

Microsoft 365 Copilot Studio vs. OpenAI ChatGPT Workspace Agents: An Enterprise Platform Comparison

A decision-grade analysis of integration depth, model choice, shared state, governance, and ROI — prepared July 2026.

Executive Summary

For enterprises choosing between Microsoft 365 Copilot Studio agents and OpenAI ChatGPT Workspace Agents, the decision has become less about raw model quality and more about *where the platform lives relative to your data and identity fabric*. Two structural facts dominate the analysis, and one of them overturns a common assumption.

First, the framing of "GPT-4 vs. Claude" as the axis that separates these platforms is now inverted. As of 2026, Copilot Studio is the **model-agnostic platform** — it offers both OpenAI GPT-5-family models *and* Anthropic Claude (Sonnet and Opus), plus Grok and Mistral, selectable per agent ([Microsoft Learn — Select a primary AI model](#)). ChatGPT Workspace Agents are **OpenAI-model-only**; no primary OpenAI documentation exposes Claude, Gemini, or any non-OpenAI model as a selectable option inside the ChatGPT product ([OpenAI Developers — External model support](#)). Teams that prioritize model-agnostic orchestration will, counterintuitively, find deeper multi-model flexibility inside Microsoft's platform today.

Second, the integration architectures differ at a foundational level. Copilot Studio grounds natively on the **Microsoft Graph semantic index**, inheriting tenant permissions automatically across SharePoint, Teams, Outlook, OneDrive, and Dataverse ([Microsoft Learn — Semantic indexing](#)). ChatGPT reaches the same data only through **external OAuth/Graph-API connectors** configured per-source, with practical gaps such as an all-or-nothing OneDrive sync setting and no synced-connector support for non-US data-residency customers ([OpenAI — SharePoint synced connectors setup](#)).

The net verdict: **Microsoft-entrenched teams get materially better ROI from Copilot Studio** because agent runtime, Graph grounding, and Claude/GPT access are largely bundled into the existing \$30/user/month Copilot license, with no incremental charge for internal, in-tenant usage by licensed users. Teams prioritizing **lightweight, self-serve, cross-cloud workflow orchestration** — especially those not standardized on Microsoft 365 — will find ChatGPT Workspace Agents faster to stand up, cheaper at the entry point (\$20–25/user/month), and less encumbered by Power Platform's administrative surface, at the cost of model lock-in and shallower native data grounding.

Platform Overview

The two products occupy adjacent but architecturally distinct categories. Understanding what each is clarifies why their trade-offs diverge.

Microsoft 365 Copilot Studio is a graphical, low-code environment for building agents and agent flows, built on the same foundation as Power Platform ([Microsoft Learn — Copilot Studio overview](#)). It supports three agent archetypes: **declarative agents** (configured on Copilot's own model/orchestrator, no extra hosting, adhering to Microsoft 365 security and compliance), **custom engine agents** (fully customized assistants that can use specialized models and pro-code frameworks like Semantic Kernel or LangChain, publishable to external apps and websites), and **autonomous agents** (event-triggered agents that act without a user-initiated conversation) ([Microsoft Learn — Extensibility decision guide](#)). It sits alongside the lighter-weight Microsoft 365 Copilot Agent Builder, and agents can be promoted from Agent Builder into Copilot Studio to unlock richer governance and connectors ([Microsoft Learn — Choose between M365 Copilot and Copilot Studio](#)).

OpenAI ChatGPT Workspace Agents are the enterprise-facing agent-builder surface inside ChatGPT Business and Enterprise. Teams build agents for repeatable tasks by choosing a model and reasoning effort, connecting apps/tools/MCP servers/files, sharing to a team directory, invoking from Slack, running on a schedule, or triggering via API ([OpenAI — ChatGPT Workspace Agents](#)). This surface sits within a broader OpenAI agent ecosystem that includes no-code **Custom GPTs**, the autonomous **ChatGPT Agent mode** (a merger of the former Operator and Deep Research), and developer-facing APIs. Notably, OpenAI is **winding down its low-code AgentKit "Agent Builder" and Evals products on November 30, 2026**, explicitly steering natural-language agent building toward Workspace Agents and code-based workflows toward the Agents SDK ([OpenAI — Introducing AgentKit](#)). This consolidation makes Workspace Agents the strategic center of OpenAI's enterprise agent story.

Dimension	Copilot Studio	ChatGPT Workspace Agents
Category	Low-code agent + workflow platform on Power Platform	Agent builder embedded in ChatGPT workspace
Agent types	Declarative, custom engine, autonomous	Workspace Agents, Custom GPTs, ChatGPT Agent mode
Primary builder persona	Makers + pro-code developers	Business users + light builders
External publishing	Yes (Teams, web, custom endpoints, external sites)	Team directory, Slack, API trigger
Default enablement	Included with Copilot license	Off by default; admin must enable (OpenAI — Workspace Agents)
Strategic trajectory	Multi-agent orchestration GA (A2A, Fabric, Agents SDK)	Consolidating around Workspace Agents; AgentKit wind-down Nov 2026

Integration Depth with Internal Data

This is the most consequential dimension for the Microsoft-vs-model-agnostic decision, because it determines how much engineering effort stands between an agent and an organization's real knowledge.

Copilot Studio: native Graph grounding with inherited permissions

Copilot Studio's defining advantage is that Microsoft 365 Copilot "maps your organization's data into an advanced lexical and semantic index" built directly from **Microsoft Graph**, giving agents contextually precise retrieval across tenant content without bespoke connector wiring ([Microsoft Learn — Semantic indexing](#)). Knowledge sources — SharePoint (via GraphSearch), Dataverse (with RAG), and Microsoft Search-indexed enterprise connectors — all require the querying user's Microsoft Entra ID authentication, so an agent "only surfaces content that the specific user can access" ([Microsoft Learn — Knowledge sources summary](#)). This permission inheritance is automatic: if a user lacks access to a SharePoint site, Teams channel, or Outlook mailbox, the agent will not reveal its content ([Microsoft Learn — Choose between M365 Copilot and Copilot Studio](#)).

For data beyond core M365, Microsoft offers layered reach: **more than 100 prebuilt Copilot (Graph) connectors** for indexing content from Box, Salesforce, ServiceNow, SAP, Jira, and others into the Graph ([Microsoft Learn — Copilot connectors overview](#)); the **1400+ Power Platform connectors** usable as callable tools/actions inside agents ([Microsoft Learn — Power Platform connector](#)); and full **Model Context Protocol (MCP)** support for tools and resources when generative orchestration is enabled ([Microsoft Learn — Extend your agent with MCP](#)). The critical distinction is that external connectors are needed only for *non-Microsoft* systems — core M365 content is reachable natively and permission-safe out of the box.

ChatGPT: external connectors, no native graph

ChatGPT's integration model is built entirely on **external, OAuth/Graph-API connectors** to third-party systems rather than a native unified knowledge index ([OpenAI — SharePoint synced connectors setup](#)). Its connector catalog is broad and cross-cloud — Box, Dropbox, GitHub, Gmail, Google Calendar/Contacts/Drive, HubSpot, Linear, Notion, and the Microsoft set (OneDrive, Outlook Email/Calendar, SharePoint, Teams) — with OpenAI marketing "60+ connectors" ([OpenAI — Connectors in ChatGPT](#)). Full write-capable MCP is in beta for Business/Enterprise/Edu, and a "Secure MCP Tunnel" enables private/on-prem server connections ([OpenAI — Developer mode and MCP apps](#)).

However, the depth of reach into Microsoft 365 specifically carries practical friction that matters for the entrenched-Microsoft buyer:

- **All-or-nothing OneDrive:** within the admin-managed SharePoint connector, OneDrive is supported *only* if the admin selects "sync all" (full site-collection permissions); scoped syncs to specific sites/folders exclude OneDrive entirely ([OpenAI — SharePoint synced connectors setup](#)).
- **Data-residency gap:** synced connectors are supported only for US-based customers or those who have not elected at-rest data residency; in-region storage for non-US residency

configurations is not yet supported ([OpenAI — SharePoint synced connectors setup](#)).

- **File limits:** 100 MB max file size, restricted to .txt/.pdf/.docx/.pptx/.xlsx ([OpenAI — SharePoint synced connectors setup](#)).
- **Frozen MCP snapshots:** approved custom MCP tool definitions do not auto-update when the live server changes; a manual refresh is required ([OpenAI — Developer mode and MCP apps](#)).
- **Agent-mode exclusion:** ChatGPT Agent mode will not use custom MCP apps, and Deep Research can use them only in read/fetch mode ([OpenAI — Developer mode and MCP apps](#)).

Integration attribute	Copilot Studio	ChatGPT Workspace Agents
Core M365 grounding	Native Graph semantic index	External per-source OAuth connectors
Permission model	Automatic Entra ID inheritance per user	Matched by email domain / user OAuth; SharePoint role-respecting
Non-Microsoft reach	100+ Graph connectors, 1400+ Power Platform connectors	60+ connectors (cross-cloud strength)
MCP support	Tools + resources (gen. orchestration on)	Beta, write-capable (Business/Ent/Edu); excluded from Agent mode
Notable M365 gaps	Declarative agents unavailable in Gov tenants	All-or-nothing OneDrive; non-US residency unsupported for sync; 100MB file cap

Bottom line on integration: for an organization whose knowledge already lives in SharePoint, Teams, Outlook, and OneDrive, Copilot Studio delivers deeper, lower-effort, permission-native grounding. ChatGPT's connector model is more *cross-cloud egalitarian* — equally comfortable with Google Drive, Slack, and GitHub — which favors heterogeneous shops over Microsoft-standardized ones.

Availability of LLM Choices

The task's original framing — "GPT-4 vs. Claude" — reflected a 2024–early-2025 reality that no longer holds. The 2026 picture is the reverse of what most buyers assume.

Copilot Studio: genuinely model-agnostic

Copilot Studio lets makers pick an agent's primary model from a dropdown, with separate settings for deep reasoning, generative responses, and the prompt builder ([Microsoft Learn — Select a primary AI model](#)). The available roster spans multiple vendors:

Model family	Examples	Provider
OpenAI GPT	GPT-4.1 (default), GPT-5 Chat, GPT-5.5 Chat, GPT-5 Reasoning/Auto	OpenAI (via Azure)
Anthropic Claude	Claude Sonnet 4.6, Sonnet 5, Opus 4.6/4.7	Anthropic (subprocessor)
xAI Grok	Grok 4.1 Fast	xAI
Mistral	Mistral Medium 3.5	Mistral

Beyond this built-in roster, admins can enable **external models from Anthropic, xAI, and Mistral** as an agent's primary model, and makers can **bring their own model from the Azure AI Foundry catalog** (DeepSeek, Phi-4, and others) for prompts and response generation ([Microsoft Learn — Select a primary AI model](#)). Anthropic Claude reached Copilot Studio via the Microsoft–Anthropic partnership: Claude is enabled by default for most commercial-cloud tenants (excluding EU/EFTA and UK, where it is off by default and admin-opt-in), and creators must explicitly select it during agent creation ([Microsoft Copilot Blog — Anthropic joins the multi-model lineup](#); [Microsoft Learn — Anthropic models in Microsoft Online Services](#)). Critically, Claude access through Copilot is **included in the existing Copilot license at no extra cost**, avoiding a separate Anthropic subscription ([Directions on Microsoft — M365 Copilot Adds Choice with Claude](#)).

ChatGPT: OpenAI-only

ChatGPT is confirmed **OpenAI-model-only**. No primary OpenAI documentation exposes Claude, Gemini, or other non-OpenAI models as selectable inside ChatGPT's chat picker, Custom GPTs, Workspace Agents, or Agent mode. The sole place OpenAI supports non-OpenAI models is the **Evals** developer tool — a benchmarking utility for developers, not a chat surface ([OpenAI — External model support](#)). Within OpenAI's own lineup, Business/Enterprise users choose among GPT-5.5 Instant, Thinking, and Pro via the model picker, and Workspace Agent builders select the model and reasoning effort per agent ([OpenAI — ChatGPT Business models & limits](#); [OpenAI — Workspace Agents](#)).

Implication: if "model-agnostic workflow orchestration" is a hard requirement, Copilot Studio is — as of 2026 — the more model-agnostic platform. ChatGPT offers best-in-class access to OpenAI's frontier models and tight vertical integration with them, but zero cross-vendor optionality inside the product. A team wanting to route reasoning-heavy steps to Claude Opus and cost-sensitive steps to a smaller model, all within one agent framework, can do so in Copilot Studio but not in ChatGPT.

Management of Shared State

Shared state — how agents remember, how context flows between agents, and how teams collaborate around a common knowledge surface — reveals a philosophical split: Copilot Studio treats state as a *governed data-platform concern*, while ChatGPT treats it as a *product-collaboration concern*.

Copilot Studio: variables, connected agents, and Dataverse

Within an agent, **variables** capture user responses and can be promoted to **global variables** available across all topics for a single user session ([Microsoft Learn — Work with global variables](#)). For multi-agent scenarios, a main agent can invoke **connected agents** in the same environment, and by default the conversation history is passed to the connected agent (with an option to restrict it to a single delegated task) ([Microsoft Learn — Connect to an existing Copilot Studio agent](#)). As of April 2026, Copilot Studio reached **general availability for multi-agent orchestration**, including Microsoft Fabric integration, Microsoft 365 Agents SDK orchestration, and open **Agent-to-Agent (A2A) protocol** support for delegating across first-, second-, and third-party agents ([Microsoft Copilot Blog — Multi-agent orchestration](#)). Durable, structured state is anchored in

Dataverse, which is a first-class, non-blockable connector and serves double duty as a knowledge source with RAG ([Microsoft Learn — Connector classification](#)).

ChatGPT: memory, shared projects, and the Conversations API

ChatGPT's shared-state story centers on three mechanisms. **Memory** (auto-updating and legacy saved-memory modes) can be toggled workspace-wide by Enterprise owners — but the "Reference Chat History" capability is **not yet available to Enterprise/Edu customers**, a notable gap versus Plus/Pro ([OpenAI — Memory FAQ](#)). **Shared Projects** (GA since October 2025) are the primary team-knowledge mechanism, grouping chats, files, and instructions into a shared context hub with Edit/Chat roles, though capped at 100 individuals and 40 files per project for Business/Enterprise, and once shared, a project switches irreversibly to project-only memory ([OpenAI — Projects in ChatGPT](#)). For programmatic persistence, OpenAI is retiring the **Assistants API on August 26, 2026**, migrating durable conversation state to the new **Responses API + Conversations API** stack ([OpenAI — Assistants API migration guide](#)).

State capability	Copilot Studio	ChatGPT Workspace Agents
In-session variables	Topic + global variables	Implicit within agent/thread
Cross-agent context	Connected agents pass history; A2A protocol (GA)	Agent orchestration within workspace
Team shared context	Dataverse, shared knowledge sources	Shared Projects (100-person cap)
Durable/programmatic state	Dataverse (governed, non-blockable)	Conversations API (Assistants retiring Aug 2026)
Notable gap	State tied to Power Platform environment model	"Reference Chat History" unavailable to Enterprise

Interpretation: Copilot Studio's Dataverse-backed, A2A-enabled state model is more robust for complex, auditable, multi-agent enterprise systems. ChatGPT's Shared Projects are simpler and faster for teams to adopt but hit collaboration ceilings and carry an Enterprise memory gap that undercuts long-horizon personalization.

Governance and Admin Controls

For enterprise deployment, governance is often the deciding factor — and here the two platforms reflect their heritage. Copilot Studio inherits Microsoft's decades-deep compliance apparatus; ChatGPT offers a cleaner, more modern console that is nonetheless narrower in surface.

Copilot Studio: Power Platform depth plus Purview

Governance is split across two consoles: the **Microsoft 365 admin center** (agent inventory, enable/disable/block, pay-as-you-go billing, Purview compliance) and the **Power Platform admin center** (finer-grained enterprise controls) ([Microsoft Learn — Choose between M365 Copilot and Copilot Studio](#)). The centerpiece is **Data Loss Prevention (DLP)**, enforced tenant-wide since early 2025, which classifies connectors into Business/Non-business/Blocked groups and can require Entra ID authentication, block specific knowledge sources, block connectors-as-tools

(which also blocks dependent MCP tools), and restrict publishing channels ([Microsoft Learn — Configure data policies for agents](#)). Compliance runs through **Microsoft Purview DSPM for AI** (sensitivity labels, DLP on labeled documents, unified audit logging), and every agent can receive a distinct **Microsoft Entra Agent ID** for identity governance ([Microsoft Purview Blog — Safeguard custom Copilot agents](#); [Microsoft Learn — Entra agent identities](#)). A structured ALM strategy (dev/test/prod environments, solutions, CI/CD via Pipelines/Azure DevOps/GitHub) governs the agent lifecycle, and the broader **Agent 365** control plane (GA May 2026) inventories agents across first- and third-party platforms ([Microsoft Learn — ALM strategy](#); [Microsoft Security Blog — Agent 365 GA](#)).

ChatGPT: a clean console with modern primitives

ChatGPT Enterprise offers four workspace roles (Member, Analytics viewer, Admin, Owner), **SCIM** and **SSO** across both ChatGPT and the API org (SAML/OIDC, ~25 IdPs including Entra ID), and a **Compliance API** with 13 pre-built partner integrations including Microsoft Purview, Netskope, Palo Alto, and Zscaler for eDiscovery/DLP/SIEM ([OpenAI — Workspace roles](#); [OpenAI — Compliance APIs](#)). **Data residency** spans ten regions (with true inference residency in Europe, US, and UAE), and OpenAI commits to **no training on business data by default** for Business/Enterprise/Edu ([OpenAI — Data residency](#); [OpenAI — ChatGPT Pricing](#)). Agent-specific governance includes write-action confirmation defaults ("Always ask"), Connector Action Constraints, publishing-permission tiers, and a specific risk warning when publishing agents that rely on agent-owned (shared) connections ([OpenAI — Workspace Agents](#)).

Two governance caveats favor Microsoft for regulated buyers: Business-tier SSO is ChatGPT-only and **excludes SCIM** (requiring an Enterprise upgrade for directory sync), and a **HIPAA BAA reportedly requires the Enterprise tier**, not Business ([OpenAI — SSO for ChatGPT Business FAQ](#); [Beam Cloud — ChatGPT Enterprise Pricing Guide 2026](#)).

Governance attribute	Copilot Studio	ChatGPT Workspace Agents
Admin surface	M365 admin center + Power Platform admin center	Single ChatGPT admin console
Data policy engine	DLP connector classification (tenant-wide)	Action control, parameter constraints, domain allow-lists
Identity for agents	Entra Agent ID per agent	Agent-owned vs. end-user account connections
Compliance tooling	Purview DSPM for AI, unified audit log	Compliance API + 13 partner integrations (incl. Purview)
Directory sync (SCIM)	Native Entra	Enterprise only (not Business)
Data residency	Azure regions / EU Data Boundary (Claude excluded)	10 regions; inference residency in EU/US/UAE
Lifecycle mgmt	Full ALM (dev/test/prod, CI/CD)	Drafts/publish flow, versioning

Interpretation: Copilot Studio's governance is deeper and more granular but demands Power Platform administrative expertise — an operational tax that is trivial for Microsoft-mature IT

organizations and burdensome for others. ChatGPT's console is materially simpler to operate and sufficient for most mid-market deployments, but its full governance (SCIM, HIPAA BAA, custom retention) is gated behind the unpublished-price Enterprise tier.

Pricing and Total Cost of Ownership

Pricing is where the "already-entrenched" calculus becomes concrete. The comparison is asymmetric: Microsoft bundles agent capability into a per-user Copilot license with metered overage, while OpenAI sells workspace seats with token-metered agent runs.

Copilot Studio economics

Microsoft 365 Copilot (Enterprise) is **\$30/user/month** (annual), as an add-on to a qualifying M365 base plan; a Business tier runs **\$18–21/user/month** for organizations up to 300 users ([Microsoft — M365 Copilot Pricing](#)). Copilot Studio access is **included** with the \$30 Copilot license, and crucially, internal in-tenant agent usage by licensed users incurs **no additional charge** — classic answers, generative answers, agent actions, and Graph grounding are all "no charge" for Copilot-licensed users ([Microsoft — Copilot Studio pricing](#); [Microsoft Learn — Billing rates and management](#)). Usage beyond that (external/customer-facing agents, unlicensed users) is billed via **Copilot Credits** — either a pre-purchased pack (**25,000 credits for \$200/month**, tenant-wide) or pay-as-you-go at **\$0.01/credit** ([Microsoft — Copilot Studio pricing](#)). Claude and GPT model access is included at no incremental model fee ([Directions on Microsoft — Claude in M365](#)). An Azure subscription is required to deploy agents ([Microsoft — Copilot Studio pricing](#)).

ChatGPT economics

ChatGPT Business is **\$25/user/month monthly or \$20/user/month annual** (2-seat minimum), including ChatGPT + Codex, connectors, SSO/MFA, and no-training-by-default ([OpenAI — What is ChatGPT Business?](#)). A \$5/seat price cut took effect April 2, 2026 ([OpenAI — Business Release Notes](#)). **Enterprise pricing is not published** — "Contact sales" — with third-party estimates clustering around **\$60/user/month** (~\$45–75 range) and a reported ~150-seat minimum, implying a rough \$108,000/year floor (these figures are unverified market reports, not OpenAI list pricing) ([OpenAI — ChatGPT Pricing](#); [Beam Cloud — Enterprise Pricing Guide 2026](#)). Workspace Agent runs are **token-metered**, drawing from the Codex agentic usage credit pool on Business plans ([OpenAI — ChatGPT Workspace Agents](#)). ChatGPT subscription fees do **not** include API token usage — AgentKit/API consumption is billed separately ([OpenAI — API Pricing](#)).

Cost element	Copilot Studio	ChatGPT Workspace Agents
Entry per-user	\$18–21/mo (Business) / \$30/mo (Enterprise)	\$20/mo annual, \$25/mo monthly (Business)
Enterprise tier	\$30/user/mo (published)	Custom / unpublished (~\$60 est., 3rd-party)
Agent runtime for licensed internal use	Included, no charge	Token-metered from credit pool
Overage model	Copilot Credits (\$200/25k pack or \$0.01/credit)	Purchasable credits / flexible pricing

Cost element	Copilot Studio	ChatGPT Workspace Agents
Multi-model access	Claude + GPT included in license	OpenAI-only, included
Prerequisite	Qualifying M365 base license + Azure sub	2+ seats

TCO insight: For a Microsoft-licensed organization, the *marginal* cost of Copilot Studio agents is low — the runtime, Graph grounding, and both Claude and GPT are bundled into a license many such organizations are already evaluating or paying for. For a team with no Microsoft 365 footprint, ChatGPT's \$20–25 entry point is cheaper and faster to activate, with no base-license or Azure prerequisite, though heavy agent usage introduces token-metering variability that requires budget monitoring.

ROI Verdict by Scenario

The two buyer profiles in the brief map cleanly onto the evidence.

Teams entrenched in the Microsoft ecosystem → Copilot Studio

For organizations already standardized on Microsoft 365, Copilot Studio delivers superior ROI on four reinforcing grounds:

1. **Bundled runtime and grounding.** Agent execution and native Microsoft Graph grounding are included in the \$30 Copilot license with no per-message charge for internal use, eliminating the token-metering variability of ChatGPT's agent runs ([Microsoft Learn — Billing rates](#)).
2. **Zero integration engineering for core data.** SharePoint, Teams, Outlook, OneDrive, and Dataverse are reachable natively with automatic permission inheritance — no per-connector OAuth setup, no all-or-nothing OneDrive constraint, no non-US residency gap ([Microsoft Learn — Knowledge sources](#)).
3. **Model choice is a bonus, not a trade-off.** These teams get Claude *and* GPT inside the same platform at no extra model cost — the model-agnostic capability without leaving the Microsoft fabric ([Microsoft Copilot Blog — Anthropic multi-model](#)).
4. **Governance reuses existing muscle.** Entra ID, Purview, and DLP are tools these organizations already operate; the Power Platform admin surface is an incremental extension, not a new discipline ([Microsoft Learn — DLP for agents](#)).

The primary caveats: Copilot Studio carries real Power Platform administrative complexity, declarative agents are unavailable in government tenants, and Claude usage sits outside the EU Data Boundary (disabled by default in EU/EFTA/UK) ([Microsoft Learn — Anthropic subprocessor](#)).

Teams prioritizing custom, model-agnostic workflow orchestration → nuanced

If "model-agnostic" is read literally as *multi-vendor model choice*, Copilot Studio is paradoxically the stronger fit today, since ChatGPT is OpenAI-only. But if the priority is *custom, cross-cloud, low-friction orchestration* — fast self-serve agent building, Slack-native invocation, scheduling, API triggers, and connectors that treat Google Drive, GitHub, and Slack as first-class alongside Microsoft — ChatGPT Workspace Agents win on speed-to-value and operational simplicity

([OpenAI — Workspace Agents](#)). ChatGPT is the better choice for heterogeneous, non-Microsoft-centric organizations, teams that value a single clean admin console over granular DLP, and those who want OpenAI's frontier models with minimal setup overhead.

The trade-offs to price in: model lock-in to OpenAI, shallower native M365 grounding, the Enterprise-tier gating of SCIM and HIPAA BAA behind unpublished pricing, the Enterprise "Reference Chat History" memory gap, and the November 2026 AgentKit wind-down that concentrates OpenAI's roadmap on this exact surface ([OpenAI — Introducing AgentKit](#)).

Recommendation Framework

If your organization...	Choose	Because
Runs on Microsoft 365 (SharePoint/Teams/Outlook)	Copilot Studio	Native Graph grounding + bundled runtime + Claude & GPT included
Needs multi-vendor model routing in one agent	Copilot Studio	Only platform offering Claude, GPT, Grok, Mistral + BYOM
Is cross-cloud / not Microsoft-standardized	ChatGPT	Connector parity across Google, Slack, GitHub; faster setup
Wants minimal admin overhead, fast self-serve	ChatGPT	Single console; no Power Platform learning curve
Requires deep, granular DLP + Purview compliance	Copilot Studio	Tenant-wide DLP, sensitivity labels, Entra Agent ID
Prioritizes lowest entry price with no base license	ChatGPT	\$20/user/mo annual, 2-seat minimum, no M365/Azure prereq
Operates in government/sovereign clouds	Copilot Studio (with caveats)	Broader gov support, though Claude & declarative agents restricted

The strategic through-line: the platform decision now hinges primarily on data gravity and identity fabric, not on which frontier model you can access. Both platforms offer excellent models; Copilot Studio uniquely offers *choice among* them. The organization's existing Microsoft footprint — or lack of one — is the dominant ROI variable.

Prepared from primary Microsoft Learn and OpenAI Help Center documentation, official pricing pages, and vendor blogs as of July 2026. Enterprise ChatGPT list pricing is not published by OpenAI; figures cited are third-party market estimates and flagged as such. Model availability, pricing, and feature rollouts in this category change frequently — verify against live vendor pages before finalizing procurement decisions.