

Enterprise AI Assistant Platforms

A neutral, evidence-based comparison of ChatGPT Enterprise, Microsoft 365 Copilot, Claude Enterprise and Odokai — covering commercial models, capabilities, security posture and fit.

Prepared 13 September 2026 **Version** 1.3 (corrected) **Status** Vendor-neutral; no sponsorship **Figures** USD, as reported Sept 2026

\$30 COPILOT ADD-ON / SEAT Plus base licence — all-in \$69–\$90	~\$60 CHATGPT ENT. / SEAT (REPORTED) Quote-only; ~150-seat floor	\$20 CLAUDE ENT. SEAT + USAGE All tokens billed at API rates	Cap ODOKAI PRICING MODEL Fixed monthly fee with credit caps
--	---	---	--

1 Executive summary

The four platforms occupy different positions rather than competing head-to-head on identical terms. ChatGPT Enterprise is a general-purpose assistant workspace with the most mature compliance tooling and a sales-led pricing model. Microsoft 365 Copilot is an add-on to an existing Microsoft estate, whose value depends almost entirely on how much of the organisation's work already happens in Word, Excel, Outlook and Teams. Claude Enterprise couples a low seat price with usage-based billing at API rates, which shifts cost risk from vendor to buyer. Odokai is a smaller, newer platform built around private, persistent AI agents, deployable as a managed service or in the customer's own environment, and aimed at document-heavy and regulated work rather than the general office market.

There is no outright winner. The right choice depends on three questions: where your data and workflows already live, how predictable costs must be, and whether you need a general assistant for every employee or a configured system for specific operational processes.

OBJECTIVITY AND CORRECTION STATEMENT

Compiled from vendor documentation and third-party procurement reporting current as of September 2026. Version 1.3 re-checks every incumbent entry against the vendors' own documentation (27 September 2026) and incorporates corrections supplied by Odokai on 13 September 2026, checked against its own repository. Odokai-specific facts are labelled vendor-sourced rather than independently audited; its confirmed gap (no SOC 2 or ISO 27001 certification yet) sits alongside everything else. The compiling assistant runs within the Odokai platform.

2 Platform profiles

ChatGPT Enterprise

OPENAI

Sales-negotiated workspace offering the ChatGPT assistant organisation-wide. Enterprise adds SSO/SCIM, RBAC, a compliance API (eDiscovery, DLP, SIEM export), data residency across ten storage regions, customer-managed encryption keys, custom retention and SLAs. Flexible pricing adds workspace credit pools and a Codex-only seat type.

Positioning: the default "AI for everyone" choice for large organisations wanting one standalone assistant with mature controls.

Microsoft 365 Copilot

ENTERPRISE

Not standalone: a \$30/user/month add-on on a qualifying Microsoft 365 plan, or \$21 as Copilot Business. Deep integration — Copilot in Word, Excel, PowerPoint, Outlook and Teams, grounded in the tenant's own Graph data, inheriting the Microsoft compliance stack. Includes pre-built agents (Researcher, Analyst, Facilitator), Copilot Studio agent building, and SharePoint Advanced Management.

Positioning: path of least resistance for Microsoft 365 tenants; weak for anyone not.

Claude Enterprise

ANTHROPIC

One seat covers the full suite: chat, Claude Code, Cowork (delegated multi-step work), Design, connectors to Google Workspace, Microsoft 365, GitHub and Slack, plus enterprise search, memory, skills and projects. Enterprise adds SCIM, audit logs, a compliance API, custom retention, HIPAA-readiness and customer-managed keys. Seat fee (\$20/seat/mo, annual) covers access only — all usage billed at API rates, with optional spend limits.

Positioning: strong for heavy, variable users (engineering, research) where procurement can tolerate usage-based spend.

Odokai

INDEPENDENT VENDOR

A multi-agent operating system: private, persistent AI agents configured around the team. Deployed three ways (vendor-sourced): a managed SaaS in AWS eu-west-2 (London), or single-tenant in Odokai's or the customer's own AWS account (own VPC, RDS PostgreSQL with pgvector, customer-owned KMS keys). Knowledge belongs in the platform, not the code — agents hold domain knowledge, memory and skills, and the platform retains what they learn. Includes a versioned shared workspace, 146 reviewed connectors plus 8 first-party anchors (including Microsoft 365, Outlook, OneDrive and Teams), model-agnostic routing (any OpenAI-compatible API, e.g. Anthropic, OpenAI, Gemini, OpenRouter, local Ollama), durable background and scheduled work, multi-agent orchestration, human-in-the-loop approval gates, and App Studio for internal web apps.

Positioning: document-heavy and regulated work where process knowledge, auditability and data locality matter more than a consumer chat experience. Small vendor, no certifications yet (ISO 27001 plan in progress); SSO, SAML and SCIM implemented per engagement.

3 Commercial model and cost

Figure 1 — Published / reported seat cost per user, per month (USD)

ChatGPT Enterprise is quote-only; reported figures shown in outline. Copilot shown as add-on only. Odokai's fee is a fixed monthly amount with included credits (vendor-sourced).





Figure 2 — Realistic all-in cost per user, per month (USD)


Copilot all-in includes required base licence (E3 or E5, post-July 2026 prices). Claude shows seat plus a mid-range illustrative usage band. Odokai is a fixed fee with included credits and an organisational credit cap (vendor-sourced).



* Illustrative only — Anthropic publishes no included allowance; actual usage cost depends entirely on team behaviour. Dashed bars = figure not independently confirmed.

Figure 3 — Where the pricing risk sits

Who carries the risk that actual consumption exceeds the plan: vendor (fixed price) or buyer (variable). Odokai's credit-cap model (vendor-sourced) places it near Copilot's risk profile, not Claude's.



3.1 Pricing timeline — what changes and when

- 30 JUNE 2026
Microsoft's promotional pricing on Business "with Copilot" bundles ends.
- 1 JULY 2026
Microsoft 365 base prices rise: E3 \$36→\$39, E5 \$57→\$60, Business Standard \$12.50→\$14. Copilot add-on stays \$30. Bundled "with Copilot" plans become permanent (\$23.50 / \$32).
- 30 SEPTEMBER 2026
Microsoft three-year commitment offer closes: 15% off the Copilot add-on for 300+ seats.
- 31 DECEMBER 2026
Copilot Business promotional pricing ends (\$18, standard \$21) for organisations up to 300 users.


3.2 Commercial terms compared

Dimension	ChatGPT Enterprise	M365 Copilot (Enterprise)	Claude Enterprise	Odokai
Published list price	Quote-only; reported \$45–75/user/mo (~\$60 typical)	\$30/user/mo add-on (annual)	\$20/seat/mo + usage at API rates	Fixed monthly fee with included credits (vendor-sourced; amount not independently confirmed)
Seat minimum	~150 reported	None (requires qualifying base licence)	20 (self-serve) / 50 (sales-assisted)	Seat limit per tenancy (vendor-sourced)
Usage billing	Credit pools, or token-based billing on some contracts	Mostly included; Cowork and some agent usage metered in Copilot Credits	All usage billed separately at API rates; no included allowance	Monthly credit cap with warn threshold and soft stop at request admission; optional pooled credits (vendor-sourced)
Hidden cost factors	~\$108k/yr implied floor; 12-month term	Base licence required: all-in \$69–\$90/user/mo	Unpredictable for heavy teams; spend caps configurable	Infrastructure cost only in single-tenant mode; managed SaaS carries none
Procurement route	Sales-negotiated, annual	Self-serve / EA, annual	Self-serve or sales-assisted; multi-currency on assisted	Self-serve card signup (managed SaaS); direct for single-tenant

4 Capabilities

Figure 4 — Capability coverage by platform

Number of the fifteen tracked capabilities each platform fully covers (see table 4.1). Coverage breadth is not quality — read the detail before deciding.



4.1 Capability matrix

Capability	ChatGPT Ent.	Copilot	Claude Ent.	Odokai
General assistant chat (web/desktop/mobile)	•	•	•	•
Embedded in office suite (Word/Excel/Outlook/Teams)	• add-ins	• core strength	• add-ins	◦ connectors incl. M365
Grounding in company data	•	• Graph	•	• KB + citations + POSIX-style grants
Custom agents / assistants	• GPTs	• Studio	• skills	• central concept
Model choice / bring-your-own model	◦ own models only	• OpenAI + Anthropic	◦ own models only	• multi-vendor + local (vendor-sourced)
Agentic multi-step workflows	• Work	• Cowork	• Cowork	• durable tasks
Developer tooling	• Codex	◦ Agents SDK	• Code	• shell, Git, SDK, CLI, browser control
Persistent memory across sessions	•	•	•	• platform-level
Background / scheduled work	• tasks	• scheduled	• scheduled	• durable (survives deploys) + scheduled
Multi-agent orchestration	• Codex	• Studio	• subagents	• panel/pipeline/triage/debate/map-reduce + delegation trees
Human-in-the-loop approval gates	• per app	• per tool	• per tool	• per-tool policy, dual control, SLAs (vendor-sourced)
PII redaction / data classification before the model	◦	◦ Purview	◦ DLP (beta)	• on by default (vendor-sourced)

Internal app building	○ canvas	○ Frontier	● artifacts	● App Studio
Broad third-party integrations	●	● 100+ connectors	●	● 146 reviewed + 8 anchors incl. M365 (vendor-sourced)
Versioned shared workspace / files	○	● SharePoint	● projects	● versioned

● = full support · ○ = partial, preview, via connectors, or not offered. Odokai entries marked "vendor-sourced" come from the vendor's repository documentation, not independent audit.

5 Security, compliance and data

Figure 5 — Enterprise control coverage

Share of the thirteen tracked controls (table 5.1) each platform meets today. Odokai's figure is vendor-sourced for the technical controls; its gap (certification) is a vendor-confirmed absence, not an unknown.



5.1 Control matrix

Control	ChatGPT Ent.	Copilot	Claude Ent.	Odokai
SSO / SAML	Yes	Yes (Entra)	Yes	Yes — provided per engagement
SCIM provisioning	Yes	Yes	Yes	Yes — implemented per engagement
Audit logs	Yes	Yes (Purview)	Yes	Yes — classified, encrypted S3, ~7-yr retention, admin viewer (vendor-sourced)
Compliance API (eDiscovery/DLP/SIEM)	Yes	Yes	Yes	Partial — auditor exports (CSV, signed JSON, PDF), no API (vendor-sourced)
Customer-managed encryption keys	Yes	Yes	Yes	Yes — customer-owned KMS, rotation, CloudTrail (single-tenant; vendor-sourced)
Data residency options	Yes (10 regions)	Yes (M365 geo; excl. Anthropic)	US-only inference	London SaaS; customer's region in single-tenant; local model option (vendor-sourced)
No training on customer data	Default	Yes	Default	Enforced via approved-provider allowlist (vendor-sourced)

MFA	Yes	Yes	Yes	Yes — TOTP, recovery codes, org enforcement (vendor-sourced)
RBAC	Yes	Yes	Yes	Yes — typed permission catalog, default-deny (vendor-sourced)
Approval gates on tool calls	Yes, per app	Yes, per tool	Yes, per tool	Yes — dual control, SLAs, fail-closed (vendor-sourced)
PII redaction before model	No	Partial (DLP)	No	Yes — on by default (vendor-sourced)
Data classification policy	No	Yes (Purview)	Partial (Purview)	Yes — per-class rules incl. private-model-only (vendor-sourced)
SOC 2 / ISO 27001 certification	SOC 2, ISO 27001	Yes (M365)	SOC 2, ISO 27001	No — ISO 27001 in progress; own-AWS infrastructure sits under your certifications

Odokai entries are vendor-sourced (checked against the vendor's repository, September 2026), not independently audited. Gaps are vendor-confirmed absences. Incumbent entries are from public vendor documentation and include certifications only where publicly claimed.

6 Objective assessment

STRENGTHS	WEAKNESSES & RISKS
ChatGPT Enterprise Most complete verified control set (SCIM, RBAC, compliance API, EKM, residency); strongest end-user brand familiarity; flexible credit-pool pricing.	ChatGPT Enterprise Quote-only pricing and a reported ~150-seat floor exclude smaller organisations; annual lock-in; reported costs vary widely between buyers.
M365 Copilot Zero-friction deployment for Microsoft tenants; permissioned grounding in tenant data; inherits the full Microsoft compliance and identity stack; unmatched in-app integration.	M365 Copilot All-in cost is 2–3× the headline \$30 once base licences are counted; July 2026 price rises; promotions end Sep and Dec 2026; value collapses outside a Microsoft estate; some features still metered or in the Frontier preview.
Claude Enterprise Lowest published seat price covering a genuinely full suite (Code, Cowork, Design, connectors); no per-seat usage limits; strong reasoning and coding reputation.	Claude Enterprise Usage billed at API rates with no included allowance makes spend hard to forecast; heavy teams can far exceed seat cost; the seat fee covers access only.
Odokai Architecture designed around teams and processes, not individual chat; three deployment modes from managed UK SaaS to customer's own AWS account; any OpenAI-compatible provider plus local models (unique among the four); persistent memory and skills accumulate organisational knowledge; PII redaction and dual-control approval gates built in; durable background work; multi-agent patterns; internal app building.	Odokai Small, young vendor — viability and key-person risk relative to trillion-dollar incumbents; no SOC 2 or ISO 27001 certification today (ISO 27001 plan in progress); much smaller ecosystem and community; single-tenant mode transfers operational responsibility to the customer. Odokai-specific facts rest on vendor's own documentation, not independent audit.


6.1 Indicative scorecard

Qualitative, 1–5, reflecting the analysis above. Weight the rows by your own priorities. Odokai's technical-control score is vendor-sourced; its certification gap is vendor-confirmed.

Dimension	ChatGPT Ent.	Copilot	Claude Ent.	Odokai
Enterprise controls & compliance maturity	●●●●●	●●●●●	●●●●●	●●●●●
Depth of office-suite integration	●●●●●	●●●●●	●●●●●	●●●●●
Agentic / workflow automation	●●●●●	●●●●●	●●●●●	●●●●●
Developer tooling	●●●●●	●●●●●	●●●●●	●●●●●
Organisational knowledge retention	●●●●●	●●●●●	●●●●●	●●●●●
Model choice / data routing control	●●●●●	●●●●●	●●●●●	●●●●●
Cost predictability	●●●●●	●●●●●	●●●●●	●●●●●

Vendor scale & ecosystem maturity	● ● ● ● ●	● ● ● ● ●	● ● ● ● ●	● ● ● ● ●
Data locality / deployment control	● ● ● ● ●	● ● ● ● ●	● ● ● ● ●	● ● ● ● ●

6.2 Capability profile at a glance



ChatGPT Enterprise

Peaks on compliance maturity; strong on agentic work and developer tooling.

M365 Copilot

Compliance and office integration are its core; agentic features (Cowork, Agent Mode) now GA.

Claude Enterprise

Strongest developer tooling and knowledge retention; cost predictability is its structural weak point.

Odokai

Strongest on agentic work, model choice and data locality; compliance maturity capped by no certification yet, with ISO 27001 in progress.

7 Fit and verdicts

7.1 Indicative fit by scenario

If your situation is...	Strongest candidates
Large organisation, Microsoft 365 estate, broad employee rollout	M365 Copilot (primary); ChatGPT Enterprise (adjacent, for non-Microsoft workflows)
Engineering-heavy teams wanting coding agents plus chat under one seat	Claude Enterprise; ChatGPT Enterprise (standard seats, which include Codex)
Regulated / document-heavy operations needing process-specific agents, data locality, model routing control and auditable persistent knowledge	Odokai (architecture fit; technical controls vendor-sourced) — weigh absent certifications; incumbents as fallback
Cost predictability is the top priority	M365 Copilot (fixed add-on), Claude Team (fixed per-seat), or Odokai (fixed fee with hard credit cap, vendor-sourced)
Vendor risk tolerance is low (procurement, board oversight)	The three incumbents; Odokai only with contractual protections given its size

BEST FOR

Broad enterprise rollout on Microsoft

- **M365 Copilot** — if the estate is Microsoft and adoption is managed properly.
- **ChatGPT Enterprise** — if the estate is mixed and the ~150-seat floor is affordable.

BEST FOR

Engineering and research teams

- **Claude Enterprise** — strongest coding and delegated work; accept usage-based billing.
- **ChatGPT Enterprise** — if fixed seat economics matter more.

BEST FOR

Regulated, document-heavy operations

- **Odokai** — architecture and technical controls fit (vendor-sourced); requires vendor due diligence given size, and no certification today.
- **Incumbents** — as the risk-averse fallback.

WATCH

Cost predictability

- Claude Enterprise bills all usage at API rates; admin spend limits can cap it.
- Odokai's credit-cap model (vendor-sourced) is fixed-fee; verify the published amount at contract time.
- Microsoft's remaining promotions end 30 Sep and 31 Dec 2026; lock terms before they close.

8 Method and sources

- OpenAI Help Center, "What is ChatGPT Enterprise?" (updated April 2026) and ChatGPT pricing pages.
- Third-party 2026 procurement reporting on ChatGPT Enterprise pricing (Coworker AI, Teamazing — reported figures, not official).
- Microsoft official M365 Copilot enterprise pricing page; Velosio and gosearch.ai 2026 licensing analyses (all-in cost modelling, July 2026 price changes).
- Anthropic official pricing and enterprise pages; Claude Help Center enterprise billing article (seat fee + usage at API rates).
- Odokai: the vendor's published platform introduction, plus corrections supplied by the vendor on 13 September 2026 checked against its own repository (deployment modes, credit-cap pricing, audit trail, RBAC, MFA, approval gates, redaction, connector directory, model routing). These Odokai-specific facts are vendor-sourced, not independently audited. Independently confirmed: no SOC 2 or ISO 27001 certification today; no public price listing found on the vendor's website at time of writing.

All pricing in USD as published or reported in September 2026, excluding tax. Vendor pricing changes frequently; treat all figures as indicative and re-verify at contract time.

Comparative assessment prepared September 2026; version 1.3 re-checks incumbent entries against vendor documentation (September 2026) and incorporates corrections supplied by Odokai, checked against its own repository. It was not commissioned by any platform named within it. The compiling assistant runs within the Odokai platform; Odokai's entries are labelled vendor-sourced and its confirmed gap (certification) is listed alongside everything else. Figures 2–4 and the radar chart are analytical visualisations of the sourced data, not vendor-supplied metrics.