




Securing ChatGPT Teams & Enterprise

Risk Management for Alternative Investment Firms



By Richard Itri, Chief Innovation Officer, ECI



Executive Summary

In my 28+ years in the alternative investments industry, no technology has captured management and C-level attention like generative AI (GenAI). As a result, firms are rapidly adopting AI, often without the usual due diligence or risk analysis for products involving sensitive data. Currently, the two most widely used OpenAI SaaS platforms are ChatGPT Enterprise and ChatGPT Teams. Let's take a closer look.

The adoption of generative AI – in particular, ChatGPT Enterprise – is accelerating rapidly across financial services. Hedge funds, private equity and investment banking firms are eager to leverage it for research automation, legal and compliance, reporting and analytics, as well as for improved investor, client and LP engagement.

However, with this transformative potential comes a unique and complex risk landscape. For alternative investment firms that routinely handle proprietary investment strategies, investor data, deal flow, and compensation information, these risks are amplified by regulatory expectations and reputational stakes.

“



The adoption of generative AI is accelerating rapidly across financial services. Hedge funds, private equity and investment banking firms are eager to leverage AI for improved investor and client engagement. However, with this transformative potential comes a unique and complex risk landscape. For alternative investment firms these risks are amplified by regulatory expectations and reputational stakes.

Richard Itri, Chief Innovation Officer, ECI

”



Security, Compliance and Data Governance

“

Residual chat histories and shared prompts pose unseen risks unless actively purged.

”

The term “Enterprise” may misleadingly suggest the ChatGPT platform is secure and compliant by default, without much setup or maintenance. While OpenAI has equipped the tool with solid security and compliance features, organizations are still responsible for implementation, integration and ongoing management.

In this paper, my goal is to highlight some of the key security, compliance and data governance risks of ChatGPT Teams and ChatGPT Enterprise adoption and provide a pragmatic roadmap for mitigation – emphasizing how a managed service approach can simplify, accelerate and safeguard your AI journey.

Note: Although this paper focuses specifically on ChatGPT Enterprise, many of the risks and mitigations can apply to any SaaS or 3rd party platform where confidential and regulatory compliant data is being stored or processed.

Bottom line, know your data and AI estate and implement best practices, as outlined in this document.






ChatGPT Teams vs Enterprise: Key Differences to Understanding Risks

Currently, the two widely used OpenAI GenAI platforms are ChatGPT Teams and Enterprise. Both offer the power of OpenAI models, grounding and research capabilities. However, they differ on price, security, compliance and agentic workflows.

Key Difference	Teams	Enterprise
SSO	✗	✓
SCIM Provisioning	✗	✓
Compliance APIs	✗	✓
Advanced Admin	Basic	Full/Granular
User Capacity	2–149	150+ (scalable)
SLA	✗	✓
Enterprise Support	✗	✓
Regulatory Readiness	Basic	Advanced (SOC 2, DPA, eDiscovery)



As you can see, although the ChatGPT Teams platform costs less and supports smaller user counts, the lack of security and compliance features makes it a difficult choice for regulated organizations.



Comprehensive Risk Landscape

Alternative investment firms handle highly sensitive, regulated data, including trading positions, investor identities, deal memos and compensation records – all of which carry regulatory, financial and reputational risks. The digitization of these workflows, and now the integration of generative AI, exposes this data to new vectors of risk, requiring a security posture that is both proactive and adaptive.

“

The term ‘Enterprise’ doesn’t make ChatGPT secure by default – implementation is everything.

”

Regulatory frameworks such as the SEC, FINRA, GDPR, and CCPA only heighten the obligation for robust controls, data residency, and demonstrable governance.

Key Risks to Consider

Identity and Access Management

- **Unauthorized Access:** Without robust IAM, SSO, and MFA controls, sensitive data can be exposed to unauthorized users – either internally or through credential compromise.
- **API and Endpoint Exposure:** Insecure API integrations or unprotected endpoints introduce risk of data exfiltration, lateral movement or malicious access.
- **Credential Leakage:** Phishing or weak credential hygiene can result in loss of access control, placing proprietary information at risk. Given the popularity of the platform, threat actors are actively phishing for ChatGPT credentials.

Compliance Risks

- **Regulatory Non-Compliance:** Failure to archive, retain and audit AI interactions can result in violations of SEC, FINRA and global data protection regulations.
- **Inadequate Audit Trails:** Incomplete logging or gaps in data lineage undermine a firm’s ability to respond to regulatory inquiries or legal holds.
- **Cross-Border Data Residency:** AI workflows can inadvertently move or process data across jurisdictions, risking non-compliance with GDPR and other data sovereignty requirements.

“

*For regulated firms,
AI risk isn't theoretical
– it's reputational,
regulatory and real.*

”

Data Privacy Risks

- **Sensitive Data Leakage:** Employees may input confidential LP lists, compensation figures, or strategic plans into ChatGPT, risking accidental exposure through sharing features or internal misuse.
- **Insider Threats:** Malicious or careless insiders may use AI to aggregate, analyze, or transmit sensitive data outside approved channels.
- **Consent Management Gaps:** Lack of centralized oversight on how data is used and shared through AI interfaces exposes the firm to contractual and regulatory breaches.

Data Security Risks

- **Encryption and Key Management:** OpenAI does not offer BYOK, so it controls the encryption keys. As a result, internal users or threat actors could potentially access those keys and your data.
- **Data Segregation:** Inadequate logical isolation between tenants or groups can result in cross-team or cross-client data exposure.
- **Third-Party Add-ons:** Unvetted plugins or integrations can introduce vulnerabilities or additional data sharing pathways.

Model Learning and Vendor Access Risks

- **Model Learning Concerns:** While OpenAI asserts that enterprise data is not used for model training, ambiguity or policy changes could create future risks.
- **Admin access:** OpenAI retains admin access, and not all the data is stored in an encrypted format, meaning OpenAI can technically access unencrypted information or logs, leaving some residual risk.

Mitigation Strategies and Controls

A successful AI adoption strategy must align each risk with strong, auditable mitigations – ideally through a unified managed service. Here's how leading firms are tackling the challenge:

Identity and Access Management

- **Authorized Access Management:** Use SSO, MFA, and role-based access with platforms like Microsoft Entra ID or Okta. Restrict AI feature access to authorized users as needed and ensure SSO only permits compliant endpoints. Limit access to corporate devices to reduce data leakage risks.
- **Endpoint and API Protection:** Deploy advanced endpoint detection (e.g. SentinelOne, CrowdStrike, Defender), and secure all API endpoints with monitoring, rate-limiting and alerts.
- **Credential Hygiene:** Automate credential management, enforce passwordless or strong authentication, and deliver continuous security awareness training.

“

AI is moving faster than most IT and compliance teams can govern. And that's the problem.

”

“

It's not about saying no to AI – it's about saying yes, safely.

”



“

What you put into ChatGPT could end up somewhere you never intended – unless controls are in place

”

Compliance

- **Retention and Archiving:** Integrate AI interactions into compliance archives (Smash, Global Relay, Microsoft Purview) with retention policies mapped to SEC Rule 17a-4 and FINRA Rule 4511.
- **Comprehensive Audit Logging:** Leverage OpenAI's Compliance API and third-party SIEMs (ECI XDR, Sentinel, Splunk) to create immutable logs, supporting e-discovery and regulatory response.
- **Data Residency Enforcement:** Use region-specific deployments (Azure OpenAI, AWS) and configure geofencing and data residency controls to keep data in approved jurisdictions.

Data Privacy

- **DLP and Prompt Filtering:** Deploy data loss prevention and prompt filtering tools (Microsoft Purview DLP, Cyera) to automatically flag or block sensitive keywords and patterns before data is submitted to ChatGPT.
- **Disable Sharing and Add-on Features:** Turn off chat sharing, third-party add-ons, and browsing capabilities via admin controls for high-risk user groups (e.g. HR, executive compensation). This is an important line item, and one often missed.
- **Group-Based Segmentation:** Use RBAC to ensure sensitive data is only accessible to specific teams and functions, with group-specific GPTs and connectors.



“
Global regulatory
frameworks
demand more
than encryption,
they require
demonstrable
governance.”

Data Security

- **Encryption and BYOK:** While native ChatGPT Enterprise doesn't currently support Bring Your Own Key (BYOK), consider hybrid models leveraging Azure OpenAI or ECI's ELLA platform with customer-managed keys for highly sensitive workflows.
- **Tenant Isolation:** Deploy logical and technical controls to guarantee strict data segregation between departments and client entities.
- **Security audits:** Conduct regular security audits through scheduled penetration testing and independent third-party evaluations of AI environments and integrations. Utilize Purview or an equivalent platform to forward audit reports to compliance teams, ensuring robust access management and data control measures are in place.

Model Learning and Vendor Access

- **Contractual Protections:** Ensure enterprise agreements explicitly prohibit use of firm data for model training, mandate audit rights, and clarify penalties for unauthorized access.
- **Continuous Vendor Monitoring:** Regularly review OpenAI's certifications (SOC 2, ISO 27001), request audit logs, and perform compliance checks as part of ongoing vendor governance.
- **Transparency and Notification:** Monitor for changes in vendor policy and ensure your organization is promptly alerted to any updates affecting data privacy or access.





Governance and Operationalizing Controls

Managing these risks at scale requires more than just technical controls. It demands a governance framework that brings together policy, process and people.

- **Centralized Admin Management:** Assign dedicated admin roles for AI governance, with authority to manage features, monitor usage, and respond to incidents.
- **Continuous Management Evolvement:** Maintain an AI risk register, schedule regular reviews, and adapt controls as threats and regulatory expectations evolve.
- **Employee Training and Accountability:** Provide ongoing training on acceptable AI use and data handling, and tie compliance to performance and disciplinary policies.
- **Incident Response and Investigation:** Establish clear protocols for investigating and remediating potential data incidents, leveraging integrated logging and forensic tools.
- **Reporting:** Circulate usage and access reports monthly to AI committee members and stakeholders. Ensure ownership over platform use and leverage to improve training and adoption.





Offboarding Risks

Offboarding is a critical, often overlooked risk and warrants its own section in the risks and mitigation factors. Keep in mind, many of these are manual tasks and should be added to your offboarding policies, procedures and controls.

Residual Data Exposure

- **Chat Histories:** Unless proactively purged, the departing user's chat logs, prompts and shared conversations may remain accessible within ChatGPT Enterprise.
- **Shared Content:** If the user shared chats or outputs with other users or groups, that data could persist even after account deprovisioning.
- **Sensitive Attachments:** Uploaded files, code snippets, or sensitive outputs may remain in storage, potentially accessible by others with sufficient permissions.

Incomplete Account Deactivation

- **Delayed Offboarding:** If IT or HR lags in deactivating accounts, departing users may retain access for some period.
- **Lack of Automated Purge:** By default, deleting a user doesn't always guarantee their data (especially shared content) is immediately deleted from the platform.

Regulatory and Legal Risks

- **Data Retention Conflicts:** Regulatory requirements may mandate retention for a period, but also require that only authorized personnel have access.
- **E-Discovery and Audit:** Unmanaged data from departed users could surface in audits or legal discovery, complicating compliance.

“

*Simply deactivating
a user is not enough.*

”

What Happens to User Data in ChatGPT Enterprise?

- **User Deactivation:** Disabling or deleting a user account via SSO or admin console typically removes login access.
- **Data Persistence:** However, chat histories, shared GPTs, and uploaded content may remain in the system unless specifically deleted.
- **Admin Controls:** Some platforms allow admins to reassign, archive, or purge user content – but this may require manual intervention or use of Compliance APIs.

Note: OpenAI's own documentation indicates that user data is not automatically purged upon account deletion. Chat content may be retained according to organizational and regulatory policy.

“

Failing to archive AI interactions could land you in breach of SEC or GDPR rules ... even unintentionally.

”





“
ChatGPT Enterprise
needs more than
a license – it
needs a layered
risk management
strategy.”

Best Practices to Ensure Data is Purged

Strong Offboarding Process

- **Immediate Deactivation:** Automate user deactivation via SSO/SCIM as part of the HR exit checklist.
- **Audit Access Logs:** Review last access and ensure sessions are terminated.

Manual or Automated Data Purge

- **Purge Chat Histories:** Use ChatGPT Enterprise admin tools or Compliance APIs to delete or export and then purge user data, including chats, custom GPTs, and file uploads.
- **Reassign or Archive Needed Content:** For regulatory retention, archive required data to a secure, access-controlled repository and delete it from ChatGPT's live environment.

Shared Content Review

- **Shared Links/GPTs:** Identify and review any shared chats or custom GPTs the user created. Disable, archive, or transfer ownership as appropriate.

Integrate with DLP and Archiving

- **Automated Scanning:** DLP systems (like Microsoft Purview, Smarsh, or Global Relay) can detect and trigger alerts if sensitive content remains accessible after offboarding.
- **Archival Policy Enforcement:** Retain only what's required; purge everything else, and document the process for auditability.

Policy and Training

- **Clear Policy:** Document the purge process and communicate it to HR, IT, and InfoSec.
- **Regular Review:** Periodically test offboarding to ensure all residual access and data retention risks are closed.

Summary Table

Risk	Mitigation Steps	Tools/Solutions
Residual Chat Data	Purge via admin tools or Compliance API	ChatGPT Admin Console, Compliance API
Shared Content	Review and disable or reassign shared chats/GPTs	Microsoft Purview, Cyera, ChatGPT Enterprise, Admin Review
Account Deactivation	Automate via SSO/SCIM, terminate sessions	Microsoft Entra ID, Okta,
Regulatory Conflicts	Archive as required, purge from production	Smarsh, Global Relay, Microsoft Purview
Auditability	Document and log all offboarding & purge actions	SIEM, Compliance Log, Microsoft Purview, Cyera

Recommendations for Firms

- **Automate offboarding:** Integrate HR, IT and ChatGPT Enterprise admin processes.
- **Use compliance APIs:** Where possible, script data export and deletion as part of the exit process.
- **Monitor and test:** Periodically review offboarding effectiveness and close any gaps.
- **Work with your MSP/MIP:** Engage your managed service provider to automate and enforce these controls, providing assurance and audit support.

Bottom line

Simply deactivating a user is not enough. Firms must implement a robust process for identifying, reviewing and purging all AI-related data for departing users to eliminate exposure and meet regulatory requirements.



Conclusion

ChatGPT Enterprise presents both opportunities and risks for alternative investment firms, making robust risk management essential.

A managed service led by AI and financial regulation experts can provide comprehensive protection covering:

- Configuration
- Monitoring
- Policy enforcement
- Employee training
- Vendor due diligence
- Audits.

With the right guardrails, ChatGPT Enterprise can become a powerful, secure enabler for innovation – provided firms embed governance, compliance and oversight from the outset.

Action Plan

Don't leave AI adoption to chance – secure it from day one. The right controls today mean fewer risks tomorrow.

For more information on implementing and managing ChatGPT Enterprise, contact [ECI: sales@eci.com](mailto:sales@eci.com).



ECI Managed Security for ChatGPT Enterprise

Service Category	Description	Examples / Deliverables
Implementation & Onboarding	• SEC, FINRA, GDPR compliance • Data residency management • Archiving & data retention policies	• Project kickoff & plan • Environment validation
Identity & Access Management	• Single Sign-On (SSO) setup via Azure Entra ID, Okta, Duo • Custom Role-Based Access Control (RBAC) • Custom security groups management- Conditional access & MFA enforcement • User lifecycle management (onboarding/offboarding) • Regular access reviews & audits	• SSO implementation with Azure Entra ID or Okta • Custom RBAC and group configuration • MFA & conditional access policies • Quarterly IAM audit & access review reports
Security Management	• Real-time monitoring of access & usage • Data protection (DLP) • Audit logging • Incident detection and response	• Microsoft Purview DLP policies • Sentinel integration & alerting • Monthly security reporting
Compliance & Governance	• SEC, FINRA, GDPR compliance • Data residency management • Archiving & data retention policies	• Quarterly compliance audit reports • Regulatory alignment documentation
Reporting & Insights	• Usage analytics & adoption reporting • Compliance reporting • Security event reporting	• Power BI dashboards & insights • Scheduled monthly or quarterly reports
Training & Enablement	• User onboarding & adoption support • Advanced scenario workshops • Compliance & security training	• Interactive workshops/webinars • White-labeled Skillable training modules
Integration Services	• Integration with CRM, Document Management, Trading Systems • Workflow automation	• Salesforce, Dynamics 365, DealCloud integrations • Azure Logic Apps custom workflows
Custom Agent Development	• Agents customized for Investment Banking, Hedge, PE workflows • Automating financial modeling, due diligence, RFP response	• Custom-built financial research & summarization agents • Agent integration & deployment documentation
Ongoing Support	• Business hours Helpdesk support • Dedicated Customer Success Manager • Platform optimization • Offboarding	• Incident management SLA adherence reports • Regular strategic optimization reviews
Platform Innovation	• Continuous improvement roadmap • Industry-specific AI capabilities updates	• Quarterly innovation workshops • New workflow & agent releases