SEPTEMBER 2026 · PLATFORM COMPARISON

ChatGPT, Copilot, Claude or Odokai: which one will do the work?

Most UK SMEs choosing an AI platform end up with the same shortlist. Copilot is already on the Microsoft bill, ChatGPT is the name everyone knows, and someone is usually trying Claude. The question is which of them will take work off your team, and which is just another subscription.

If you want an assistant for every employee, one of the incumbents is likely the right buy, and the tables show which. If you want agents and apps built around how your business runs, Odokai gives you your choice of hundreds of models, your data where you put it, knowledge that belongs to your business, and a fixed monthly bill.

3	15	13	4
comparison tables	capabilities assessed	enterprise controls	platforms compared

THE FOUR PLATFORMS

Four different assumptions about where work happens.

Each platform starts from a different idea of where your team's work happens, and that shapes everything else about it.

ChatGPT Enterprise

A sales-negotiated workspace putting the ChatGPT assistant in front of the whole organisation. SSO and SCIM, role-based access, a compliance API covering eDiscovery,

Microsoft 365 Copilot

Not a standalone product: a $30 per user per month add-on to a qualifying Microsoft 365 plan, or $21 as Copilot Business for up to 300 users. In exchange, AI sits inside Word, Excel,

DLP and SIEM export, data residency across ten storage regions, customer-managed encryption keys, custom retention and SLAs. The default choice for a large organisation that wants one standalone assistant, independent of any office suite, with mature controls.

PowerPoint, Outlook and Teams, grounded in the tenant's own data through Graph, inheriting the entire Microsoft compliance stack. The path of least resistance inside a Microsoft estate, and weak outside one.

Claude Enterprise

One seat covers the full suite: chat, Claude Code, delegated multi-step work, design, connectors to Google Workspace, Microsoft 365, GitHub and Slack, plus enterprise search, memory, skills and projects. The $20 seat fee covers access only; all usage is billed separately at API rates, with optional spend limits for the organisation and each user. Strong for heavy, variable users.

Odokai

A multi-agent operating system: private, persistent agents set up around your team and the processes it runs. Deployed as a shared managed service in AWS London, or single-tenant in an AWS account we manage or in your own AWS account. Built for document-heavy and regulated work, with a consultancy service that gets your first workflow live.

METHOD

How this comparison was compiled.

Compiled in September 2026 from vendor documentation and third-party procurement reporting. No vendor commissioned it. It was produced by an assistant running on the Odokai platform, so Odokai's entries are vendor-sourced.

Bias. Produced on the Odokai platform. Odokai's entries are marked vendor-sourced throughout.

Evidence. Odokai's entries were checked against its own product documentation, not verified by an outside party. Incumbent entries come from published vendor documentation, with certifications listed only where publicly claimed. None of the figures has been independently audited.

Gaps. Odokai has no SOC 2 or ISO 27001 certification today; an ISO 27001 plan is in progress. SSO, SAML and SCIM provisioning are implemented as part of an engagement.

COMMERCIAL MODEL

What each platform costs.

Seat prices only tell part of the story. Check what each price leaves out, and who pays when usage runs higher than planned.

	ChatGPT Enterprise	M365 Copilot	Claude Enterprise	Odokai
Published price	Quote only. Reported at $45–75 per user per month, around $60 typical.	$30 per user per month, annual, as an add-on to a qualifying base licence. Copilot Business is $21 for up to 300 users.	$20 per seat per month, annual, plus all usage at API rates.	**Fixed monthly fee with included credits.** No figure published; the pricing model is set out separately.
Realistic all-in	~$60 per user per month, reported.	$69–$90 per user per month once the E3 or E5 licence underneath is counted.	~$37 light use, ~$87 heavy use. Variable, with optional spend limits.	The fixed fee. Infrastructure cost applies only in single-tenant mode.
Seat minimum	Around 150, reported. An implied floor near $108k a year.	None, but the base licence is compulsory.	20 self-serve, 50 sales-assisted.	A seat limit per tenancy rather than a floor.
Usage billing	Credit pools, or token-based billing on some contracts.	Mostly included. Cowork and some agent usage are metered in Copilot Credits.	Everything billed separately at API rates. No included allowance.	Monthly credit cap with a warning threshold and a soft stop when requests are admitted.
Who carries the overrun	Buyer, moderate. Higher on token billing.	Buyer, low.	Buyer, high, unless spend limits are set.	Buyer, low. The cap is buyer-set and enforced.
Procurement route	Sales-negotiated, annual, twelve-month term.	Self-serve or Enterprise Agreement, annual.	Self-serve or sales-assisted, multi-currency on assisted.	Sales-assisted: a call to scope the first workflow, then a fixed fee in writing.

All figures in USD as published or reported in September 2026, excluding tax. Microsoft's SMB bundle promotions ended on 30 June 2026 and base prices rose on 1 July 2026 (E3 $36 to $39, E5 $57 to $60), with the add-on itself held at $30. Partners report

enterprise Copilot promotions running to 30 September 2026. Vendor pricing changes often; re-verify at contract time.

CAPABILITIES

What each platform does.

Fifteen capabilities, assessed the same way for all four platforms.

	ChatGPT Enterprise	M365 Copilot	Claude Enterprise	Odokai
General assistant chat	Yes	Yes	Yes	Yes
Embedded in the office suite	Yes, Word, Excel and PowerPoint add-ins	**Yes, its core strength**	Yes, Excel, Word and PowerPoint add-ins; Outlook in beta	Via connectors, Microsoft 365 included
Grounding in company data	Yes	Yes, through Graph	Yes	Yes, with citations and per-file grants
Custom agents or assistants	Yes, GPTs and workspace agents	Yes, Copilot Studio	Yes, skills	Yes, the central concept
Model choice and bring-your-own model	Own models only	Yes, OpenAI and Anthropic models; bring your own through Foundry in Copilot Studio	Own models only	**Any provider with an OpenAI-compatible API**, including Anthropic, OpenAI, Gemini, OpenRouter or a local model
Agentic multi-step workflows	Yes	Yes, Cowork, Researcher and Agent Mode	Yes	Yes, as durable tasks
Developer tooling	Yes, Codex	Partial, Agents Toolkit and SDK; GitHub Copilot sold separately	Yes, Claude Code	Yes: shell, Git, SDK, CLI, browser control
Persistent memory across sessions	Yes	Yes	Yes	Yes, at platform level rather than per user
Background and scheduled work	Yes, tasks	Yes, scheduled prompts and Copilot Studio triggers	Yes, Cowork scheduled tasks and Claude Code routines	Yes, durable across deploys, and scheduled
Multi-agent orchestration	Yes, subagents in ChatGPT Work and Codex	Yes, child and connected agents in Copilot Studio	Yes, subagents, workflows and agent teams	**Panel, pipeline, triage, debate, map-reduce, delegation trees**
Human approval gates on tool calls	Yes. Admins set per-app action limits and ask-before-write permissions	Yes. Per-tool approval in Copilot Studio agents	Yes. Admins set "Needs approval" per connector and tool, by role	**Per-tool policy, dual control, SLAs**

	ChatGPT Enterprise	M365 Copilot	Claude Enterprise	Odokai
PII redaction before the model sees it	No	Partial, DLP blocks sensitive content; no redaction	No redaction; DLP blocking in beta	**On by default**
Internal app building	Partial, canvas	Partial, Frontier	Yes, artifacts with storage and connectors	Yes, App Studio
Broad third-party integrations	Yes	Yes, 100+ connectors, Power Platform and MCP	Yes	146 reviewed connectors plus 8 first-party anchors
Versioned shared workspace	Partial, shared projects	Yes, SharePoint	Partial, shared projects, not versioned	Yes, versioned

Odokai entries come from Odokai's own documentation and have not been independently audited. "Partial" covers preview features, connector-based support, and anything not offered as a built-in capability.

SECURITY & COMPLIANCE

Which enterprise controls each platform meets.

Thirteen controls. The incumbents lead on certification. All four offer approval gates on tool calls. Odokai is the only one that redacts PII before the model sees it.

	ChatGPT Enterprise	M365 Copilot	Claude Enterprise	Odokai
SSO and SAML	Yes	Yes, Entra	Yes	**Yes.** Provided as part of your engagement, on the shared platform or your own instance.
SCIM provisioning	Yes	Yes	Yes	**Yes.** Implemented as part of your engagement.
SOC 2 or ISO 27001	SOC 2, ISO 27001 and ISO 42001	Yes, via Microsoft 365	SOC 2 Type II, ISO 27001 and ISO 42001	**No.** ISO 27001 in progress. In your own AWS account, the infrastructure sits under your certifications and AWS's.
Compliance API for eDiscovery, DLP, SIEM	Yes	Yes	Yes	Partial. Auditor exports as CSV, signed JSON and PDF, but no API.

	ChatGPT Enterprise	M365 Copilot	Claude Enterprise	Odokai
Audit logs	Yes	Yes, Purview	Yes	Yes. Classified, encrypted, roughly seven-year retention, admin viewer.
Customer-managed encryption keys	Yes	Yes	Yes	Yes, in single-tenant. Customer-owned KMS keys, rotation and CloudTrail.
Data residency options	Yes, ten regions	Yes, Microsoft 365 geo and EU Data Boundary, excluding Anthropic models	US-only inference option; no EU or UK	London for the managed service, the customer's own region single-tenant, or local models.
No training on customer data	Default	Yes	Default	Enforced through an approved-provider allowlist.
Multi-factor authentication	Yes	Yes	Yes	Yes. TOTP, recovery codes, organisation-wide enforcement.
Role-based access control	Yes	Yes	Yes	Yes. Typed permission catalogue, default deny.
Approval gates on tool calls	Yes, per connector	Yes, per tool in Copilot Studio	Yes, per connector and tool	**Yes.** Dual control, SLAs, fail-closed.
PII redaction before the model	No	Partial, DLP blocks sensitive content; no redaction	No redaction; DLP blocking in beta	**Yes.** On by default.
Data classification policy	No	Yes, Purview sensitivity labels and DLP	Partial, via Purview and DLP hooks	**Yes.** Per-class rules, including private-model-only.

Odokai's entries are vendor-sourced. Incumbent entries come from public vendor documentation, with certifications listed only where publicly claimed.

POINTS OF DIFFERENCE

Where Odokai is different.

Six ways Odokai works differently from the other three platforms.

It runs where the data already sits

A shared managed service in AWS London, single-tenant in an AWS account we manage, or single-tenant inside your own AWS account with its own VPC, database and KMS keys. The other three are SaaS services. If your data

Any model, including ones that never leave your environment

Any provider with an OpenAI-compatible API, including Anthropic, OpenAI and Microsoft models, plus hundreds more through OpenRouter and Gemini, and thousands of

can't leave your environment, Odokai can run inside it.

open models through Ollama, routed per workload and per data class. ChatGPT Enterprise and Claude Enterprise use their own models; Copilot offers OpenAI and Anthropic models. Only Odokai runs local models, so sensitive work never has to leave your environment.

The knowledge belongs to the organisation

Agents hold domain knowledge, memory and skills, and the platform keeps what they learn. All four have persistent memory, but in the other three it belongs to a user or a project. In Odokai it builds up as shared organisational knowledge, so it stays when people leave.

Governance built into the product

Approval gates with dual control, PII redaction on by default, per-class data rules including private-model-only, and permissions that deny by default. The incumbents offer approval gates, and Copilot classifies data through Purview. None of them redacts PII before the model sees it.

A fixed fee with a hard cap

A fixed monthly fee with included credits, a warning threshold, and a cap you set. Claude Enterprise bills all usage at API rates with no included allowance; spend limits can stop it, but not reduce the rate.

Structured multi-agent work, built in

Panels, pipelines, triage, debate, map-reduce and delegation trees, plus background work that survives a deployment and scheduled work that runs without anyone present. The incumbents run subagents and scheduled tasks; Odokai gives you these patterns ready to use.

INDICATIVE FIT

Which platform suits which situation.

A Microsoft estate and a rollout to everyone

Engineering-heavy teams wanting coding agents

Microsoft 365 Copilot, comfortably. The work already happens in Word, Excel, Outlook and Teams, the identity and compliance stack is already in place, and deployment friction is close to zero. ChatGPT Enterprise alongside it where a meaningful part of the organisation works outside Microsoft. If what you want is an assistant for everyone, this is the assistant to buy.

Claude Enterprise. The strongest developer tooling of the four, one seat covering chat, coding and delegated work. The trade is usage-based billing; ChatGPT Enterprise's standard seats, which include Codex, are the alternative where fixed seat pricing matters more.

Vendor risk an oversight committee will not carry

One of the three incumbents, if your board needs the vendor scale and certification those platforms hold today. Odokai is a small UK company, live for over a year, with no SOC 2 or ISO 27001 certification yet; an ISO 27001 plan is in progress. Source code escrow and other contractual protections are available for an Odokai engagement.

Cost predictability above everything else

Copilot's fixed add-on, a fixed per-seat Claude Team plan, or Odokai's fixed fee with a buyer-set cap. Each avoids Claude Enterprise's usage billing, where every token is billed at API rates with no included allowance.

Regulated, document-heavy operational work

Odokai. Agents hold the context of your business, apps are built around your processes, and work runs on a schedule whether or not anyone is watching. You choose from hundreds of models, your data stays where you put it, and there is an auditable record of everything the agents know and produce. The incumbents give every employee an assistant. Odokai gives your business a workforce.

Where did these figures come from?

Assistants and operating systems are different buys.

If you want the best assistant, buy Copilot where the work runs on Microsoft, Claude where engineers want coding agents, and ChatGPT Enterprise where familiarity matters most.

If you want AI that becomes part of how your team works, Odokai is an agentic operating system with a consultancy service to get you started. You build your own agents, apps and workflows on it, using the models you choose, in AWS London or your own AWS account, with a fixed monthly bill.

Source: https://odok.ai/compare. Printed from the web page on 19 October 2026.