IntuitionLabs

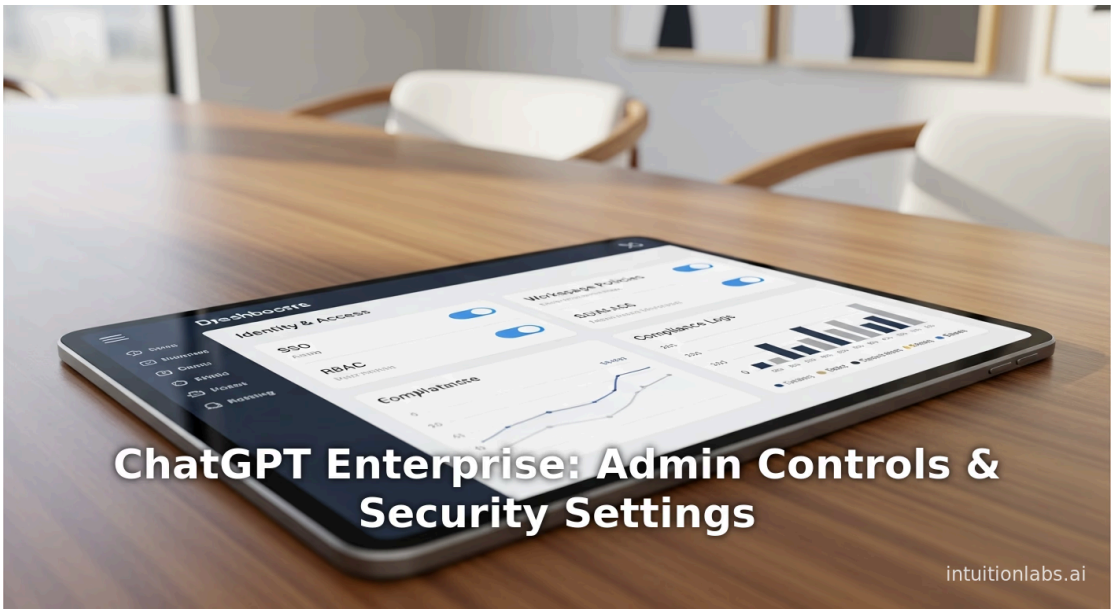
INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

ChatGPT Enterprise: Admin Controls & Security Settings

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-04-18 · chatgpt enterprise · admin controls · ai governance · data residency · enterprise security · rbac · sso integration · compliance logs · workspace management

Original article: https://intuitionlabs.ai/articles/chatgpt-enterprise-admin-controls-security



In this article

1. Executive Summary
2. Introduction and Background
3. ChatGPT Enterprise and Plan Comparison
4. Identity and Access Management
5. Workspace Administration and RBAC
6. Data Security, Privacy, and Compliance Controls
7. Integrations, Customization, and Tools
8. Case Studies and Real-World Examples
9. Security, Governance, and Risk Perspectives
10. Implications and Future Directions

Executive Summary

ChatGPT Enterprise (as of April 2026) is OpenAI's premium offering designed to meet the needs of large organizations, providing robust administrative controls, security, and compliance features beyond those in consumer and small-business plans. It builds on the capabilities of the base ChatGPT platform (powered by GPT-4 and GPT-5 models) but adds **enterprise-grade security** (e.g. SOC 2 compliance, AES-256 encryption at rest and TLS 1.2+ encryption in transit (OpenAI), data residency, and optional customer-managed encryption keys), **identity and access management** (domain verification, SSO/SAML, SCIM provisioning, IP allowlists), and **fine-grained governance** (workspace policies, audit logs, and role-based access control) (openai.com). Workspace Owners can assign custom roles. Administrators can centrally manage workspaces via a console, enable or disable specific apps (integrations) and actions, and export detailed usage/compliance logs for analysis (help.openai.com). For example, OpenAI's own documentation emphasizes that with Enterprise "you do not train on your business data or conversations," all conversations are "encrypted in transit and at rest," and a specialized **Global Admin Console** provides "domain verification, SSO, and usage insights" for large-scale deployment (openai.com) (help.openai.com). Industry reports confirm the rapid enterprise adoption of ChatGPT: by 2026 ChatGPT reportedly has over **50 million paying users** and ~**900 million weekly active users**, with enterprise contracts contributing roughly **40% of OpenAI's revenue** (www.techradar.com). Real-world evidence suggests significant productivity impact – for instance Asana's data team reported that ChatGPT Enterprise "cut down research time by an average of an hour per day" (openai.com).

In this report, we present an **in-depth survey** of ChatGPT Enterprise's options and controls. We cover its architecture and features, the configurable security and data governance settings, integration options (such as third-party "apps" and custom GPTs), and administrative workflows (identity, roles, billing). We include tables comparing enterprise vs other plans, links to the latest OpenAI documentation, illustrative case studies (e.g. enterprise deployments at Asana, Klarna), and analysis of security/compliance considerations. Wherever possible, claims are backed by official sources or independent analysis (openai.com) (www.techradar.com) (beyondscale.tech). We also consider future directions (e.g. the new Global Admin Console, evolving compliance logs API, broader data-residency) and discuss implications for enterprise IT and governance.

Introduction and Background

Since its public release in late 2022, ChatGPT has seen explosive growth across consumer and professional settings. By early 2024, OpenAI reported **millions of daily users** noting widespread use for tasks from drafting emails to coding. Enterprises quickly adopted the technology informally – one analysis claims roughly *92% of Fortune 500 companies* had employees using ChatGPT by late 2025 (www.humai.blog) – raising concerns about data privacy, compliance, and control within regulated organizations. In response, OpenAI launched **ChatGPT Enterprise** in August 2023 as "another step towards an AI assistant for work that helps with any task" and "protects your company data" (openai.com). The Enterprise edition was explicitly marketed as delivering "enterprise-grade security & privacy" and the **"most powerful version of ChatGPT yet"**, with unlimited higher-speed GPT-4 (and later GPT-5) access, 32k context windows, and advanced data analysis (Code Interpreter) (openai.com).

Key differentiators of ChatGPT Enterprise (versus consumer or basic plans) include **data controls and guarantees**: for example, OpenAI commits that "we do not train our models on your data by default," and that users "own and control" their inputs/outputs (openai.com). Enterprise customers can specify **data residency** (choosing one of 10 global regions for data storage) and even use **customer-managed encryption keys (EKM)** to encrypt all content (help.openai.com) (www.techradar.com). Administratively, the Enterprise plan introduces a secure multi-user environment (workspaces) with granular permissions. A new **Global Admin Console** (announced in late 2025) allows an organization to centrally manage multiple workspaces and deployments across ChatGPT and OpenAI's API platform, including domain verification and SSO configuration spanning all accounts (help.openai.com).

The result is that, by April 2026, ChatGPT Enterprise has become a major force in AI adoption. OpenAI reports ~50 million paying customers and the company's *$2 billion/month* revenue reflects this enterprise demand – about 40% of revenue comes from business clients (www.techradar.com). Tech Radar notes that OpenAI's APIs now process **15 billion tokens per minute**, underscoring massive usage by enterprises and developers (www.techradar.com). This popularity has led to an ecosystem of third-party consulting, security auditing, and integration with enterprise compliance tools (e.g. Microsoft Purview (learn.microsoft.com)). In the sections that follow, we explore **all aspects** of ChatGPT Enterprise configuration and control, with liberal citation of official documentation (mainly OpenAI's Help Center and blog) and industry analysis.

ChatGPT Enterprise and Plan Comparison

ChatGPT Enterprise builds on OpenAI's core ChatGPT platform but adds features tailored to organizational needs.Table 1 contrasts the *key administrative and security features* available in ChatGPT Enterprise relative to the standard Business (formerly "Team") plan and individual/Plus plans. (Note: OpenAI renamed "ChatGPT Team" to "ChatGPT Business" in Aug 2025; this table treats Business as the mid-range subscription.) Only the Enterprise plan offers the full suite of governance controls (domain SSO, SCIM, RBAC, compliance logs, etc.) identified below (openai.com) (help.openai.com).

Feature / Capability	ChatGPT Enterprise	ChatGPT Business (Team)	Free/Plus Personal
SSO & Domain Verification	Yes – full SAML/OIDC support; domain ownership required for SSO (help.openai.com)	Yes – supports SAML SSO (with domain verification)	No
SCIM User Provisioning	Yes – SCIM directory sync (Okta, Azure AD, etc.) to auto-provision users/groups (help.openai.com)	Limited – manual/user-managed (no SCIM)	No
Role-Based Access Control (RBAC)	Yes – create custom roles/groups, granular feature permissions (help.openai.com)	No (single default admin/member roles only)	No
IP Allowlisting	Yes – restrict login access to whitelisted IP addresses (help.openai.com)	No	No
Regional Data Residency	Eligible Enterprise workspaces can select regional storage for in-scope customer content; some metadata and processing may remain outside that region (OpenAI Help Center)	No Business regional-storage entitlement is established by the cited eligibility documentation	No
Customer-Managed Encryption Keys (EKM)	Available to Enterprise and Edu workspaces with a named OpenAI account representative (EKM overview)	No	No
Privacy/Data Handling	Full enterprise commitments: no training on your data by default; AES-256/TLS encryption; SOC 2, ISO 27001 compliance (openai.com)	Business inputs and outputs are excluded from training by default; data is encrypted in transit and at rest (OpenAI Help Center)	Data may be used (unless user opts out)
Analytics & Audit Logs	Yes – workspace usage analytics; Compliance Logs API (export immutable chat/auth logs) (help.openai.com)	No – no built-in audit/log export beyond basic usage stats	No
Model Access & Context Limits	Verify enabled models in workspace settings and consult model-specific limits (OpenAI Help Center)	Check the selected model's documentation and the applicable plan terms	Model defaults and limits can change; consult the Free Tier FAQ and the documentation for the selected plan
Custom GPTs & Plugins	Yes – create and manage internal GPTs; enterprise may vet or restrict public GPTs; enterprise code/plugins isolated (help.openai.com)	Yes – also available, but fewer admin controls to review public GPTs	Yes (Plus has custom GPTs)
3rd-Party "Apps" & Connectors	Enabled/Disabled by admin (off by default); full control per-app and per-action (help.openai.com)	Enabled by default; limited admin blocking	Enabled by default (user opt-in through store)
Workspace Roles	Multiple built-in roles (Owner, Admin, custom) with seats controls (help.openai.com); allocate seats via SCIM	Owners/Admins vs members; no custom roles; seats by invite only	N/A

Feature / Capability	ChatGPT Enterprise	ChatGPT Business (Team)	Free/Plus Personal
Seat Types & Spend Controls	Yes – dedicated “ChatGPT” vs “Codex” seat types; admins set defaults and budget/spend controls (help.openai.com)	April 2026: standard ChatGPT subscription seats and usage-billed Codex-only seats (Business release notes)	N/A
AI Policy Modal	Yes – define an organization-wide AI usage policy; shown to users on login update (help.openai.com)	No	No

Table 1: Comparing key administrative, security, and feature differences between ChatGPT Enterprise and other plans (data from OpenAI documentation and announcements). Sources: OpenAI help docs and blogs (openai.com) (help.openai.com) (www.techradar.com).

The table highlights that **only the Enterprise tier** offers true enterprise-grade controls like domain SSO, SCIM provisioning, RBAC, IP whitelisting, dedicated data-residency/keys, and compliance log exports. Non-Enterprise plans (Business or personal) lack these features or offer them in much more limited form. For example, SCIM auto-provisioning is an Enterprise-only feature (help.openai.com), and while Plus or Business users can create chat logs, only Enterprise customers can export detailed **immutable audit logs** through the Compliance API (help.openai.com). Check model availability against the workspace’s enabled options and the documentation applicable to the deployment date (OpenAI Help Center).

Identity and Access Management

A cornerstone of enterprise administration is **centralized access control**. ChatGPT Enterprise provides *single sign-on (SSO)* and *user provisioning* to integrate with corporate identity systems. Administrators must verify their organization’s domain (via a DNS TXT record) before enabling SSO (help.openai.com). Once verified, the Global Admin Console (or local Workspace “Identity & access” settings) allows configuration of SAML/OIDC connections to an IdP (e.g. Okta, Azure AD, Google Workspace) (help.openai.com). This means users log in with corporate credentials, and administrators can force SSO enforcement or make it optional. OpenAI’s docs warn that “users will be locked out if SSO is not set up correctly” and recommend testing carefully during rollout (help.openai.com).

Once SSO is established, **SCIM provisioning** can be enabled to automate user and group management (help.openai.com). ChatGPT’s SCIM integration (via WorkOS) supports leading providers (Okta, Entra ID/Azure AD, Google Workspace, OneLogin, etc.) (help.openai.com). With SCIM turned on, assigning or removing a user in the company directory automatically invites or deactivates the account in ChatGPT. Administrators can sync directory groups as ChatGPT “Projects” or to assign RBAC roles (see below) (academy.openai.com) (help.openai.com). Slack/governance tools: Company’s SCIM sync typically runs every ~30–40 minutes; and owners can distinguish SCIM-provisioned accounts from manually added users in the workspace settings (help.openai.com). Importantly, SCIM is only available on high-tier (Enterprise/Custom) plans (help.openai.com), reinforcing its enterprise focus.

In addition, Enterprise workspaces can enforce **IP allowlisting**. This lets administrators restrict ChatGPT access to corporate networks or VPN IP ranges. The release notes state that “Enterprise and Edu workspaces can now enable IP allowlisting to control which IP addresses can access ChatGPT and the Compliance API” (help.openai.com). When turned on, only logins from approved networks are permitted. (IP allowlisting is always enforced on Compliance API access and is specific to ChatGPT web access; it does not cover the OpenAI API endpoints (help.openai.com).) This control is particularly valuable for highly regulated industries to ensure that only known internal networks can reach the ChatGPT service.

Finally, identity controls include the ability to create organizational **groups** and **custom roles**. As described below, these interoperate with SSO/SCIM. For example, synced directory groups can populate ChatGPT user groups, and those groups can be assigned roles or projects in ChatGPT. Administrators can even require Multi-Factor Authentication (MFA) via their IdP for ChatGPT logins, although OpenAI’s console does not manage MFA itself (the IdP does that). In sum, ChatGPT Enterprise integrates tightly with existing corporate identity providers to bring ChatGPT logins under enterprise IT policies (help.openai.com).

Workspace Administration and RBAC

Within ChatGPT Enterprise, users belong to a **Workspace** – an organization-wide environment managed by Owners and Admins. The Workspace Settings interface (accessed via chatgpt.com/admin) gives fine control over how ChatGPT is used. As OpenAI documentation explains, from the Settings area “workspace owners and admins can manage members, seat defaults, feature access, GPT governance, apps, analytics, and identity controls” (help.openai.com). Key components include:

- **Workspace Profile:** Owners can update the workspace’s name, display image, and domain identifier. A policy modal (editable under “General”) can be configured to present organization-specific AI usage guidelines to all users every 30 days (help.openai.com). This helps enforce internal AI policies or compliance reminders.
- **Members and Groups:** Owners/administrators can invite or remove users, assign them *seat types* (e.g. ChatGPT seat vs Codex seat), and view pending invites. If SCIM is on, most provisioning is automatic into a default seat type (configurable by the admin); the docs note that admins should set a default seat type *before* enabling SCIM, since “SCIM-provisioned users follow the workspace’s default seat type” (help.openai.com). Users can be organized into **groups** (either created in ChatGPT or synced from the directory). Groups simplify assigning permissions to teams – for example a “Marketing” group might have access to certain GPTs or apps, while “Developers” have others.
- **Permissions & Roles:** Under the “Permissions & roles” tab, workspace Owners set feature access controls. There is a workspace-wide “default role” that applies to all users without a custom role. More importantly, ***custom roles*** can be defined and assigned to one or more groups. The RBAC documentation (updated Jan 2026) explains: *“Workspace Owners will gain RBAC functionality...where they can create custom roles with granular permissions”* (help.openai.com). Workspace Owners can create roles (e.g. “Engineer”), toggle which ChatGPT features those roles allow, and then assign roles to groups (help.openai.com). Permissible toggles include: enabling use of the ChatGPT agent mode, the Canvas coding tool, Codex (code assistant) and its internet access, creating/reusing **Custom GPTs**,

sharing knowledge via Memory, and even data/actions like web search or sending emails via apps (help.openai.com). In effect, RBAC lets the organization safely enable sophisticated features for vetted teams while disabling them for others.

For example, a workspace Owner could create a role permitting only *read-only* actions in certain connected apps, or restrict GPT creation to a "Dev" group. The RBAC FAQ confirms this: *"Permissions include: Canvas (code execution and network access), ChatGPT agent, Codex use, Apps, GPTs (creation, editing, sharing), Project creation, Record (memory notes), Search (web/agent mode), [and] Shared projects"* (help.openai.com). Role inheritance is additive, so a user in multiple groups has the union of those permissions. By default, thread/actions access can also be controlled: admins "can control access to apps on a per-app basis and for select apps, on a per-action basis" (help.openai.com). Workspace Admins may view or update existing roles where supported, but they cannot create, delete, assign, or unassign custom roles (current RBAC permissions).

- **Seat Defaults and Spend Controls:** For this April 2026 comparison, distinguish standard ChatGPT seats from Codex-only seats. OpenAI's April 2 announcement offered Codex-only seats to both Business and Enterprise without a fixed seat fee; their usage was billed by token consumption. Standard Business seats included Codex subject to usage limits. Check the applicable seat terms and usage charges before setting a budget (April 2 announcement).

In summary, the Workspace settings UI exposes nearly all administrative controls in a business-friendly interface. The end-to-end flow is: corporate IT sets up identity (SSO/SCIM, domain), then workspace Owners invite or import users under that domain, define groups and roles, and then toggle feature flags to enforce company policy. The flexibility of RBAC and Groups ensures large enterprises can scale ChatGPT usage securely. (For an example workflow, see the Q&A in OpenAI's RBAC doc (help.openai.com) and enterprise SSO guide (help.openai.com).)

Data Security, Privacy, and Compliance Controls

A crucial appeal of ChatGPT Enterprise is its **data governance and security assurances**. OpenAI's enterprise privacy commitments emphasize that **customer data** (inputs like prompts and outputs like generated text) remains under customer control. Key guarantees include "We do not train our models on your data by default" and "you own your inputs and outputs" (openai.com). All content is encrypted in transit (TLS 1.2+) and at rest (AES-256) (openai.com) (www.techradar.com). Conversations and files can be stored in customer-chosen regions to meet jurisdictional requirements (www.techradar.com). In practice, this means that neither prompts nor responses from ChatGPT Enterprise enter OpenAI's public training pipeline, and admins can adjust how long logs are kept.

Enterprise Owners can set a custom chat-retention policy with a minimum of 90 days (OpenAI Academy). Chat and file lifecycles differ, and deletion schedules can include security or legal exceptions; review the applicable policy for each content type (chat and file retention). Compliance-log retention is separate: the Compliance

Logs Platform retains data for 30 days, and customers needing longer retention should continuously export and retain the logs under their own policies (Compliance Platform).

Enterprise Key Management (EKM) gives eligible Enterprise/Edu workspaces control over an external key management system; implementation requires a named OpenAI account representative. Supported providers are AWS KMS, Google Cloud, and Azure Key Vault (EKM overview). OpenAI's authorized services request that the customer's KMS unwrap a data encryption key, then use that key to decrypt customer content. The customer's master key remains outside OpenAI; this arrangement permits authorized service-side decryption (EKM technical FAQ).

From a compliance standpoint, ChatGPT Enterprise also provides **detailed audit logs**. All admin actions (workspace changes, role edits, SSO logins) can be traced. In mid-2025 OpenAI integrated ChatGPT logs with its broader "OpenAI Compliance Logs Platform" for enterprises (help.openai.com). Released Dec 2025, this consolidation means workspace modifications, authentication activities, and even Codex usage become queryable audit data (help.openai.com). These logs can be fed into SIEM and eDiscovery tools. For example, Microsoft Purview (for customers on Azure AD/Entra) captures ChatGPT interactions in the unified audit log (learn.microsoft.com). Purview's DSPM (Data Security Posture Management) can then analyze ChatGPT prompts and responses alongside other enterprise data flows (learn.microsoft.com). This cross-tool integration lets legal/compliance teams search for AI usage (even including AI-generated emails stored in Exchange) and apply governance policies (retention or deletion) on AI-generated content (learn.microsoft.com).

Finally, admins can disable or restrict advanced features that pose risk. For instance, the **Web Browsing (Agent) mode** is powerful but can pull in external content. ChatGPT Enterprise/Edu lets organizations request domain blocklists for the agent; an admin can contact OpenAI to block browsing or actions to specific sites or subdomains (help.openai.com). This prevents the agent from accessing classified intranets or competitor websites. Similarly, admins can review and disable insecure plugins or app integrations (e.g. apps that send data to third parties). In release notes, OpenAI explicitly notes "Available to [Enterprise] admins: review app actions, set RBAC, enable/disable new actions or connectors" (for Slack, SharePoint, etc.) (help.openai.com). Treat these settings as inputs to a compliance review, and assess the deployment's legal, contractual, and operational requirements separately.

Integrations, Customization, and Tools

Beyond basic chat, Enterprise editions provide extensive integration options and toolkits:

- **Premium Models and Tools**: Advanced data analysis, previously called Code Interpreter, is distinct from Codex (Enterprise launch announcement). Codex supports agentic coding; OpenAI introduced GPT-5-Codex on September 15, 2025 (Codex announcement). Check the seat and billing terms discussed under "Seat Defaults and Spend Controls" before estimating included usage or additional charges.
- **Custom GPTs and Plugins**: Custom GPTs ("GPT-powered apps") allow teams to build specialized assistants with their own instructions and behavior. In Enterprise, these are first-class objects: admins can see all workspace GPTs (with metadata) and must approve any external actions they trigger. The

RBAC system in Workspace Settings governs who can create, share or delete these GPTs (help.openai.com). Plugins (integrations packaged as GPTs) from the ChatGPT plugin store are available but can be restricted. Because plugins may call external APIs, enterprises treat them as potential data exfiltration pathways – policy must govern which plugins are trusted. Admins can disable plugins entirely at the workspace level, or selectively allow only IT-vetted plugins.

- **First-Party "Apps" (Connectors)**: ChatGPT Enterprise comes with a library of productivity app connectors to fetch or act on data from business systems. Examples include Google Drive/Sheets/Docs, Microsoft Outlook/Teams/SharePoint, Slack, Zoom, Atlassian, Salesforce, Shopify, Zendesk, and many others. Administrators control these connectors in the "Apps" tab of workspace settings. By default (unlike Business, where apps are on), Enterprise starts with connectors **disabled**; admins selectively enable needed ones (help.openai.com). For each enabled connector, admins review and approve the specific **actions** it can take. OpenAI's release notes describe a refreshed UI where admins can "enable actions, including sync where supported" for each connector (help.openai.com). For example, in late 2025 admins could deploy a new **Atlassian Rovo** connector, which brings Jira/Confluence context into ChatGPT (and even allows creating Jira issues via chat) (help.openai.com). Google's connector was unified under "Google Drive" so that Docs/Sheets/Slides are accessed via one interface (help.openai.com). In all cases, admins decide who can link and use each app (via roles) and can turn off any undesired write actions.
- **Company Knowledge and MCP**: A notable feature for enterprises is **Company Knowledge**. This allows loading organizational documents (via approved connectors) into ChatGPT as background knowledge. In November 2025 OpenAI updated Company Knowledge to support **custom MCP connectors**: organizations can bring in internal data sources (e.g. proprietary databases) via self-built connectors (help.openai.com). Once enabled, allowable content from these sources is indexed and cited in ChatGPT answers. Administrators govern which connectors' data is available and to whom. There is also a developer mode where an organization's data scientists can build and test new connectors using OpenAI's Model Context Protocol (MCP) and then have them reviewed and published to their workspace. In short, enterprises gain a searchable knowledge repository integrated into ChatGPT, expanding its utility for company-specific Q&A.
- **Canvas & Projects**: New collaboration tools in ChatGPT, like the Canvas coding environment and shared Projects, are also under policy. Admins can disable the Canvas (which executes code) for certain roles under RBAC (help.openai.com). Shared Projects (file-and-instruction sets sharable among users) were introduced in 2025; in Enterprise these can be shared across teams, but admins could limit sharing settings.
- **Memory and Fine-Tuning**: Enterprise customers can use ChatGPT's "Memory" to store organization-specific facts and preferences. Admins have the option to disable memory altogether for privacy, or moderate what types of memory can be written. Similarly, although not widely offered as of April 2026, any future fine-tuning of models would be restricted: custom models created via the OpenAI API are private to the customer (openai.com).

Table 2 (below) summarizes the main **administration panes and settings** in a ChatGPT Enterprise workspace, and the scope of what each controls. It highlights how virtually every aspect of the AI assistant's behavior and integrations can be adjusted by administrators.

Workspace Setting	Function	Administration Scope	Documentation
General	Workspace name/image; AI policy modal	Set org ID, logo, and configure a modal shown every 30 days with usage policy/terms (help.openai.com)	OpenAI Help (Workspace Settings) (help.openai.com)
Members & Groups	Manage user invitations, removals, group membership	Add/remove users or admins; assign seats (ChatGPT vs Codex); create groups (manual or via SCIM) (help.openai.com)	OpenAI Help (Members)
Permissions & Roles	Enable/disable core features; set RBAC	Toggle access to ChatGPT agent, Canvas, Codex, GPT creations, memory, tool use; create custom roles for groups (help.openai.com)	OpenAI RBAC Guide (help.openai.com)
Billing & Seats	View invoices and seat usage; purchase add-ons	Add/change ChatGPT or Codex seat types; set pass/fail rate limits or budgets (spend controls) (help.openai.com)	Release Notes (Codex seats) (help.openai.com)
GPTs	Manage custom GPT assistants	Review/transfer/delete GPTs; view system prompt and actions for internal/custom GPTs (help.openai.com)	OpenAI Help (GPT management)
Apps & Connectors	Enable/disable integrations; configure actions	Select which apps/connectors (e.g. Slack, Google Drive) are available; assign to roles; enable sync; create custom MCP apps (help.openai.com)	OpenAI Help (Apps management)
Workspace Analytics	Usage metrics and trends	View daily/weekly active users, messages/user, seat count, GPT usage, etc. (help.openai.com)	OpenAI Help (Workspace Analytics) (help.openai.com)
Identity & Access	Authentication and provisioning settings	Configure SSO (link to IdP), domain verification, SCIM provisioning, IP allowlists (help.openai.com)	OpenAI Help (SSO/SCIM)
Compliance & Logs	Data export and auditing options	Enable Compliance API; download ChatGPT audit logs; set retention time; integrate with SIEM tools (help.openai.com)	OpenAI Help (Compliance API) (help.openai.com)

Table 2: Key ChatGPT Enterprise workspace settings relevant to administrative control. Each pane in Workspace Settings (as highlighted in the OpenAI help docs) corresponds to strategic control over features, user access, data retention, and integrations (help.openai.com).

Case Studies and Real-World Examples

While OpenAI does not publicly detail enterprise deployments, several customer testimonials and reports illustrate ChatGPT Enterprise in action. OpenAI's launch announcement quotes executives from major adopters: **Klarna**'s CEO praised ChatGPT Enterprise as empowering employees to serve 150 million

customers, while **Asana** reported it "cut down research time by an average of an hour per day" for its data teams (openai.com). Other early enterprise users include Block (Square), Canva, Carlyle Group, Estée Lauder, PwC, and Zapier (openai.com), indicating use cases across finance, design, consulting, and tech.

Independent analyses suggest broad savings and ROI. For example, a recent survey (HumAI Blog, Dec 2025) claims Fortune 500 savings of *40+ minutes daily per employee* and "millions in ROI" via productivity gains across drafting and brainstorming tasks (www.humai.blog).

In practical terms, companies often start by restricting the model's scope. A common pilot might enable ChatGPT for non-PII tasks under a monitored channel, while IT departments analyze logs for leaks or compliance issues. Security firms recommend that enterprises *explicitly define which data must not be entered into ChatGPT prompts* – for example, forbidding sharing of customer data or proprietary code in prompts (beyondscale.tech).

On the analytics side, administrators report using the Workspace Analytics dashboard to monitor adoption. Typical metrics under watch include monthly active users, peak usage times, and which features (like file uploads or code assistant) are most used. These insights can guide internal training or detect bottlenecks in AI usage.

Overall, available anecdotes suggest that when governed properly, ChatGPT Enterprise can accelerate workflows (for research, code review, content revision) without sacrificing security. As one tech buyer magazine noted: ChatGPT has become "a catch-all copilot for email, knowledge management, customer support, and development" in enterprises (www.techradar.com), with OpenAI's enterprise tools moving beyond proof-of-concept to mission-critical applications.

Security, Governance, and Risk Perspectives

While OpenAI adds many controls, enterprise security teams emphasize that **governance cannot rely solely on technical safeguards**. A detailed security guide by the firm BeyondScale (April 2026) argues that ChatGPT Enterprise's infrastructure security is strong (SOC 2, encryption, SSO), but the real attack surface is "the prompt channel" – i.e. what users type (beyondscale.tech).

To mitigate, BeyondScale recommends enterprises implement clear **AI usage policies** and education. They list action items such as: "Is there a written acceptable use policy that explicitly defines what categories of data cannot be entered into ChatGPT prompts?" and "Is any AI-aware DLP deployed to monitor or block sensitive data entering AI interfaces?" (beyondscale.tech). In practice, this has led some companies to route ChatGPT web traffic through supervised proxy or DLP systems that scan for corporate data. Modern DLP solutions (firewalls or endpoint agents) are starting to add "ChatGPT-aware" modules to catch sensitive information in prompts or outputs.

Another risk vector is **API key sprawl**: employees may create personal OpenAI API keys (outside the enterprise subscription) and build stealth integrations, bypassing corporate controls (beyondscale.tech). The BeyondScale report points out that this undermines governance – such usage would not appear in the

ChatGPT Enterprise logs. The cure, they say, is to inventory all OpenAI usage across the company and maybe restrict external API calls. Secret-cracking research found AI-assisted code commits leak secrets at higher rates, underscoring the need to manage keys centrally (beyondscale.tech).

Review custom assistants and connected tools before enabling them, and define which external services users may send organizational data to.

In summary, from a governance standpoint, ChatGPT Enterprise equips IT teams with many tools, but **policy and oversight remain crucial**. Most experts agree that technical controls must be complemented by organizational measures: explicit user training, classification of sensitive categories, regular auditing of logs (via the Compliance API), and eventual incident response plans for AI misuse. The combination of OpenAI's built-in controls (SSO, RBAC, encryption, logs) and a disciplined security posture is what enables safe scaling of generative AI in business.

Implications and Future Directions

ChatGPT Enterprise represents a pivotal shift in how organizations adopt AI. The breadth of administrative controls and integrations shows that enterprises want both power and accountability from AI. **Looking forward**, several trends emerge:

- **Unified Admin Management:** In 2026 OpenAI introduced the **Global Admin Console** (at admin.openai.com) to address large organizations' needs for a tenant-level control plane. A Tenant can include multiple ChatGPT workspaces and API organizations, all under one SSO and verified domains (help.openai.com). Global Admins can manage domains and SSO for all sub-accounts, and add other Global Admin users (help.openai.com). This development signals that OpenAI is moving toward enterprise IT norms (e.g. Azure's multi-tenant management) and will likely expand central policies. Future implications: multi-workspace analytics, corporate-wide security dashboards, and delegating privileges across business units.
- **Expanded Compliance Features:** OpenAI has steadily added features: data residency (10 regions as of late 2025) (www.techradar.com), compliance log integrations (e.g. connectors with CrowdStrike, GlobalRelay) (help.openai.com), and EKM (help.openai.com). We expect more: for instance, after decoupling the stateful audit route, they may add more integrations (e.g. Splunk, and real-time alerting), as well as advanced DLP integrations. The move to align with tools like Microsoft Purview (learn.microsoft.com) suggests future tighter coupling with enterprise compliance platforms.
- **Enhanced AI Collaboration:** The integration of advanced AI agents, code tools, and knowledge graphs will continue. OpenAI's mention of building an "AI superapp" combining ChatGPT, Codex, browsing, and other agentic AI (www.techradar.com)points to a future where ChatGPT Enterprise might unify all work-related AI tasks. For example, one might imagine a unified interface where a single agent can read emails, query databases, generate documents, and write code, all within corporate controls. This will necessitate even finer-grained controls (e.g. per-AI-agent policies) and raises questions about human-AI teaming and oversight.

- **Market and Competition:** The enterprise AI market remains dynamic. In a16z's January 30, 2026 survey of 100 Global 2000 companies, 78% of surveyed enterprise CIOs reported using OpenAI models in production, directly or through cloud providers. This measures adoption within that sample, and respondents could use multiple vendors; it is not OpenAI's share of all enterprise AI usage (original survey). Enterprises should evaluate data flows across the LLM providers they use. The breadth of ChatGPT Enterprise's tooling sets a high bar, so competing products will need comparable admin ecosystems.

In conclusion, ChatGPT Enterprise offers a rich platform for organizations to leverage generative AI while maintaining control. The combination of official documentation and real-world analysis shows a continuously evolving system. Today's capabilities – from SSO and RBAC to data encryption and compliance APIs – address many of the pitfalls that general-purpose ChatGPT lacks. For IT leaders, staying current with OpenAI's updates (via the help center release notes and admin console) is critical. The trajectory suggests that by 2026 and beyond, managing AI will become a core part of IT governance, with ChatGPT Enterprise as a leading example.

All claims and data above are drawn from OpenAI's official documentation and credible analyses. Key sources include OpenAI's Help Center articles (linked throughout this report) (help.openai.com), OpenAI's product announcements (openai.com) (www.techradar.com), financial/industry reporting (www.techradar.com), and security research (beyondscale.tech). Readers should consult the latest OpenAI documentation for up-to-date details on enterprise features (links provided above) and consider independent expert guidance (e.g. security assessments) when deploying ChatGPT in a corporate environment.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://openai.com/business-data
2. https://openai.com/index/introducing-chatgpt-enterprise
3. https://help.openai.com/en/articles/11750701-managing-feature-access-with-role-based-access-control-in-chatgpt
4. https://help.openai.com/en/articles/8411955-managing-workspace-settings-in-chatgpt-enterprise
5. https://help.openai.com/en/articles/12289294-global-admin-console
6. https://www.techradar.com/pro/we-are-growing-revenue-four-times-faster-than-the-companies-who-defined-the-internet-and-mobile-eras-openai-says-its-making-usd2-billion-a-month-mostly-from-enterprise-users
7. https://beyondscale.tech/blog/chatgpt-enterprise-security-assessment
8. https://www.humai.blog/how-top-companies-use-chatgpt-enterprise-real-case-studies-results-and-what-you-can-learn-from-their-success
9. https://openai.com/policies/api-data-usage-policies?%25253Butm_content=17950474&%25253Butm_medium=email&%25253Butm_source=hs_email
10. https://help.openai.com/en/articles/10128477-chatgpt-enterprise-edu-release-notes%23.webm
11. https://www.techradar.com/pro/openai-now-lets-business-customers-choose-where-their-chatgpt-data-is-hosted
12. https://learn.microsoft.com/en-us/purview/ai-chatgpt-enterprise
13. https://help.openai.com/articles/11750701

14. https://help.openai.com/en/articles/10472980-enabling-sso-on-chatgpt
15. https://help.openai.com/en/articles/10011769
16. https://help.openai.com/en/articles/9903489-data-residency-and-inference-residency-for-chatgpt
17. https://help.openai.com/en/articles/20000943-openai-enterprise-key-management-ekm-overview
18. https://help.openai.com/en/articles/8798634-managing-data-sharing-and-privacy-in-chatgpt-business
19. https://help.openai.com/es-es/articles/9261474-openai-compliance-platform-for-enterprise-customers
20. https://help.openai.com/en/articles/11165333-chatgpt-enterprise-and-edu-models-and-limits
21. https://help.openai.com/en/articles/9275245-chatgpt-free-tier-faq
22. https://help.openai.com/en/articles/11509118-admin-controls-security-and-compliance-in-connectors-enterprise-edu-and-team?hsPreviewerApp=page
23. https://help.openai.com/en/articles/11391654-chatgpt-business-release-notes
24. https://academy.openai.com/public/clubs/admins-6o6xf/resources/scim
25. https://openai.com/index/codex-flexible-pricing-for-teams
26. https://academy.openai.com/en/public/clubs/admins-6o6xf/resources/data-governance-and-compliance
27. https://help.openai.com/en/articles/8983778-chat-and-file-retention-in-chatgpt
28. https://help.openai.com/en/articles/9261474-openai-compliance-platform-for-enterprise-and-edu-customers
29. https://help.openai.com/en/articles/20000945-ekm-technical-faq
30. https://openai.com/index/introducing-upgrades-to-codex
31. https://help.openai.com/en/articles/11391654-chatgpt-team-release-notes
32. https://www.techradar.com/pro/openai-models-are-the-favorites-for-enterprise-users-right-now-but-anthropic-isnt-far-behind
33. https://a16z.com/leaders-gainers-and-unexpected-winners-in-the-enterprise-ai-arms-race

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.