

FRONTIER.

AGENT AGENTS ENTITY

FRONTIER / NEW ANALYSIS

Microsoft Gives Copilot Studio Agents Their Own Identities. CIOs Inherit the Liability.

Microsoft's agent accounts, process memory, and proactive triggers move Copilot Studio toward persistent organizational labor. The real contest is who controls their authority, memory, and cost.

NEW ANALYSIS

Built to change a decision.

Microsoft's agent accounts, process memory, and proactive triggers move Copilot Studio toward persistent organizational labor. The real contest is who controls their authority, memory, and cost.

EDITION	New Analysis
AUTHOR	Frontier Research
PUBLISHED	October 20, 2026

This PDF is a publication rendering of the canonical Frontier News article. The web article remains the live source of truth; analysis, figures, and source citations travel with this edition.

What changed

Microsoft's October 15 Copilot Studio announcement crosses an important boundary: an AI assistant that answers an employee is not the same thing as an agent that belongs to a continuing business process. Microsoft is introducing agent accounts, business process memory, and event- or schedule-driven triggers intended to make agents durable participants in organizational work. The shift is not better chat. It is work delegated to software that can persist when its creator is offline—or has left the company.[1]

Agent accounts give these systems their own credentials and scoped permissions rather than borrowing a worker's. Employees can delegate via Teams or Outlook and @mention an agent in comments across Word, Excel, or PowerPoint. With a dedicated identity, actions can be attributed to an agent, and administrators can manage access and lifecycle. Microsoft is extending familiar collaboration patterns to a non-human process participant; its enterprise identity products become a potential enforcement point.[1][2]

Business process memory is designed to retain recurring instructions, process exceptions, and institutional preferences across cycles. Triggers in the GitHub Copilot harness can start an agent on an external event or schedule without a new prompt. Microsoft cites Ooredoo Qatar's executive business-case review agent, operating through Outlook, Teams, and Word, as an example. This is a vendor-described use case, not independently audited evidence of operating returns.[1]

Availability matters. Microsoft's announcement describes a rollout through its Frontier program. Additional administration and cost controls, automated agent refinements, and browser-enabled cloud-computer capability are described as planned or upcoming. The illustrative weekly budget-review agent is hypothetical. CIOs must distinguish generally usable features from controlled previews and roadmap claims before permitting unattended production work.[1]

Why it matters

The old chatbot inherits a person's identity, performs a bounded exchange, and largely stops when the session ends. A process agent has a durable identity, memory, triggers, and potentially a continuous workload. It becomes a machine principal capable of changing business state. Traditional SaaS governance is centered on named employees, application roles, human approval chains, and service accounts with predictable functions. It cannot automatically explain who owns an agent that alters its plan, receives hostile email, or executes after a manager changes roles.

The first failure mode is orphaned authority. Distinct agent credentials improve attribution but also create an expanding non-human workforce that needs provisioning, business ownership, least privilege, budget controls, access reviews, and decommissioning. Reading a forecast is not the same permission as notifying a supplier or initiating payment. Microsoft Agent 365 documentation correctly identifies distinct identity as a prerequisite for enforceable policies and traceability. An identity alone, however, does not establish a legitimate purpose or a durable delegation boundary.[2]

The second failure mode is memory becoming a shadow system of record. An agent may recall that a finance analyst waived an approval in July and reuse that exception in December after policy changed. It may mistake a suggestion, outdated document, or attacker-controlled instruction for durable operating guidance. Shared process memory multiplies a mistake across users. Every consequential remembered fact therefore needs provenance, sensitivity classification, retention, versioning, correction, and deletion. Microsoft's idea of designating authoritative records is important; buyers must demonstrate what happens when a remembered instruction conflicts with the latest approved record.[1][3]

The third is unreliable long-running execution. External events may arrive twice. A connector can time out after the receiving system has committed an action. Human approval might arrive days after a deadline. A successful retry must know whether the effect already happened. Google's Agent Executor work emphasizes durable event logs, snapshots, resumability, and interruptions. An agent without idempotency keys, checkpoints, approval expiry, and compensation steps can become an expensive duplicate-transaction machine.[4]

These features sit squarely inside the second wave examined in our analysis, *The Agents Have Left the Chat. Now They Want the Keys to the Company*. OpenClaw showed that a capable model combined with messaging, tools, persistence, and memory can behave like an operator. The enterprise response is to wrap that agency in policy, identity, isolation, and accountable execution. Microsoft is not alone: Oracle's Fusion Claw pairs reasoning with deterministic enterprise computation and outcome receipts; Google Cloud's October Gemini agent claims durable cloud work, coworker identities, dedicated mailboxes, and orchestration; OpenAI dots operate on cloud computers, with specialist enterprise pilots and plans to integrate identity governance through Microsoft Agent 365.[3][5][6][7]

Microsoft's advantage is powerful distribution: Entra identity, Microsoft 365, Teams, Outlook, Office documents, and Power Platform are already woven into enterprise routines. The same embeddedness is a lock-in risk when business-process memory, permissions, and approval history are encoded in platform-specific forms. The UK competition regulator is examining Microsoft's position in business software and questions about choice. That is an ongoing inquiry, not a finding of misconduct, but it demonstrates why portability and third-party interoperability belong in procurement diligence.[8]

The economics also move beyond a flat assistant subscription. Always-on work consumes models, cloud runtime, connectors, monitoring, approvals, failed retries, and sometimes external transaction fees. Our earlier analysis of Microsoft's metered Copilot economics called the seat a potential cover charge. Our AWS agent-stack coverage described the rush to own orchestration, governance, and commercial metering. The purchase metric should shift from cost per seat or token to cost per correctly completed, governed business outcome.[9][10]

Independent reporting from October's Oxford AI summit raised questions about agents undermining one another's protections and the adequacy of approval and logging systems. That is ecosystem evidence, not proof of a Copilot Studio flaw. It does make hostile input, cross-agent escalation, and complete action receipts mandatory tests before organizations scale delegated authority.[11]

Frontier take

The essential enterprise infrastructure is the delegation ledger, not another agent interface. A defensible record identifies the business goal, delegating owner, machine identity, current policy version, authorized tools, relevant evidence, approvals, spending, actions, and outcomes. The organization must be able to reconstruct that chain and revoke authority without relying on the model to cooperate.

That is the dividing line between the competing platforms. Microsoft approaches agency through workplace identity and collaboration. Google offers a cloud-persistent agent with context and model routing. Oracle roots autonomous execution in governed transactions. OpenAI gives agents broad cloud-computer capability while describing its Enterprise dots beta as disabled by default, with proactive background research read-only under its default protections. Different interfaces and runtimes are competing over the same elements: identity, state, policy, action, and audit.[1][6][5][7][12]

The catastrophic failure is not a slightly inaccurate paragraph. It is a machine that follows the wrong retained instruction while possessing valid credentials. Prompt injection, permission drift, stale shared memory, trigger storms, and unbounded spending become more dangerous when an agent can change systems of record. Our analysis of NVIDIA's agent security architecture argued that enforcement must move below the model and into an independent runtime boundary. The same principle applies here: the agent may recommend a supplier payment, but a deterministic control must decide whether it can initiate it.[13]

Microsoft deserves credit for making the organizational agent identity explicit. Yet the unanswered production questions are harder: can an administrator inspect and correct memory, stop a pending trigger instantly, prove why a decision was made, replay a task without duplicating external effects, and move a process when the model or vendor changes? Buyers should demand demonstrated behavior in their own tenant, not promises implied by the product family. Some relevant controls remain on Microsoft's roadmap.[1]

The Frontier thesis: persistent agents will turn identity systems into the operating systems of digital labor. The winning CIO will not have the largest agent roster. The winner will delegate the most valuable work under authority that remains measurable, auditable, reversible, and portable.

Three moves for CIOs

1. — Register each persistent agent as a governed non-human worker Create a machine-workforce register before permitting recurring work. For each Copilot Studio agent document its business owner, purpose, credential, approved tools and data, spending cap, runtime, access-review cadence, and expiry. Separate read-only research and drafting from external communication, financial decisions, and production modifications. Enforce deterministic policy gates outside the model for consequential actions. Demonstrate that suspending an agent revokes tokens and stops both active tasks and queued triggers.
 - Decision trigger: Require risk and IAM sign-off whenever an agent receives dedicated credentials, serves multiple employees, or can initiate a business action without a prompt.
 - Why now: Microsoft is introducing organizational agent accounts while additional administrative controls are still arriving. The right time to establish this contract is before departmental agents become unowned long-running processes.[1][2]
2. — Treat memory, events, and recoverability as production data systems For one real recurring process, map authoritative records and document which categories of memory may be retained. Require provenance, time limits, inspection, correction, and deletion. Capture unique task IDs, triggers, checkpoints, approvals, tool responses, transaction receipts, and compensating actions. Red-team malicious email, poisoned persistent memory, duplicate triggers, delayed approvals, a failed connector, and the departure of the original process owner.
 - Decision trigger: Block expansion if a duplicate event can repeat an external commitment, stale memory can override current policy, or auditors cannot reconstruct and reverse a material decision.
 - Why now: Persistent agents move the primary failure mode from an inaccurate answer to a durable side effect. Google's execution research and independent agent-safety reporting make resumability and memory integrity explicit production requirements.[4][11]
3. — Compete the control plane and price actual completed work Run a bounded 60–90-day pilot of a Microsoft 365-native business process against an alternative suited to the workload, such as Google, Oracle, or a governed independent runtime. Measure accurately completed cases, cycle time, human rework, exception rate, cloud and model spending, reviewer hours, and recovery time. Require exports of workflow definitions, memory provenance, action receipts, and identity mappings. Negotiate hard budgets, event-rate caps, approval escalation, and a credible exit route.
 - Decision trigger: Withhold broad procurement when identity or memory cannot be exported, material actions cannot be independently audited, or the supplier cannot show cost per correctly completed business outcome.
 - Why now: The enterprise-agent market is being assembled now and vendors are racing to own durable execution and its meter. Our analyses of enterprise OpenClaw successors and Microsoft Copilot pricing show why an early convenience decision can become an expensive architectural dependency.[3][9]

Sources

- Microsoft Copilot Studio announcement, October 15, 2026 [1]
- Microsoft Agent 365 agent-identity documentation [2]
- Frontier Research: The Agents Have Left the Chat [3]
- Google Cloud Gemini agent announcement, October 8, 2026 [6]
- Google Cloud Agent Executor technical disclosure [4]
- Oracle Fusion Claw announcement, September 29, 2026 [5]
- OpenAI introducing dots, September 29, 2026 [7]
- OpenAI Help Center: Enterprise dots permissions [12]
- Reuters: UK regulator examines Microsoft business software, October 16, 2026 [8]
- Reuters Breakingviews: Oxford summit and corporate agents, October 19, 2026 [11]
- Frontier News: Microsoft metered Copilot Business pricing [9]
- Frontier News: AWS fills in the agent stack [10]
- Frontier News: NVIDIA's agent security stack moves below the model [13]

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

- 01** **Microsoft Copilot Studio announcement, October 15, 2026**
<https://www.microsoft.com/en-us/copilot/blog/copilot-studio/delegate-shared-business-processes-to-agents-that-are-more-persistent-and-proactive-with-microsoft-copilot-studio/>
- 02** **Microsoft Agent 365 agent-identity documentation**
<https://learn.microsoft.com/en-us/microsoft-agent-365/guidance/agent-identity>
- 03** **Frontier Research: The Agents Have Left the Chat**
<https://thisisfrontier.ai/blog/the-agents-have-left-the-chat>
- 04** **Google Cloud Agent Executor technical disclosure**
<https://cloud.google.com/blog/products/ai-machine-learning/agent-executor-googles-distributed-agent-runtime/>
- 05** **Oracle Fusion Claw announcement, September 29, 2026**
<https://www.oracle.com/news/announcement/oracle-extends-fusion-agentic-applications-with-introduction-of-fusion-claw-2026-09-29/>
- 06** **Google Cloud Gemini agent announcement, October 8, 2026**
<https://cloud.google.com/blog/products/ai-machine-learning/welcome-to-gemini-at-work-2026>
- 07** **OpenAI introducing dots, September 29, 2026**
<https://openai.com/index/introducing-dots/>
- 08** **Reuters: UK regulator examines Microsoft business software, October 16, 2026**
<https://www.reuters.com/legal/litigation/uk-competition-regulator-targets-microsofts-business-software-grip-2026-10-16/>
- 09** **Frontier News: Microsoft metered Copilot Business pricing**
<https://thisisfrontier.ai/news/microsoft-will-default-new-csp-copilot-business-sales-to-metered-billing-making-ai-cost-controls-an-architectu>
- 10** **Frontier News: AWS fills in the agent stack**
<https://thisisfrontier.ai/news/aws-agent-stack-road-guardrails-tollbooth>
- 11** **Reuters Breakingviews: Oxford summit and corporate agents, October 19, 2026**
<https://www.reuters.com/commentary/breakingviews/oxford-summit-rings-alarm-corporate-ai-2026-10-19/>
- 12** **OpenAI Help Center: Enterprise dots permissions**
<https://help.openai.com/en/articles/20001554-manage-dots-in-chatgpt-workspaces>
- 13** **Frontier News: NVIDIA's agent security stack moves below the model**
<https://thisisfrontier.ai/news/nvidia-open-agent-safety-platform-runtime-security>

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.