if Defender flags risky behavior, that might be reflected as a trend).
- **Support & Escalation**

Does Agent 365 observability require Entra Agent ID?

No. Observability is independent of identity. Agents can emit telemetry using the Agent 365 Observability SDK without an Entra Agent ID. However, identity-anchored observability—correlating behavior, risk, lifecycle, and policy enforcement—is only possible once an agent has an Entra Agent ID.

Licensing & Eligibility

What is the business model and pricing structure for Microsoft 365 E7?

Microsoft E7 is priced at $99 per user per month. It can be purchased on a month-to-month, 1-year or 3-year license. The list price reflects the value of ME5, Copilot, Entra suites and A365 discounted as a bundle ($99 versus $117 for the standalone components cost – a 15% savings).

Do I need Microsoft 365 E7 to use Agent 365?

Microsoft 365 E7 is designed for organizations deploying AI broadly, where users, Copilot, agents, and enterprise security need to work together as a single, governed system. While ME7 provides the most complete experience for scaling agents securely across an organization, Agent 365 is also available as a standalone license,

Is Agent 365 included in Microsoft 365 E7?

Microsoft 365 E7 includes Agent 365. ME7 brings together user, Copilot, AI agents, and enterprise grade security, governance, and management to operationalize AI at scale. Agent 365 is the control plane within ME7 that provides centralized governance, security, and observability for agents, so they can run safely across the organization.

Will Agent 365 be included in Microsoft 365 E5 or E3?

Microsoft 365 E7 is the recommended option for organizations looking to scale AI broadly and securely. ME7 includes Agent 365 alongside Microsoft 365 E5, Microsoft 365 Copilot, and the Microsoft Entra Suite—bringing productivity, AI, identity, security, and agent governance together as a single, integrated system. This unified approach reduces complexity, ensures consistent controls across users and agents, and provides the strongest foundation for enterprise-wide AI adoption.

That said, Agent 365 can be purchased as a standalone SKU. This option is designed for organizations that want centralized visibility, governance, and security for AI agents without changing their existing Microsoft 365 licensing.

Can companies buy Agent 365 as a standalone SKU?

Microsoft 365 E7 is the recommended option for organizations looking to scale AI broadly and securely. ME7 includes Agent 365 alongside Microsoft 365 E5, Microsoft 365 Copilot, and the Microsoft Entra Suite—bringing productivity, AI, identity, security, and agent governance together as a single, integrated system. This unified approach reduces complexity, ensures consistent controls across users and agents, and provides the strongest foundation for enterprise-wide AI adoption.

That said, Agent 365 can be purchased as a standalone SKU for **$15 per month/per user**. This option is designed for organizations that want centralized visibility, governance, and security for AI agents without changing their existing Microsoft 365 licensing.

How is Agent 365 licensed?

Agent 365 is available in the new Microsoft 365 E7 license with list price $99/user/mo. It can also be available standalone as a per user subscription at $15/user/mo. The per user Microsoft 365 E7 or standalone Agent 365 license will cover all On Behalf Of (OBO) agents of a licensed user. There are no consumption-based/agent specific costs.

Is a Microsoft 365 E3 or E5 license a prerequisite for using Agent 365 to govern and secure agentic AI workloads?

No. While certain features, such as label inheritance, may require an E3 or E5 license to activate, it is not a commercial prerequisite. Agent 365 can be used to govern and secure agentic AI workloads without an existing M365 E3 or E5 subscription

Does an Agent 365 license grant E5-level Purview or Defender capabilities to customers currently on an M365 E3 tier?

No. An Agent 365 license does not grant E5-level security features to a user or elevate their seat's capabilities. It provides no additional user SCIM value, meaning it does not upgrade the underlying human user's identity or access rights to premium security tools.

If customers already have Microsoft 365 Copilot, do they need Agent 365 for full protection and compliance?

Microsoft 365 Copilot provides AI capabilities to end-users, but it does not fully govern agents – Agent 365 is needed to ensure those agents (and any beyond Copilot) are properly managed and secured. Think of it this way: **Copilot is about what AI can do; Agent 365 is about controlling what AI is allowed to do.**

How does Agent 365 licensing work for non-Microsoft agentic platforms (e.g. Manus, Genspark, BNY, GS proprietary frameworks)?

Third party agents require both:

- An Agent 365 license (for observability, governance, tooling access), and
- Their own native licensing (Marketplace-managed or Bring Your Own License).
- Agent 365 does not replace the 3P vendor's licensing.

Purchase Channels & How to Buy

What is the pricing structure for Microsoft 365 E7?

Microsoft E7 is priced at $99 and can be purchased on a Annual/Annual, Annual/Monthly, and Monthly/Monthly terms. The list price reflects the value of ME5, Copilot, Entra Suites and A365 discounted as a bundle. For a limited time, this list price is eligible for 10%, 15% 20% promotional offers for 10+, 100+, 1,000+ seat purchases, respectively on annual licenses. These promotional offers are available through the end of the calendar year.

Microsoft 365 E7 will be available to purchase through all channel: EA, CSP, and Web Direct.

Where can customers buy ME7?

Via Microsoft Admin center and through authorized CSP partners.

Is ME7 purchasable through CSP and visible in price lists/offer matrix?

Yes, ME7 is available through CSP starting 5/1, with pricelist available on Partner Center /offer matrix starting 4/1.

Mid-term Upgrades

If a customer already has a Microsoft 365 Copilot license and Productivity License, can combine to a single ME7 License?

Yes, customers can avail of a many-to-one mid-term upgrade inclusive of a Microsoft 365 Copilot plus a base Productivity License (e.g., ME3, ME5) to ME7.

Do equivalent upgrade paths exist for SKUs without Teams?

Yes, equivalent upgrade paths exist for SKUs without Teams.

Non-profit, EDU, Government

Is Microsoft 365 E7 available to different audiences (government, EDU, nonprofit?

At this point, ME7 will not be available in sovereign clouds (i.e. USGOV), EDU, or nonprofit at launch.

Promotions & Offers

What new promotions are available for Microsoft 365 E7 (with and without Teams)?

- **Annual (1-year) Term**
    - **10% off** 10+ seats
    - **15% off** 100+ seats

Who is eligible for these promotions?

- All customers transacting via CSP

What is the promo window and term?

- Valid **May 1, 2026 through Dec 31, 2026**
- Annual commitment required (A/A or A/M billing)
- There is no Monthly commitment option available for the promotions

What are the promo prices for each offer?

- **Annual (1-year) Term**
    - **10% off** 10+ seats
        - **Net Partner Price: $71.28 (ME7 with Teams)**
    - **15% off** 100+ seats
        - **Net Partner Price: $67.32 (ME7 with Teams)**
- Please see additional details in our Global Promo Readiness Guide (new offers to be reflected **4/1**)

Can customers add seats during the promo term?

Yes, customers can add seats anytime during the annual term at the promo price.

Which customers qualify for promotional pricing?

Promotional pricing is available to all customers who meet the minimum purchase requirements. Eligibility applies regardless of whether the purchase occurs during a mid-term upgrade or at renewal. Renewals qualify only if the subscription renewal occurs during the active promotional period and does not automatically apply at renewal. To receive the promotion at renewal, you must apply the promotion ID at the time of renewal purchase.

Are mid-term upgrades eligible for promotional pricing?

Yes, mid-term upgrades are allowed starting May 1. The commitment term may restart at the time of upsell.

Are these promos available globally?

Yes, they are available worldwide through CSP New Commerce Experience.

What are the partner incentives for ME7?

Incentive rates are the same (for Direct and Indirect providers) across Microsoft 365 E7 Copilot Business and the existing Microsoft 365 E SKUs, inclusive of base rate, strategic product accelerator, and growth accelerator (for first time purchase). *Incentive structures may change in FY27 based on business planning and Microsoft's priorities. Partners should check the Microsoft Partner Center and incentive guides for the latest details. Please see additional details on incentive rates in our Launch Kit*.