Enterprise AI Assistant Comparison: Data Controls & Integrations

IntuitionLabs.ai · Adrien Laurent · 2026-09-05 · enterprise ai assistant · data controls · integrations comparison · claude enterprise · chatgpt enterprise · microsoft 365 copilot · gemini enterprise · ai governance · fedramp compliance · enterprise chatbot

Executive Summary

This report compares four enterprise AI assistants, Anthropic's **Claude Enterprise**, OpenAI's **ChatGPT Enterprise**, **Microsoft 365 Copilot**, and Google's **Gemini Enterprise**, specifically on data controls and integrations, using each vendor's own trust and pricing documentation, government compliance registries, and independent benchmarks accessed as of September 2026. All four now state that customer data is **not used to train their models by default** (anthropic.com) (support.google.com), but the depth of control underneath that promise varies sharply by plan tier. Claude Enterprise offers customer-managed encryption keys and a US-only inference option at **$20 per seat per month** plus metered usage (support.claude.com); ChatGPT Enterprise offers data residency across ten regions with Business seats from **$20 to $25 per month** and custom Enterprise pricing ; Microsoft 365 Copilot bundles Purview-based sensitivity labels, DLP, and no-extra-cost audit logging into a **$30.00 per user per month** license (microsoft.com); and Gemini Enterprise restricts customer-managed encryption keys to its top edition while documenting third-party connectors (docs.cloud.google.com) (Google Cloud Documentation).

Government authorization must be evaluated against the exact cloud service offering and certified boundary named in the FedRAMP Marketplace. On independent benchmarks, an Anthropic-based agent led SWE-bench Verified at **79.2%**, ahead of a Gemini-based agent at 77.4% and the top OpenAI submission at 72.8% (swebench.com), while Google was named a Leader in Gartner's 2026 Magic Quadrant for Conversational AI Platforms, a category where neither Anthropic, OpenAI, nor Microsoft appears as a Leader (cxtoday.com).

Adoption is large and growing across all four: Anthropic reports over 300,000 business customers (anthropic.com), OpenAI reports over 1 million (openai.com), Microsoft 365 Copilot has surpassed 30 million paid seats (microsoft.com), and Gemini Enterprise has sold over 8 million paid seats in its first four months (blog.google). That growth outpaces governance maturity: only one in five companies has a mature model for governing autonomous AI agents, per Deloitte (deloitte.com), and 64% of privacy and security professionals still worry about inadvertently sharing sensitive data through a generative AI tool, per Cisco's 2025 Data Privacy Benchmark Study (investor.cisco.com). No single assistant leads on every documented data-control or integration dimension; the right choice depends on which specific control, key ownership, residency, bundled governance, or verified government authorization, matters most for the buying organization's workflow.

Introduction and Background

Enterprise buyers evaluating an AI assistant, sometimes procured under the label enterprise chatbot or conversational AI platform, in September 2026 face four dominant options: **Anthropic's Claude Enterprise**, **OpenAI's ChatGPT Enterprise**, **Microsoft 365 Copilot**, and **Google's Gemini Enterprise**. Each vendor now publishes broadly similar marketing language, no training on customer data by default, encryption at rest and in transit, SOC 2 and ISO 27001 attestations, single sign-on (SSO). What differs, and what a procurement or security team actually needs to evaluate, is the fine print underneath that language: default retention windows, which plan tiers include audit logs, whether encryption keys can be customer-managed, which government-authorization tier a vendor has actually reached, and which third-party systems a native connector can reach without a middleware layer. This report compares the four assistants specifically on data controls and integrations, treating general capability and writing quality, the subject of an earlier IntuitionLabs comparison, as settled ground it does not repeat (intuitionlabs.ai).

Method. This comparison was built by pulling each vendor's own trust, privacy, security, and pricing documentation, plus government and independent compliance registries, and reading every page for the exact language it uses, not for a vendor's marketing summary of itself. Facts were organized into a fixed set of **data-control dimensions** (default training use of customer data, retention period, encryption and key management, access control model, audit logging, compliance certifications, and government-cloud authorization) and a fixed set of **integration dimensions** (native first-party connectors, the underlying connector protocol, identity-provider federation, and third-party ecosystem breadth). Every claim below is dated to when it was accessed or to the date the source itself states, because retention windows, certification status, and pricing change frequently. The **Performance and Benchmarks** section is different in kind: it reports scores independent evaluators published, not anything IntuitionLabs measured, and vendor-claimed figures are labeled as such throughout.

Enterprise demand for this scrutiny is not abstract. Enterprise generative AI spending grew from **$1.7 billion in 2023 to $37 billion in 2025**, according to Menlo Ventures' annual survey of enterprise decision-makers (menlovc.com), and Stanford's AI Index found the share of surveyed organizations using generative AI in at least one business function more than doubled from **33% in 2023 to 71% in 2024** (hai.stanford.edu). At the same time,

Cisco's 2025 Data Privacy Benchmark Study, based on 2,600 privacy and security professionals across 12 countries, found **64% worry about inadvertently sharing sensitive information** with a generative AI tool (investor.cisco.com). Data controls and integration depth, not headline model quality, are increasingly what decides a procurement outcome.

Claude Enterprise (Anthropic)

Capabilities

Anthropic states that **prompts, data, and results are not used to train its models by default** on commercial plans (anthropic.com). On the direct Anthropic API and other backend systems, inputs and outputs are automatically deleted within **30 days** of receipt or generation absent a special arrangement (privacy.anthropic.com). **Zero data retention (ZDR)** agreements, where nothing is stored at all, apply only to eligible Anthropic APIs, products that use a commercial organization's API key (including Claude Code), and Claude Code for Enterprise, not the full Claude Enterprise chat surface (privacy.claude.com). Enterprise customers can instead configure **custom data retention**, with a 30-day minimum where each configured month is counted as 30 days (support.claude.com).

Claude Enterprise supports **customer-managed encryption keys (CMEK)**: organizations provision a key in their own cloud provider that Anthropic then uses to protect chats, projects, and files, though enabling CMEK disables the manual "Export logs" button, pushing audit access to the Compliance API instead (support.claude.com). **Audit logs are restricted to Enterprise organizations** and are unavailable on Team or Pro plans (support.claude.com). Anthropic's SSO architecture links **each parent organization to exactly one identity provider** across all linked sub-organizations, a constraint multi-brand or multi-subsidiary customers should plan around (support.claude.com). Enterprise plans also support **custom roles** for group-level feature access beyond the standard User, Admin, Owner, and Primary Owner tiers (support.claude.com).

On compliance, Anthropic lists **SOC 2, ISO 27001, GDPR, and CCPA** on its enterprise page (anthropic.com) and separately holds **ISO/IEC 42001:2023**, the AI management-systems standard (support.claude.com). For government use cases, buyers should confirm the exact product, boundary, and authorization record before relying on a regulatory claim.

Integrations run through native connectors to **Google Drive, Gmail, Google Calendar, GitHub, Microsoft 365, and Slack** for retrieving workplace context without manual uploads (support.claude.com), and Anthropic's broader connector directory is explicitly built on the **Model Context Protocol (MCP)**, the open standard it originated (claude.com). Enterprise plans also offer a **US-only inference** option that keeps processing within the United States. One limitation for regulated buyers: Anthropic's Business Associate Agreement (BAA) for HIPAA explicitly **excludes Claude Console and Claude Cowork**, and data sent to third parties through connectors is **not covered** by the BAA even when the base product is (support.claude.com) (support.claude.com).

Adoption

Anthropic said in September 2025 that it serves **over 300,000 business customers**, and that its count of large accounts, those representing more than $100,000 in annual run-rate revenue each, had grown **nearly 7x year over year** (anthropic.com). Pricing for Claude Enterprise is **$20 per seat per month** plus separate, metered token

consumption. Notably, the Enterprise seat fee **covers platform access only**, with all Chat, Claude Code, and Cowork usage billed at standard API rates on top (support.claude.com), a cost structure buyers should model explicitly rather than compare on seat price alone.

Strengths and Limitations

Claude's documented strengths are concentrated in encryption control (self-managed CMEK), an explicit US-only inference option, and an MCP-native connector architecture that other vendors have since converged toward. Its main documented limitations for regulated buyers are the BAA carve-outs for Console and Cowork and the single-identity-provider SSO ceiling.

ChatGPT Enterprise (OpenAI)

Capabilities

OpenAI states that **by default it does not use business data to train its models** for ChatGPT Business, Enterprise, Edu, Healthcare, Teachers, or API usage after March 1, 2023. Deleted conversations are removed from OpenAI's systems within **30 days**, unless legally required to retain them ; the OpenAI API separately allows organizations to request **zero data retention (ZDR)** for eligible endpoints and qualifying use cases (openai.com). Data is encrypted with **AES-256 at rest and TLS 1.2 or higher in transit** (openai.com), and **Enterprise Key Management (EKM)** lets customers bring their own encryption keys for content stored at rest. **Data residency**, storing customer content at rest in-region, is available across **ten regions**: Europe, the UK, the US, Canada, Japan, South Korea, Singapore, India, Australia, and the UAE.

On compliance, OpenAI states its practices adhere to **CSA STAR, SOC 2 Type 2, and ISO/IEC 27001**, alongside ISO 27017, 27018, and 27701. OpenAI can execute a **HIPAA Business Associate Agreement** in support of customer compliance (openai.com). Workspace admins can pull an **audit log of conversations and GPTs** through the Enterprise Compliance API (openai.com), whose Compliance Logs Platform retains data for **30 days** unless a customer configures continuous export (help.openai.com) and integrates with third-party eDiscovery, DLP, and SIEM tools including **Microsoft Purview, CrowdStrike, Netskope, Palo Alto Networks, Zscaler, and Varonis** (help.openai.com).

SCIM provisioning is available for eligible Enterprise and Edu workspaces, syncing with **Okta, Microsoft Entra ID, Google Workspace, PingFederate, OneLogin, and Rippling**, but a **standalone ChatGPT Business plan and ChatGPT for Teachers do not include SCIM** (help.openai.com) (help.openai.com). Integrations center on **Company knowledge**, which lets ChatGPT pull context from connected apps such as Slack, SharePoint, Google Drive, and GitHub while **respecting each user's existing permissions**, so ChatGPT can only surface what that user is already authorized to view (openai.com), with Enterprise and Edu admins able to manage which connected apps are available at the group level (openai.com).

Adoption

OpenAI said in November 2025 that **more than 1 million business customers** now use OpenAI directly, describing it as the fastest-growing business platform in history (openai.com), with total **ChatGPT for Work seats surpassing 7 million**, up 40% in two months, and Enterprise seats growing **9x year over year** (openai.com).

Pricing for **ChatGPT Business** is **$20 per month billed annually** ($25 billed monthly) for a standard seat, with a premium seat at $100 ($125 monthly), while **ChatGPT Enterprise pricing is custom** and quoted directly through sales rather than published per seat (openai.com). The Enterprise plan bundle is described as including **SCIM, EKM, user analytics, domain verification, and role-based access controls** (openai.com), and even the lower Business tier connects to **Google Workspace, Slack, GitHub, and Microsoft 365** out of the box (openai.com).

Strengths and Limitations

ChatGPT Enterprise's documented strengths are its ten-region data residency footprint, the breadth of its Compliance Logs Platform's SIEM/DLP integrations, and its documented enterprise data controls. Its clearest documented limitation is that core identity controls, SCIM in particular, are unavailable below the Enterprise or Edu tier, meaning smaller Business-plan customers cannot enforce automated deprovisioning through their identity provider.

Microsoft 365 Copilot

Capabilities

Microsoft states that **prompts, responses, and data accessed through Microsoft Graph are not used to train foundation LLMs**, and that Copilot **only surfaces organizational data to which the individual user already has at least view permission** (learn.microsoft.com), a permission-inheritance model shared conceptually with OpenAI's Company knowledge and Google's Gemini. Tenant isolation is enforced through **Microsoft Entra authorization and role-based access control**, with content encrypted at rest and in transit using **BitLocker, per-file encryption, TLS, and IPsec**.

Copilot's security, legal, and contractual considerations should be assessed separately by product and service scope. Where Microsoft's integration with its own **Purview** governance suite is distinctive: sensitivity-labeled content requires the **EXTRACT usage right in addition to VIEW** before an AI app can return it (learn.microsoft.com), Purview **data loss prevention (DLP)** applies deep content inspection to AI interactions (learn.microsoft.com), and Copilot prompts, responses, and the files accessed are captured in the **unified audit log** (learn.microsoft.com). Unlike the compliance add-ons other vendors require, Microsoft includes audit logging for **its own** Copilot products at no extra Pay-as-you-go cost under Audit Standard, a distinction it draws explicitly against non-Microsoft AI applications (learn.microsoft.com).

Copilot Studio, Microsoft's low-code agent-building layer, is separately covered under a **HIPAA BAA** (learn.microsoft.com) and has been **independently audited for SOC** compliance (learn.microsoft.com). For government customers, FedRAMP Marketplace records identify specific Microsoft cloud service offerings rather than a blanket status for government cloud services (FedRAMP Marketplace) (FedRAMP Marketplace), and Microsoft 365 Copilot itself is available in **GCC, GCC High, and DoD** environments, operating entirely inside the customer's government tenant (learn.microsoft.com); GCC High specifically targets customers whose contracts require **FedRAMP High, DFARS, or ITAR/EAR** compliance (learn.microsoft.com).

Integrations run through **Microsoft Graph connectors**, of which Microsoft and its partners publish **more than 100 prebuilt** connectors (learn.microsoft.com), extending Copilot to data beyond Microsoft 365 via two mechanisms: **synced connectors**, which index content into Microsoft Graph, and **federated connectors**, which fetch data in

real time using the **Model Context Protocol (MCP)** without indexing it at all (learn.microsoft.com) (learn.microsoft.com).

Adoption

Microsoft said in its July 2026 momentum update that **Microsoft 365 Copilot has surpassed 30 million paid seats**, with net new seat additions **more than doubling quarter over quarter** (microsoft.com), and that the number of customers with **more than 50,000 seats grew more than 7x year over year** (microsoft.com). List pricing is **$30.00 per user per month** paid yearly for the full Microsoft 365 Copilot license (microsoft.com), while a lighter **Copilot Chat** is included at **no additional cost** for Microsoft Entra account holders on an eligible Microsoft 365 subscription (microsoft.com); **Copilot Studio** capacity is metered separately at **$200 per pack per month for 25,000 Copilot Credits** (microsoft.com).

Strengths and Limitations

Copilot's most distinctive documented strength is depth of integration with Microsoft's own governance stack, Purview sensitivity labels, DLP, and unified audit logging apply to Copilot with no separate compliance product to buy, a bundling advantage a standalone AI-only vendor cannot replicate. FedRAMP Marketplace records for Microsoft 365 Government Community Cloud & Supporting Services and Microsoft 365 Government Community Cloud-High describe separate cloud service offerings and should not be read as a general Microsoft 365 Copilot certification (FedRAMP Marketplace) (FedRAMP Marketplace).

Gemini Enterprise (Google)

Capabilities

Google states it **does not use Workspace data to train or improve the underlying generative AI and large language models** that power Gemini (support.google.com); for Google AI plan users specifically, Gemini in Gmail, Calendar, Chat, Docs, Drive, Sheets, Slides, Meet, and Vids **does not use content to train or improve Gemini or other generative AI models** (support.google.com), and it **does not store the prompt or the generated output without the user's permission**.

Google documents security and compliance information for Agent Search separately; buyers should verify the applicable product, service boundary, and contractual scope. Encryption is handled by Google by default, but **customer-managed encryption keys (CMEK)** via Cloud KMS are available for organizations that want direct control of key location, rotation, and access (docs.cloud.google.com); notably, **CMEK is available only in the Enterprise Edition** (docs.cloud.google.com), and it requires apps and data stores to sit in **US or EU multi-region locations**, not global ones (docs.cloud.google.com). Securing a deployment against data exfiltration requires configuring a **VPC Service Controls** perimeter; once enabled, Gemini Enterprise's agent "actions", such as sending an email or creating a support ticket, are **blocked by default** until Google allowlists each service, and a perimeter **cannot be applied to a project with pre-existing data stores**, which must instead be deleted and recreated (docs.cloud.google.com) (docs.cloud.google.com).

Gemini Enterprise ships in **five editions**, Business (1 to 500 users), Standard, Plus, Pay-as-you-go, and Frontline (which requires 150 or more Standard or Plus seats), each with a different pooled storage allowance and feature set (docs.cloud.google.com); notably, the entry-level **Business edition is the only one that lacks "enterprise-grade security and compliance"** on Google's own edition-comparison table, a gap the Standard, Plus, Pay-as-you-go, and Frontline tiers all close (docs.cloud.google.com). Connector coverage is broad: generally available connectors include **Box, Confluence, Dropbox, GitHub, Google's own apps, HubSpot, Jira, Microsoft Entra ID, Outlook, Teams, OneDrive, SharePoint, Monday, ServiceNow, and Slack**, with Salesforce, Notion, Zendesk, Asana, GitLab, Stripe, Dynamics 365, Docusign, and custom **MCP servers** in public preview (docs.cloud.google.com).

At the access-control layer, Workspace admins can restrict a user's access to Gemini **entirely**, or scope Gemini so it can still be used for chat while its **access to Workspace data (Drive, Gmail) is restricted** separately (support.google.com); by default, **Gemini inherits the same data access as the user invoking it**, so it cannot reach Drive files that were not shared with that user, files whose owner disabled download and print, or another person's delegated inbox (support.google.com). For public-sector buyers, the FedRAMP Marketplace identifies Google Workspace and Gemini for Government as separate cloud service offerings (FedRAMP Marketplace) (FedRAMP Marketplace), and Gemini is **included in Workspace Enterprise plans** without a separate paid AI add-on (cloud.google.com).

Adoption

On Alphabet's Q4 2025 earnings call in February 2026, CEO Sundar Pichai said Google had sold **more than eight million paid seats of Gemini Enterprise** to more than 2,800 companies, four months after launch, with Gemini Enterprise managing **over 5 billion customer interactions** in the quarter, growing **65% year over year** (blog.google). Pichai separately stated that **more than 120,000 enterprises use Gemini** overall, and that the standalone consumer **Gemini App had surpassed 750 million monthly active users** (blog.google). On the developer side, **Gemini Code Assist** is billed hourly rather than per seat, at **$0.031232877 per hour** for Standard under a monthly commitment ($0.026027397 with a 12-month commitment) and roughly double that for Enterprise (cloud.google.com).

Strengths and Limitations

Gemini Enterprise's clearest documented strength is breadth: a documented third-party connector catalog and more than eight million paid seats by Google's own account (Google). Its most consequential documented limitation is architectural rather than a missing certification: enterprise-grade security and compliance is absent from the entry Business edition, and CMEK, the encryption control regulated buyers most often require, is confined to the Enterprise edition and to US or EU multi-region deployments only.

Feature Comparison

Table 1 below summarizes the data-control and integration posture documented for each vendor's flagship enterprise plan as of September 2026, drawn from the primary sources cited throughout the sections above.

Dimension	Claude Enterprise (Anthropic)	ChatGPT Enterprise (OpenAI)	Microsoft 365 Copilot	Gemini Enterprise (Google)
Default training use of customer data	Not used by default (anthropic.com)	Not used by default, incl. connected apps (openai.com)	Prompts/responses/Graph data not used (learn.microsoft.com)	Not used to train/improve Gemini or LLMs (support.google.com)
Default backend retention	30 days (privacy.anthropic.com); custom retention configurable (30-day min) (support.claude.com)	30 days after deletion, unless legally required otherwise (openai.com)	Governed via Purview/unified audit log, tenant-configured (learn.microsoft.com)	No prompt/output storage without permission (support.google.com)
Customer-managed encryption keys	Yes, via own cloud provider (support.claude.com)	Yes, Enterprise Key Management (openai.com)	Bundled via Microsoft 365 service-side encryption (learn.microsoft.com)	Yes, Enterprise edition only, US/EU multi-region only (docs.cloud.google.com)
SSO / SCIM	SSO (one IdP per parent org) (support.claude.com)	SCIM on Enterprise/Edu only, not Business/Teachers (help.openai.com)	Entra ID-based RBAC and tenant isolation (learn.microsoft.com)	Admin console access restrictions, granular per-user (support.google.com)
Security and compliance information	SOC 2, ISO 27001, ISO 42001, GDPR, CCPA (anthropic.com) (support.claude.com)	SOC 2 Type 2, ISO 27001/17/18/27701, CSA STAR (openai.com)	GDPR, ISO 27018/27001/42001, HIPAA (learn.microsoft.com)	SOC 1/2/3, ISO 27001/17/18/27701, PCI DSS, BSI C5 (docs.cloud.google.com)
Government cloud records	Evaluate the exact FedRAMP Marketplace offering and boundary; do not infer assistant-product coverage from an offering record.	Evaluate the exact FedRAMP Marketplace offering and boundary; do not infer assistant-product coverage from an offering record.	Evaluate the exact FedRAMP Marketplace offering and boundary; do not infer assistant-product coverage from an offering record.	Evaluate the exact FedRAMP Marketplace offering and boundary; do not infer assistant-product coverage from an offering record.
Native connector protocol	Model Context Protocol (MCP) (claude.com)	Company knowledge, permission-scoped connectors (openai.com)	Synced connectors + federated MCP connectors (learn.microsoft.com)	First- and third-party connectors + preview MCP servers (docs.cloud.google.com)
Entry-tier list price	$20/seat/month + usage (support.claude.com)	Business: $20-25/month/seat; Enterprise custom (openai.com)	$30.00/user/month annual; Chat free with eligible sub (microsoft.com)	Business edition 1-500 users; edition-dependent pricing (docs.cloud.google.com)

The pattern across the table is convergence at the level of stated policy, and divergence at the level of implementation detail. All four vendors now assert the same baseline promise (no training on enterprise data by default), but they differ meaningfully in where encryption control sits (which plan tier, which region), which identity and audit features require the top-priced tier, and how far each has actually progressed through a formal government authorization rather than simply stating an intent to pursue one.

Performance and Benchmarks

The comparisons in this section are entirely **independent, third-party measurements**, not figures IntuitionLabs generated, and not vendor-claimed benchmark scores; they should be read alongside the documented-capability comparison above rather than in place of it, since a model's coding accuracy says nothing about its data-retention policy.

On **SWE-bench Verified**, an independent benchmark of 500 real-world software issues, the top-scoring submission built on Anthropic's model resolved **79.2%** of issues (Sonar Foundation Agent plus Claude 4.5 Opus), ahead of a Gemini-based agent at **77.4%** (live-SWE-agent plus Gemini 3 Pro Preview) and the top OpenAI-model submission at **72.8%** (GPT-5.2, high reasoning effort) (swebench.com) (swebench.com) (swebench.com). On **Artificial Analysis's Intelligence Index**, a composite benchmark aggregator, Anthropic's top model scored **57** and OpenAI's top model scored **55**, the two highest-rated models overall on the index, while Google's best-scoring model reached only **47** (artificialanalysis.ai) (artificialanalysis.ai) (artificialanalysis.ai). On **LMArena's** community-voted leaderboards, an Anthropic model held the **#1 position on both the Agent leaderboard** (13.74% score) and the **general Text leaderboard** (1507 rating) as of the access date (lmarena.ai) (lmarena.ai).

Independent analyst coverage is more mixed on category leadership than raw model benchmarks suggest. **Gartner's 2026 Magic Quadrant for Conversational AI Platforms**, which evaluates 14 vendors, names **Google as a Leader** alongside Salesforce, Kore.AI, and SoundHound AI; neither Anthropic, OpenAI, nor Microsoft appears as a Leader in that specific quadrant, which scores conversational-platform maturity rather than model quality (cxtoday.com) (cxtoday.com). **Forrester's** inaugural Wave evaluation of AI foundation-model providers assessed Anthropic, Google, Microsoft, and OpenAI head-to-head across **21 criteria** alongside six other vendors (forrester.com).

On transparency specifically, an area directly relevant to data-control claims, Stanford's **Foundation Model Transparency Index (FMTI)**, December 2025 edition, found the mean transparency score across 13 companies **fell to 41, a 17-point year-over-year decline** industry-wide (crfm.stanford.edu). Within that decline, **Anthropic rose to the top ranks** among frontier labs tracked since 2023, while Meta and OpenAI **fell from 1st and 2nd place in 2023 to 5th and 6th in 2025**, with OpenAI's score specifically down **14 points** year over year (crfm.stanford.edu) (crfm.stanford.edu); IBM, not one of the four assistants compared here, scored highest of all companies assessed at **95 out of 100** (crfm.stanford.edu).

FedRAMP status should be verified on the Marketplace for the exact cloud service offering and certified boundary; an offering record should not be generalized to an assistant product or a different government service.

Data Analysis and Evidence

Enterprise generative AI spending reached **$37 billion in 2025**, up from $1.7 billion in 2023, according to Menlo Ventures' survey of 495 US enterprise decision-makers, conducted in November 2025. Within that spend, **horizontal AI**, the category covering general-purpose copilots and agent platforms such as the four assistants compared here, was the **largest and fastest-growing** application segment at **$8.4 billion**, up **5.3x year over year** (menlovc.com). Buying behavior shifted sharply toward off-the-shelf tools: **76% of enterprise AI use cases were purchased rather than built internally** in 2025, up from a roughly even split in 2024 (menlovc.com).

Table 2 below summarizes the self-reported adoption metrics each vendor disclosed for its enterprise assistant, alongside the independent benchmark and authorization data from the section above; figures are vendor-disclosed unless otherwise noted and should not be treated as directly comparable across different reporting windows.

Vendor	Reported paid seats / customers	Reporting date	Entry price (list)
Claude Enterprise	300,000+ business customers; large accounts ($100k+ run-rate) up ~7x YoY (anthropic.com)	September 2025	$20/seat/month + usage (support.claude.com)
ChatGPT Enterprise / Work	1M+ business customers; 7M+ total Work seats, +40% in 2 months; Enterprise seats +9x YoY (openai.com) (openai.com)	November 2025	$20-25/seat/month (Business); Enterprise custom (openai.com)
Microsoft 365 Copilot	30M+ paid seats; 50,000+-seat customers up 7x YoY (microsoft.com) (microsoft.com)	July 2026	$30.00/user/month (microsoft.com)
Gemini Enterprise	8M+ paid seats across 2,800+ companies; 120,000+ enterprises use Gemini overall (blog.google) (blog.google)	February 2026	Edition-dependent (Business through Frontline) (docs.cloud.google.com)

At the model-provider level rather than the assistant-product level, Menlo Ventures found **Anthropic, OpenAI, and Google together account for 88% of enterprise LLM API usage**, with Anthropic at **40%**, OpenAI at **27%**, and Google at **21%**, and Anthropic specifically holding an estimated **54% share of the enterprise coding-model market** versus **21%** for OpenAI (menlovc.com) (menlovc.com). Separately, Microsoft's own connector gallery lists **more than 100 prebuilt connectors** built by Microsoft and partners, the only vendor among the four to publish a specific integration count (learn.microsoft.com).

Governance maturity has not kept pace with adoption. Deloitte's State of AI in the Enterprise research found that **only one in five companies has a mature governance model** for overseeing autonomous AI agents, even as worker access to AI tools rose **50% in 2025** (deloitte.com) (deloitte.com). Cisco's 2025 Data Privacy Benchmark Study, covering 2,600 privacy and security professionals in 12 countries, found **63% now describe themselves as "very familiar"** with generative AI, yet **64% still worry about inadvertently sharing sensitive information** publicly or with a competitor through an AI tool (investor.cisco.com) (investor.cisco.com) (investor.cisco.com). Multi-vendor deployment is now the norm rather than the exception: a16z's 2025 survey of 100 enterprise CIOs across 15 industries found **37% now use five or more AI models** in production or experimentation, up from **29%** a year earlier (a16z.com) (a16z.com), a pattern consistent with the four-way comparison this report undertakes rather than a single winner-take-all assistant choice.

Implications and Future Directions

Three trends run through the data above and are likely to shape how enterprises choose and combine these assistants over the next planning cycle. First, **data-control depth is migrating into pricing tiers, not just into the product overall**. Gemini Enterprise's Business edition explicitly excludes "enterprise-grade security and compliance" from its own comparison table (docs.cloud.google.com), OpenAI reserves SCIM for Enterprise and Edu plans only (help.openai.com), and Anthropic restricts audit logs to Enterprise organizations

(support.claude.com). A buyer comparing "Claude" to "ChatGPT" to "Copilot" to "Gemini" in the abstract is comparing marketing surfaces; the actual comparison has to happen at the specific plan tier the organization intends to purchase.

Second, **agent permissioning is becoming the primary new control surface**, ahead of the older retention-and-training question. Google's default behavior of blocking Gemini Enterprise agent actions until each connected service is explicitly allowlisted under VPC Service Controls (docs.cloud.google.com) is a visible example of vendors building safety defaults around what an agent is permitted to *do*, not just what data it can *see*. As all four vendors converge on the Model Context Protocol as a shared connector standard (claude.com) (learn.microsoft.com) (docs.cloud.google.com), the meaningful integration differentiator is shifting from "can it connect" to "what does it need explicit permission to do once connected."

Third, **federal and regulated-industry authorization remains uneven and is a legitimate tie-breaker** for buyers in government, defense, or highly regulated sectors. The FedRAMP Marketplace comparison above shows real variance, not marketing variance, between a "Not yet certified" listing, Microsoft 365 Government Community Cloud-High's Class D (High) certification with five authorizations, and the separate Microsoft 365 Government Community Cloud & Supporting Services offering's Class C (Moderate) certification with 99 authorizations. For life-sciences organizations specifically, the pattern that matters most is not FedRAMP but the HIPAA BAA carve-outs each vendor documents: Anthropic excludes Claude Console and Claude Cowork from its BAA and does not cover data sent through third-party connectors (support.claude.com) (support.claude.com), while Microsoft's Copilot Studio and OpenAI's ChatGPT Enterprise each publish their own, differently scoped, BAA terms (learn.microsoft.com). A pharmaceutical or clinical-operations team evaluating any of these assistants for a workflow that will touch protected health information needs the BAA's exact feature list, not the vendor's general compliance-certification list, before scoping a pilot.

That distinction, between what a page documents in general and what applies to a specific regulated workflow, is where an adjacent life-sciences AI advisory practice like IntuitionLabs positions its own work: connecting an AI assistant to a company's authoritative internal sources with defined identity, permissions, retrieval, and evaluation, rather than treating a vendor's default settings as sufficient for a validated environment (intuitionlabs.ai). IntuitionLabs, which publishes this comparison, is a life-sciences AI consultancy and an official Veeva Vault CRM X-Pages partner rather than a vendor of any of the four assistants compared above, and it does not appear as an option in the tables in this report (intuitionlabs.ai). That framing does not change the underlying vendor comparison above; it changes how an organization operationalizes whichever assistant it selects.

Conclusion

Data-control, certification, and BAA terms vary by product, plan, service boundary, and contract. The differences that matter to a procurement decision sit one level down: which plan tier includes audit logging and SCIM, whether encryption keys can be customer-managed and in which regions, how far each vendor has actually progressed through FedRAMP authorization rather than merely announcing intent, and which BAA exclusions apply to which product surface. On integrations, all four have converged on a similar connector philosophy, permission-scoped access to the systems a user already has, increasingly built on the shared Model Context Protocol, with Gemini Enterprise documenting third-party connectors and Microsoft and partner organizations providing more than 100 prebuilt connectors that connect to popular Microsoft and non-Microsoft services (Microsoft Learn). On independent, non-vendor performance measures, Anthropic's models currently lead the cited SWE-bench, Artificial

Analysis, and LMArena results, while Google leads the one analyst quadrant examined here that scores conversational-platform maturity rather than raw model capability, and FedRAMP classifications apply to individual cloud service offerings: ChatGPT Enterprise and API Platform is Class C (Moderate) (FedRAMP Marketplace), whereas Microsoft 365 Government Community Cloud-High is Class D (High) (FedRAMP Marketplace). None of these findings displaces the earlier, broader IntuitionLabs comparison of these four platforms; they extend it with the specific, dated evidence a security or procurement reviewer needs to defend a choice in September 2026, and that evidence will keep moving: every retention window, certification, and authorization class cited above should be re-verified against the vendor's live documentation before it anchors a contract.

Frequently Asked Questions (FAQs)

Which enterprise AI assistant is best for data security?

No single vendor leads on every documented dimension. Claude Enterprise offers customer-managed encryption keys and a US-only inference option (support.claude.com) (support.claude.com); ChatGPT Enterprise offers the broadest data-residency footprint at ten regions plus Enterprise Key Management (openai.com) (openai.com); Microsoft 365 Copilot bundles sensitivity-label enforcement and audit logging into its native Purview stack at no extra cost (learn.microsoft.com); and Gemini Enterprise provides VPC Service Controls for agent-action restriction (cloud.google.com) (docs.cloud.google.com). The right choice depends on which specific control, key ownership, residency, bundled governance, or action-level restriction, matters most for a given workflow.

How do enterprise AI assistant integrations compare?

All four now support a similar first-party connector set to common productivity and collaboration systems (Google Workspace, Microsoft 365, Slack, GitHub) and are converging on the Model Context Protocol as a shared standard for third-party and custom connectors (claude.com) (learn.microsoft.com) (docs.cloud.google.com). Microsoft and partner organizations provide more than 100 prebuilt connectors that connect to popular Microsoft and non-Microsoft services (Microsoft Learn).

What compliance and access controls should a buyer look for?

At minimum: an explicit written statement of default training use of data, a stated default retention window, whether customer-managed encryption keys are available and at which plan tier, whether audit logs are included or a paid add-on, and the vendor's exact FedRAMP authorization class and date, not just a stated intent to pursue authorization (fedramp.gov) (fedramp.gov).

Is any of these assistants suited to life sciences specifically?

BAA availability and exclusions must be verified for the specific product and intended workflow; Anthropic's excludes Console and Cowork and third-party connector traffic (support.claude.com), so a life-sciences buyer should map the intended workflow (clinical documentation, regulatory writing, commercial analytics) against the specific BAA's feature list rather than the vendor's general compliance-certification page.