ibp
Institute for Biopolymers
at Hof University

KREATIK

Data Security, Privacy, and Compliance Risks:
Comprehensive Review of 12 Enterprise AI Tools

Prepared for: IT Security and Management Decision-Making
Date: June 23, 2026
Classification: Internal Use – Management Review

Bayerisches Staatsministerium für
Wissenschaft und Kunst

Finanziert von der
Europäischen Union

Inhalt

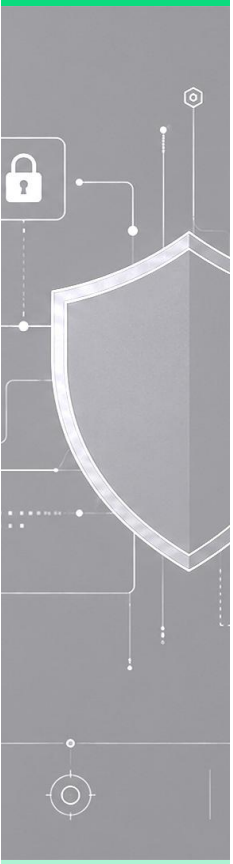

1. Overview and Purpose

This comprehensive review evaluates data security, privacy, and compliance risks associated with 12 enterprise AI tools commonly deployed for meeting transcription, document processing, and workplace productivity. The analysis examines data processing practices, model training policies, GDPR/EU[1] compliance, security certifications, and known incidents to provide evidence-based risk classifications for different organizational contexts.

Key Findings:

Tier 1 (Lowest Risk): Microsoft 365 Copilot, Google Gemini for Workspace, and Anthropic Claude Enterprise demonstrate the strongest compliance posture, offering EU Data Boundary commitments, comprehensive certifications (SOC 2 Type II, ISO 27001, ISO 42001), no model training on customer data by default, and mature enterprise governance frameworks [1], [2], [3].

Tier 2 (Moderate-Low Risk): Zoom AI Companion, Fireflies.ai Enterprise, and Perplexity Enterprise Pro provide solid enterprise security with SOC 2 Type II certification and GDPR compliance, though with varying degrees of EU data residency guarantees [4], [5], [6].

Tier 3 (Moderate Risk): ChatGPT Enterprise, Notion AI Enterprise, Fellow, and Read AI offer enterprise-grade security but face limitations in EU data residency, reliance on third-party AI subprocessors, or limited public documentation of compliance measures [7], [8], [9], [10].

Tier 4 (Elevated Risk): Otter.ai and Fathom present heightened risks due to active litigation (Otter.ai class action lawsuit for unauthorized recording), limited enterprise certifications, and insufficient public documentation of data protection measures [11], [12].

Critical Regulatory Context: OpenAI (ChatGPT) faced a €15 million fine from the Italian Data Protection Authority in January 2025 for GDPR violations, following a temporary ban in 2023 [13]. Zoom reversed controversial terms of service in 2023 that would have permitted AI training on customer content without explicit consent [14]. These incidents underscore the dynamic regulatory landscape and the importance of continuous compliance monitoring.

For EU Organizations: Only Microsoft 365 Copilot and Google Gemini for Workspace currently offer comprehensive EU Data Boundary commitments with in-region processing guarantees. All other tools process data outside the EU or rely on subprocessors not covered by EU data residency commitments [1], [2].

Recommendation: Organizations handling sensitive data should prioritize tools with ISO 42001 (AI management systems) certification, explicit no-training policies, EU data residency options, and mature incident response capabilities. Meeting transcription tools require particular scrutiny due to third-party consent requirements and the risk of capturing sensitive discussions without proper authorization.

[1] https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:02016R0679-20160504

2. Comparative Overview Table

Tool	SOC 2 Type II	ISO 27001	ISO 42001	GDPR Compliant	EU Data Residency	No Model Training (Default)	HIPAA Available	Known Incidents
Microsoft 365 Copilot	✓	✓	✓	✓	✓ Full (EU Data Boundary)	✓	✓	None
Zoom AI Companion		✗ (ISO 27701)	✗	✓	✓ Partial (EU infrastructure available)	✓	✗	2023 ToS controversy (resolved)
Otter.ai	! Limited	✗	✗	! Limited	✗	! Unclear	✗	2025 class action lawsuit (active)
Fellow	✓	✗	✗	✓	! Not documented	! Not clearly documented	! Claimed, not verified	None
Read AI	✓	✗	✗	! Claimed	! Not documented	✓ (opt-in only)	✗	None
Fireflies.ai	✓	✗	✗	✓	! Private storage option (Enterprise)	✓ (0-day retention with vendors)	✓ (Enterprise)	None
ChatGPT Business/Enterprise	✓	✗	✗	✓	✗ (US-based, no guarantee)	✓	✓ (API with BAA)	Italy ban 2023, €15M fine 2025
Google Gemini Workspace	✓	✓	✗	✓	✓ Full (with appropriate settings)	✓	✗	None
Claude Team/Enterprise	✓	✓	✓	✓	! Limited (AWS/GCP worldwide)	✓	✓ (Enterprise)	None
Perplexity Enterprise Pro	✓	✗	✗	✓	! Not documented	✓	✓	None
Notion AI/Enterprise	✓	✓	✗	✓	! Not documented	✓ (Anthropic/ OpenAI as subprocessors)	✗	None
Fathom	✓	✗	✗	✓ (claimed)	! Not documented	! Not clearly documented	✗	None

Legend:

✓ = Verified/Available

✗ = Not Available/Not Certified

! = Limited/Unclear/Requires Verification

3. Detailed Analysis per Tool

3.1 Microsoft 365 Copilot/Teams Premium

Data Processing & Storage:

Microsoft 365 Copilot operates within the Microsoft 365 service boundary, processing prompts, responses, and data accessed through Microsoft Graph without transferring data outside the existing Microsoft 365 infrastructure [1]. Customer data remains within the tenant's geographic region, with EU Data Boundary compliance ensuring that EU traffic stays within the EU for LLM processing [1]. Prompts and responses are stored as part of the user's Copilot activity history, encrypted at rest (AES-256) and in transit (TLS 1.2+), and can be managed through Microsoft Purview for eDiscovery and retention policies [1].

Model Training Policies:

Microsoft explicitly states that "prompts, responses, and data accessed through Microsoft Graph aren't used to train foundation LLMs, including those used by Microsoft 365 Copilot" [1]. Optional customer feedback may be used to improve services but not to train foundation models [1]. Azure OpenAI services used by Copilot do not cache customer content, and abuse monitoring with human review has been opted out for Microsoft 365 Copilot services [1].

GDPR/EU Compliance:

Microsoft 365 Copilot is compliant with GDPR and operates within the EU Data Boundary for European users [1]. Data residency commitments were formalized in the Microsoft Product Terms effective March 1, 2024, covering Advanced Data Residency (ADR) and Multi-Geo capabilities [1]. However, Anthropic models used within some Copilot experiences are explicitly out of scope for the EU Data Boundary and in-country LLM processing commitments, representing a notable exception [1].

Security Certifications:

Microsoft 365 Copilot holds SOC 2 Type II, ISO 27001, ISO 42001 (AI management systems), and HIPAA certifications [1]. The ISO 42001 certification is particularly significant as it demonstrates adherence to international standards for AI management systems, including risk assessment, transparency, and accountability frameworks [1].

Meeting Transcription Risks:

When used in Microsoft Teams, Copilot respects existing permission models, surfacing only data users have view permissions to access [1]. Meeting transcripts and summaries are stored within the Microsoft 365 tenant and subject to the same access controls as other Teams content [1]. The primary risk lies in the Anthropic subprocessor exception, which processes certain requests outside the EU Data Boundary [1].

Integration Risks:

Copilot integrates deeply with Microsoft Graph, accessing emails, chats, documents, calendar, and contacts [1]. While this integration is architecturally contained within the Microsoft 365 boundary, it creates a broad attack surface if tenant permissions are misconfigured [1]. Extensibility through agents and Microsoft Graph connectors introduces third-party data access risks that require admin oversight [1].

Prompt/Document Leakage Risks:

Microsoft implements multiple protections including harmful content blocking, protected material detection, and prompt injection (jailbreak) attack mitigation [1]. Semantic Index honors user identity-based access boundaries, ensuring grounding processes only access content the current user is authorized to view [1]. Sensitivity labels and Information Rights Management (IRM) encryption are respected by Copilot [1].

Enterprise Governance:

Administrators can manage Copilot through the Microsoft 365 admin center, control agent permissions, set retention policies via Microsoft Purview, and use Content search for eDiscovery [1]. Users can delete their Copilot activity history through the My Account portal [1]. Privacy controls for connected experiences in Microsoft 365 Apps can affect Copilot availability [1].

Known Incidents:

No major security incidents or data breaches specific to Microsoft 365 Copilot have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Low	Strong compliance, EU Data Boundary, no training on customer data
R&D	Low	Respects permissions, encrypted storage, mature governance
HR Data	Low-Medium	Suitable with proper permission configuration; Anthropic subprocessor outside EU boundary
Customer Data	Low	GDPR compliant, data residency commitments, encryption at rest/transit
Legal/Compliance Meetings	Low-Medium	Strong security but requires awareness of Anthropic subprocessor limitation
Public Sector/University (EU)	Low-Medium	EU Data Boundary compliant except for Anthropic models; ISO 42001 certified

3.2 Zoom AI Companion

Data Processing & Storage:

Zoom AI Companion processes meeting audio, video, chat, and transcripts to generate summaries, action items, and other AI-powered features [4]. Data is encrypted in transit (TLS) and at rest (AES-256) [4]. Zoom offers EU infrastructure options, allowing organizations to specify data routing preferences for meetings and recordings [4]. However, AI processing may occur in multiple regions depending on capacity and service availability [4].

Model Training Policies:

Following significant controversy in 2023, Zoom clarified that it does not use customer content (audio, video, chat, screen sharing, attachments, or other communications-like customer content) to train AI models without customer consent [14]. The updated terms explicitly state that customers retain ownership of their content and Zoom only uses it to provide services [14]. This policy reversal was implemented after widespread criticism and regulatory scrutiny [14].

GDPR/EU Compliance:

Zoom is GDPR compliant and offers a Data Processing Addendum (DPA) for enterprise customers [4]. The company holds ISO 27701 certification, a privacy-specific extension to ISO 27001 [4]. Zoom has completed Data Transfer Impact Assessments (DTIAs) for Zoom Meetings, Chat, Webinars, Phone, Contact Center, and Virtual Agent, available through the Zoom Trust Portal [4].

Security Certifications:

Zoom holds SOC 2 Type II and ISO 27701 certifications [4]. Notably, Zoom does not hold ISO 27001 or ISO 42001 certifications, which are common among enterprise AI platforms [4]. HIPAA compliance is available through a Business Associate Agreement (BAA) for healthcare customers [4].

Meeting Transcription Risks:

Zoom AI Companion can automatically join meetings to provide transcription and AI features. The primary risk is third-party consent: participants may not be aware that AI processing is occurring, particularly if they join from external organizations [14]. The 2023 terms of service controversy highlighted the importance of explicit consent mechanisms for AI features [14].

Integration Risks:

Zoom integrates with calendar systems, CRM platforms, and productivity tools. AI Companion can access meeting metadata, participant information, and content shared during meetings [4]. Organizations must ensure that integrations do not inadvertently expose sensitive data to AI processing without proper authorization [4].

Prompt/Document Leakage Risks:

Zoom implements content filtering and security controls, but specific details about prompt injection protection and data leakage prevention for AI Companion are not extensively documented in public materials [4]. The platform's reliance on multiple AI vendors and processing regions introduces potential data exposure risks [4].

Enterprise Governance:

Zoom provides admin controls to enable or disable AI Companion features at the account, group, or user level [4]. Administrators can configure consent notifications, data retention policies, and access controls through the Zoom Admin Portal [4]. Audit logs track AI feature usage and data access [4].

Known Incidents:

In August 2023, Zoom faced significant backlash over updated terms of service that appeared to grant the company broad rights to use customer content for AI training [14]. Following public outcry and regulatory concerns, Zoom issued clarifications and updated its terms to explicitly prohibit using customer content for AI training without consent [14]. This incident raised questions about transparency and trust in AI-powered communication platforms [14].

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Low-Medium	GDPR compliant, but 2023 ToS controversy raises trust concerns
R&D	Medium	EU data residency partial; AI processing may occur outside EU
HR Data	Medium	Third-party consent risks in interviews; limited HIPAA scope
Customer Data	Medium	DPA available but multi-region processing introduces complexity
Legal/Compliance Meetings	Medium-High	Third-party recording consent requirements; 2023 incident history
Public Sector/University (EU)	Medium	ISO 27701 certified but lacks ISO 42001; EU infrastructure available but not guaranteed

3.3 Otter.ai

Data Processing & Storage:

Otter.ai processes audio recordings, generates transcripts, and uses speaker identification to tag participants [15]. The platform collects extensive metadata including timestamps, device information, IP addresses, and usage patterns [15]. Audio recordings and transcripts are stored on Otter.ai's infrastructure, with encryption at rest and in transit [15]. However, specific details about data residency, regional storage options, and data retention policies are not comprehensively documented in public materials [15].

Model Training Policies:

Otter.ai's privacy policy states that the company may use "deidentified or aggregated data" for its purposes, including improving speech recognition technology [15]. Critically, the policy does not explicitly prohibit using customer audio recordings or transcripts to train AI models [15]. The August 2025 class action lawsuit alleges that Otter "retains conversational data indefinitely and leverages it to refine its speech recognition technology without participant permission" [11].

GDPR/EU Compliance:

Otter.ai claims GDPR compliance and offers data subject rights mechanisms [15]. However, the company does not provide a publicly available Data Processing Addendum (DPA) or comprehensive documentation of GDPR compliance measures [15]. EU data residency options are not documented, and there is no evidence of EU-specific infrastructure or data boundary commitments [15].

Security Certifications:

Otter.ai's security certifications are limited compared to enterprise-grade competitors. The company does not publicly claim SOC 2 Type II, ISO 27001, or ISO 42001 certifications [15]. This absence of third-party security audits is a significant gap for organizations with stringent compliance requirements [15].

Meeting Transcription Risks:

The August 2025 class action lawsuit (*Brewer v. Otter.ai*, US District Court, Northern District of California, Case No. 5:25-cv-06911) alleges that Otter.ai "unlawfully records conversations in popular video conferencing platforms without the consent of all participants" [11]. The complaint asserts that Otter records not only account holders but also "unsuspecting third parties involved in customer's meetings, allegedly violating federal and California wiretap laws" [11]. The lawsuit further claims that Otter shifts compliance responsibility to customers by instructing them to "make sure you have the necessary permissions" rather than obtaining proper consent itself [11].

Integration Risks:

Otter.ai integrates with Google Calendar, Microsoft Outlook, Zoom, and other platforms [15]. The lawsuit alleges that Otter's automatic meeting joining feature records participants without their knowledge or consent, creating significant legal and compliance risks for organizations [11]. Third-party participants may be recorded without being informed that an AI transcription service is active [11].

Prompt/Document Leakage Risks:

Otter.ai's privacy policy indicates that the platform collects and processes extensive user data, including "Platform Information" from connected third-party services [15]. The lack of detailed documentation about data isolation, access controls, and leakage prevention mechanisms raises concerns about inadvertent data exposure [15].

Enterprise Governance:

Otter.ai offers Business and Enterprise plans with admin controls, but public documentation of governance features is limited [15]. The platform does not provide comprehensive audit logs, data lineage tracking, or advanced permission management comparable to enterprise-grade competitors [15].

Known Incidents:

The August 15, 2025 class action lawsuit is the most significant known incident [11]. The case alleges violations of federal wiretap laws (18 U.S.C. § 2511) and California's Invasion of Privacy Act (Cal. Penal Code § 632), seeking damages for unauthorized recording and use of conversational data to train AI models [11]. As of June 2026, the lawsuit remains active and unresolved [11].

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	High	Active class action lawsuit; unclear consent mechanisms; limited certifications
R&D	High	No EU data residency; potential model training on customer data; litigation risk
HR Data	Very High	Third-party consent violations alleged in lawsuit; legal liability for employers
Customer Data	High	No DPA; limited GDPR documentation; indefinite data retention alleged
Legal/Compliance Meetings	Very High	Wiretap law violations alleged; unsuitable for privileged communications
Public Sector/University (EU)	Very High	No ISO certifications; no EU data residency; active litigation; GDPR compliance unclear

3.4 Fellow

Data Processing & Storage:

Fellow processes meeting agendas, notes, action items, and feedback within its platform [16]. The service integrates with calendar systems and video conferencing platforms to provide meeting management features [16]. Data is encrypted at rest and in transit, with server infrastructure hosted on cloud providers [16]. Specific details about data residency options and regional storage are not extensively documented in public materials [16].

Model Training Policies:

Fellow's public documentation does not provide explicit statements about whether customer data is used to train AI models [16]. The absence of a clear no-training policy is a notable gap compared to enterprise-grade competitors [16].

GDPR/EU Compliance:

Fellow claims GDPR compliance and offers data subject rights mechanisms [16]. However, the company does not provide a publicly available Data Processing Addendum (DPA) or comprehensive documentation of GDPR compliance measures [16]. EU data residency options are not documented [16].

Security Certifications:

Fellow holds SOC 2 Type II certification [16]. The company claims HIPAA compliance but does not provide public documentation of a Business Associate Agreement (BAA) or HIPAA audit reports [16]. Fellow does not hold ISO 27001 or ISO 42001 certifications [16].

Meeting Transcription Risks:

Fellow does not provide native transcription services but integrates with third-party transcription tools [16]. The primary risk lies in the integration layer, where meeting content may be shared with external services without comprehensive visibility or control [16].

Integration Risks:

Fellow integrates with Google Calendar, Microsoft Outlook, Zoom, Microsoft Teams, and other platforms [16]. These integrations access meeting metadata, participant information, and potentially meeting content [16]. The security and privacy practices of these integrations depend on both Fellow's controls and the third-party platforms' policies [16].

Prompt/Document Leakage Risks:

Fellow's documentation does not provide detailed information about data isolation, access controls, or leakage prevention mechanisms [16]. The platform's reliance on integrations with multiple third-party services introduces potential data exposure risks [16].

Enterprise Governance:

Fellow provides admin controls for user management, permissions, and access [16]. The platform supports Single Sign-On (SSO) through Google, Office365, Okta, OneLogin, and custom OIDC providers [16]. However, advanced governance features such as comprehensive audit logs, data lineage tracking, and granular permission management are not extensively documented [16].

Known Incidents:

No major security incidents or data breaches specific to Fellow have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Medium	SOC 2 certified but limited public documentation of data practices
R&D	Medium	No clear no-training policy; EU data residency not documented
HR Data	Medium-High	HIPAA claimed but not verified; limited governance documentation
Customer Data	Medium	No public DPA; GDPR compliance claimed but not detailed
Legal/Compliance Meetings	Medium-High	Limited documentation of security controls and data isolation
Public Sector/University (EU)	Medium-High	No ISO 27001/42001; EU data residency not documented; limited compliance transparency

3.5 Read AI

Data Processing & Storage:

Read AI processes meeting transcripts, email content, and data from integrated services to provide AI-powered insights and summaries [10]. The platform emphasizes a privacy-by-design architecture where "data from integrated services stays inside each user's own knowledge base by default" [10]. Sharing occurs "piece by piece" with explicit user control [10]. The internal authorization service runs "half a billion permission checks daily" to enforce data boundaries [10].

Model Training Policies:

Read AI explicitly states that "customer data is not used to train models unless a customer actively opts in" [10]. This opt-in approach is more restrictive than many competitors and aligns with privacy-by-design principles [10]. The platform's default position is no training on customer data [10].

GDPR/EU Compliance:

Read AI claims GDPR compliance but does not provide extensive public documentation of compliance measures [10]. A Data Processing Addendum (DPA) is not publicly available, and EU data residency options are not documented [10]. The company's privacy stance is articulated in blog posts and marketing materials but lacks the formal compliance documentation typical of enterprise-grade platforms [10].

Security Certifications:

Read AI holds SOC 2 Type II certification [10]. The company does not hold ISO 27001 or ISO 42001 certifications [10]. HIPAA compliance is not claimed or documented [10].

Meeting Transcription Risks:

Read AI provides meeting transcription and analysis features. The platform's user-level permission model reduces the risk of inadvertent data exposure within an organization [10]. However, the lack of comprehensive documentation about third-party consent mechanisms and recording notifications creates potential compliance risks [10].

Integration Risks:

Read AI integrates with email, calendar, and communication platforms [10]. The platform's architecture isolates data at the user level, meaning "no other user in the organization can access email from a colleague's inbox when running their own search" [10]. This design reduces integration risks compared to platforms that aggregate organizational data into shared pools [10].

Prompt/Document Leakage Risks:

Read AI's permission model is designed to prevent data leakage between users [10]. However, the platform's reliance on third-party AI providers and the lack of detailed documentation about prompt injection protection and adversarial attack mitigation are notable gaps [10].

Enterprise Governance:

Read AI provides admin controls for user management and access [10]. The platform's emphasis on user-level data control means that traditional centralized governance features (such as

organization-wide audit logs or admin access to all user data) may be limited [10]. This trade-off prioritizes privacy but may complicate compliance and eDiscovery requirements [10].

Known Incidents:

No major security incidents or data breaches specific to Read AI have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Medium	SOC 2 certified; strong privacy architecture but limited public compliance documentation
R&D	Medium	No training on customer data (opt-in only); EU data residency not documented
HR Data	Medium	User-level isolation reduces risk but no HIPAA; limited governance documentation
Customer Data	Medium	No public DPA; GDPR claimed but not extensively documented
Legal/Compliance Meetings	Medium-High	Strong privacy model but limited audit/eDiscovery features
Public Sector/University (EU)	Medium-High	No ISO 27001/42001; EU data residency not documented; limited formal compliance evidence

3.6 Fireflies.ai

Data Processing & Storage:

Fireflies.ai processes meeting audio, generates transcripts, and provides AI-powered summaries and action items [5]. The platform offers a "Private Storage" option on Enterprise plans, allowing organizations to store meeting data at their preferred location [5]. Data is encrypted at rest (AES-256) and in transit (TLS) [5]. Fireflies implements a "0-day data retention policy with all vendors and partners," meaning that data shared with third-party AI providers is not retained by those providers [5].

Model Training Policies:

Fireflies.ai explicitly states that it does not train AI models on customer data by default [5]. The 0-day retention policy with vendors ensures that third-party AI providers do not retain or use customer data for training purposes [5]. This policy is a significant differentiator among meeting transcription tools [5].

GDPR/EU Compliance:

Fireflies.ai is GDPR compliant and provides data subject rights mechanisms [5]. However, the company does not provide a publicly available Data Processing Addendum (DPA) or comprehensive documentation of GDPR compliance measures [5]. EU data residency is available through the Private Storage option on Enterprise plans, but the default infrastructure and data routing are not documented [5].

Security Certifications:

Fireflies.ai holds SOC 2 Type II certification [5]. The company also claims HIPAA compliance with a Business Associate Agreement (BAA) available for Enterprise customers [5]. Fireflies does not hold ISO 27001 or ISO 42001 certifications [5].

Meeting Transcription Risks:

Fireflies.ai automatically joins meetings to provide transcription services. The primary risk is third-party consent: participants may not be aware that an AI transcription service is recording and processing the meeting [5]. Organizations must ensure that meeting hosts obtain proper consent from all participants before enabling Fireflies [5].

Integration Risks:

Fireflies.ai integrates with Zoom, Microsoft Teams, Google Meet, and other video conferencing platforms [5]. The platform also integrates with CRM systems, project management tools, and collaboration platforms [5]. These integrations access meeting metadata, participant information, and potentially meeting content [5]. The Private Storage option on Enterprise plans mitigates some integration risks by allowing organizations to control where data is stored [5].

Prompt/Document Leakage Risks:

Fireflies.ai's documentation does not provide detailed information about prompt injection protection or adversarial attack mitigation [5]. The platform's reliance on third-party AI providers introduces potential data exposure risks, though the 0-day retention policy reduces the risk of long-term data retention by vendors [5].

Enterprise Governance:

Fireflies.ai provides a "Rules Engine" for Enterprise customers, allowing admins to automate workflows and control meeting access and privacy [5]. The platform supports "Super Admin" status, granting full meeting access to the entire workspace [5]. Single Sign-On (SSO) is available for enterprise authentication [5]. Custom data retention policies allow organizations to control how long meeting data is stored [5].

Known Incidents:

No major security incidents or data breaches specific to Fireflies.ai have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Low-Medium	SOC 2 certified; 0-day retention policy; no training on customer data
R&D	Medium	Private Storage option available (Enterprise); 0-day retention reduces risk
HR Data	Medium	HIPAA available (Enterprise); third-party consent risks remain
Customer Data	Medium	0-day retention policy; Private Storage option; no public DPA
Legal/Compliance Meetings	Medium	Strong data retention policy but third-party consent risks; no ISO 42001
Public Sector/University (EU)	Medium	GDPR compliant; Private Storage option; no ISO 27001/42001; EU data residency requires Enterprise plan

3.7 ChatGPT Business/Enterprise

Data Processing & Storage:

ChatGPT Business and Enterprise process user prompts, generate responses, and store conversation history [7]. Data is encrypted at rest (AES-256) and in transit (TLS 1.2+) [7]. OpenAI states that "by default, we do not use your business data for training our models" [7]. Conversation history is retained according to workspace admin settings, with deleted conversations removed from systems within 30 days unless legally required to retain them [7].

Model Training Policies:

OpenAI explicitly states that it does not train models on business data from ChatGPT Business, ChatGPT Enterprise, ChatGPT for Healthcare, ChatGPT Edu, ChatGPT for Teachers, or the API Platform (after March 1, 2023) by default [7]. Customers can opt in to share data for model improvement through explicit feedback mechanisms [7]. This no-training policy is a core commitment of OpenAI's enterprise offerings [7].

GDPR/EU Compliance:

OpenAI is GDPR compliant and offers a Data Processing Addendum (DPA) for enterprise customers [7]. However, OpenAI does not guarantee EU data residency, and data processing occurs primarily in US-based infrastructure [7]. This limitation is significant for EU organizations with strict data localization requirements [7].

The Italian Data Protection Authority (Garante) temporarily banned ChatGPT in March 2023, citing GDPR violations including failure to provide transparency information, absence of a legal basis for training data processing, inaccuracy in outputs, and failure to verify users' age [13]. OpenAI subsequently implemented corrective measures, and the ban was lifted [13]. In January 2025, the Italian Garante issued a €15 million fine to OpenAI for continued GDPR violations [13]. This regulatory history underscores the ongoing compliance challenges faced by OpenAI in the EU [13].

Security Certifications:

ChatGPT Enterprise holds SOC 2 Type II certification [7]. OpenAI does not hold ISO 27001 or ISO 42001 certifications [7]. HIPAA compliance is available through a Business Associate Agreement (BAA) for the API Platform and ChatGPT for Healthcare [7].

Meeting Transcription Risks:

ChatGPT does not provide native meeting transcription services. However, users may paste meeting transcripts or sensitive content into ChatGPT prompts, creating data exposure risks [7]. Organizations must implement policies to prevent employees from entering confidential information into ChatGPT [7].

Integration Risks:

ChatGPT Enterprise supports "apps" that enable integration with internal sources and third-party applications [7]. These integrations can send and retrieve information from connected systems, potentially exposing sensitive data [7]. Workspace admins control which apps are enabled, and each user must authenticate with connected applications before use [7]. By default, OpenAI does not train models on data accessed from apps [7].

Prompt/Document Leakage Risks:

OpenAI implements automated content classifiers and safety tools to detect harmful content and potential abuse [7]. However, the risk of prompt injection attacks and adversarial manipulation remains [7]. Users may inadvertently expose sensitive information through prompts, and organizations must implement training and technical controls to mitigate this risk [7].

Enterprise Governance:

ChatGPT Enterprise provides workspace admins with controls over user access, available features, data retention, and app permissions [7]. The Enterprise Compliance API allows admins to access audit logs of conversations and GPTs [7]. Single Sign-On (SSO) via SAML is available for enterprise authentication [7]. Workspace admins control how long data is retained, with deleted conversations removed within 30 days [7].

Known Incidents:

- March 2023: Italian Data Protection Authority temporarily banned ChatGPT for GDPR violations [13].
- January 2025: Italian Garante issued a €15 million fine to OpenAI for continued GDPR violations [13].

These incidents highlight the regulatory scrutiny faced by OpenAI and the importance of continuous compliance monitoring for organizations using ChatGPT [13].

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Medium	SOC 2 certified; no training on business data; but €15M GDPR fine and no EU data residency
R&D	Medium	No training on business data; but US-based processing and regulatory history
HR Data	Medium-High	No EU data residency; GDPR fine history; prompt leakage risks
Customer Data	Medium-High	DPA available but no EU data residency; €15M fine for GDPR violations
Legal/Compliance Meetings	High	No EU data residency; regulatory history; prompt leakage risks; unsuitable for privileged communications
Public Sector/University (EU)	High	€15M GDPR fine; no EU data residency; no ISO 27001/42001; Italy ban history

3.8 Google Gemini for Workspace

Data Processing & Storage:

Google Gemini for Workspace processes user prompts, generates responses, and accesses Google Workspace data (emails, documents, calendar, etc.) to provide AI-powered features [2]. Data is encrypted at rest and in transit using Google's standard encryption protocols [2]. Google states that "with appropriate settings," Workspace data is not used to train Gemini models [2]. Conversation history and AI-generated content are stored within the Google Workspace environment [2].

Model Training Policies:

Google explicitly states that it does not train Gemini models on Google Workspace data by default when appropriate settings are configured [2]. This policy applies to Workspace customers with enterprise agreements [2]. However, the conditional nature of this commitment ("with appropriate settings") requires organizations to verify their configuration to ensure no-training protections are active [2].

GDPR/EU Compliance:

Google Gemini for Workspace is GDPR compliant and offers comprehensive data protection mechanisms [2]. Google provides EU data residency options, allowing organizations to specify that Workspace data remains within the EU [2]. This capability is critical for EU organizations with strict data localization requirements [2]. Google's Data Processing Amendment (DPA) covers Gemini for Workspace [2].

Security Certifications:

Google Gemini for Workspace holds SOC 2 Type II and ISO 27001 certifications [2]. Google does not hold ISO 42001 certification for Gemini [2]. HIPAA compliance is not available for Gemini for Workspace [2].

Meeting Transcription Risks:

Google Meet provides native transcription features integrated with Gemini. Transcripts are stored within Google Workspace and subject to the same access controls as other Workspace content [2]. The primary risk is ensuring that meeting participants are aware of and consent to AI-powered transcription [2].

Integration Risks:

Gemini for Workspace integrates deeply with Gmail, Google Drive, Google Docs, Google Sheets, Google Slides, Google Calendar, and Google Meet [2]. This integration provides powerful AI capabilities but also creates a broad attack surface if Workspace permissions are misconfigured [2]. Organizations must ensure that Workspace access controls are properly configured to prevent unauthorized data access by Gemini [2].

Prompt/Document Leakage Risks:

Google implements content filtering and safety mechanisms to prevent harmful outputs and detect potential abuse [2]. However, specific details about prompt injection protection and adversarial attack mitigation for Gemini are not extensively documented in public materials [2]. The risk of users inadvertently exposing sensitive information through prompts remains [2].

Enterprise Governance:

Google Workspace admins can control Gemini feature availability at the organizational unit level [2]. Audit logs track Gemini usage and data access [2]. Data retention policies can be configured through Google Vault [2]. However, granular controls over specific Gemini features (such as disabling certain AI capabilities while enabling others) may be limited [2].

Known Incidents:

No major security incidents or data breaches specific to Google Gemini for Workspace have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Low	SOC 2 and ISO 27001 certified; GDPR compliant; EU data residency available
R&D	Low	No training on Workspace data (with appropriate settings); EU data residency
HR Data	Low-Medium	Strong compliance but requires proper Workspace permission configuration
Customer Data	Low	GDPR compliant; EU data residency; DPA available
Legal/Compliance Meetings	Low-Medium	Strong security but requires verification of no-training settings
Public Sector/University (EU)	Low	GDPR compliant; EU data residency; ISO 27001 certified; no ISO 42001

3.9 Claude Team/Enterprise

Data Processing & Storage:

Anthropic Claude Team and Enterprise process user prompts, generate responses, and store conversation history [3]. Data is encrypted at rest and in transit [3]. Anthropic states that it does not train models on "Claude for Work conversations" (Team and Enterprise plans) [3]. Conversation history is stored according to workspace settings, with deleted conversations removed from systems within 30 days [3].

Model Training Policies:

Anthropic explicitly states that it does not use Claude for Work conversations (Team and Enterprise) to train generative models [3]. This no-training policy is a core commitment of Anthropic's enterprise offerings [3]. API data is also not used for training by default [3].

GDPR/EU Compliance:

Anthropic is GDPR compliant and offers a Data Processing Addendum (DPA) for enterprise customers [3]. However, Anthropic does not guarantee EU data residency, and data processing occurs on AWS and Google Cloud Platform infrastructure worldwide [3]. This limitation is significant for EU organizations with strict data localization requirements [3].

Security Certifications:

Anthropic Claude Enterprise holds an exceptionally comprehensive set of certifications: SOC 2 Type II, ISO 27001, ISO 42001, CSA STAR, HIPAA (Enterprise), NIST 800-171, and FedRAMP High (for Claude for Government) [3]. The ISO 42001 certification is particularly notable, as it demonstrates adherence to international standards for AI management systems [3]. This certification portfolio is among the most robust in the industry [3].

Meeting Transcription Risks:

Claude does not provide native meeting transcription services. However, users may paste meeting transcripts into Claude prompts, creating data exposure risks [3]. Organizations must implement policies to prevent employees from entering confidential information into Claude [3].

Integration Risks:

Claude Enterprise supports integrations with internal knowledge bases and third-party applications [3]. These integrations can access organizational data to provide context for AI responses [3]. Anthropic's architecture respects existing permissions, and users must authenticate with connected applications before use [3]. By default, data accessed through integrations is not used for model training [3].

Prompt/Document Leakage Risks:

Anthropic implements safeguards to detect and mitigate harmful content, prompt injection attacks, and adversarial manipulation [3]. However, the risk of users inadvertently exposing sensitive information through prompts remains [3]. Organizations must implement training and technical controls to mitigate this risk [3].

Enterprise Governance:

Claude Enterprise provides workspace admins with controls over user access, available features, and data retention [3]. Audit logs track user activity for compliance [3]. Single Sign-On (SSO) and SCIM provisioning are available for enterprise authentication and user management [3].

Known Incidents:

No major security incidents or data breaches specific to Anthropic Claude have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Low-Medium	Comprehensive certifications (ISO 42001, SOC 2, ISO 27001); no training on enterprise data; but no EU data residency guarantee
R&D	Low-Medium	Strong security posture; no training on enterprise data; but AWS/GCP worldwide infrastructure
HR Data	Medium	HIPAA available (Enterprise); but no EU data residency; prompt leakage risks
Customer Data	Low-Medium	DPA available; comprehensive certifications; but no EU data residency
Legal/Compliance Meetings	Medium	Strong security (ISO 42001, FedRAMP High for Gov); but no EU data residency; prompt leakage risks
Public Sector/University (EU)	Medium	ISO 42001 certified; FedRAMP High (for Gov); but no EU data residency; GDPR compliant but data processed worldwide

3.10 Perplexity Enterprise Pro

Data Processing & Storage:

Perplexity Enterprise Pro processes user queries, searches internal knowledge bases and the web, and generates AI-powered responses with citations [6]. Data is encrypted at rest and in transit [6]. Perplexity states that it never trains LLMs on enterprise customer data [6]. The platform offers configurable file retention, with automatic file deletion available up to 1 day [6].

Model Training Policies:

Perplexity explicitly states that it does not train LLMs on enterprise customer data [6]. This no-training policy is a core commitment of Perplexity's enterprise offerings [6]. The platform's emphasis on privacy and data protection is central to its value proposition [6].

GDPR/EU Compliance:

Perplexity is GDPR compliant [6]. However, the company does not provide extensive public documentation of GDPR compliance measures, and EU data residency options are not documented [6]. A Data Processing Addendum (DPA) is not publicly available [6].

Security Certifications:

Perplexity Enterprise Pro holds SOC 2 Type II certification [6]. The company also claims HIPAA compliance [6]. Perplexity does not hold ISO 27001 or ISO 42001 certifications [6].

Meeting Transcription Risks:

Perplexity does not provide native meeting transcription services. However, users may paste meeting transcripts into Perplexity queries, creating data exposure risks [6]. Organizations must implement policies to prevent employees from entering confidential information into Perplexity [6].

Integration Risks:

Perplexity Enterprise supports integrations with internal knowledge bases and file systems [6]. These integrations can access organizational data to provide context for AI responses [6]. The platform's architecture includes user management controls to limit who can upload and download files and how answers are shared [6]. However, detailed documentation of integration security and data isolation mechanisms is limited [6].

Prompt/Document Leakage Risks:

Perplexity's documentation does not provide detailed information about prompt injection protection or adversarial attack mitigation [6]. The platform's reliance on multiple AI models (the platform "orchestrates the best models across your files and tools") introduces potential data exposure risks [6].

Enterprise Governance:

Perplexity Enterprise provides user management controls, audit logs for compliance, and Single Sign-On (SSO) with SCIM provisioning [6]. Configurable file retention allows organizations to control how long data is stored [6]. However, advanced governance features such as granular permission management and data lineage tracking are not extensively documented [6].

Known Incidents:

No major security incidents or data breaches specific to Perplexity Enterprise have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Medium	SOC 2 certified; no training on enterprise data; but limited public compliance documentation
R&D	Medium	No training on enterprise data; but EU data residency not documented
HR Data	Medium	HIPAA claimed; but limited governance documentation; prompt leakage risks
Customer Data	Medium	No public DPA; GDPR claimed but not extensively documented
Legal/Compliance Meetings	Medium-High	Limited documentation of security controls; prompt leakage risks
Public Sector/University (EU)	Medium-High	No ISO 27001/42001; EU data residency not documented; limited formal compliance evidence

3.11 Notion AI/Enterprise AI

Data Processing & Storage:

Notion AI processes user prompts, accesses Notion workspace content, and generates AI-powered responses [8]. Data is encrypted at rest and in transit [8]. Notion states that it does not use customer content to train AI models [8]. However, Notion AI relies on third-party AI providers (Anthropic and OpenAI) as subprocessors, meaning customer data is shared with these vendors to provide AI features [8].

Model Training Policies:

Notion explicitly states that it does not train AI models on customer content [8]. However, the use of Anthropic and OpenAI as subprocessors means that customer data is processed by these third-party providers [8]. Notion's agreements with these subprocessors include commitments that they will not train on Notion customer data, but this introduces an additional layer of trust and compliance complexity [8].

GDPR/EU Compliance:

Notion is GDPR compliant and provides data subject rights mechanisms [8]. The company offers a Data Processing Addendum (DPA) for enterprise customers [8]. However, EU data residency options are not documented, and the use of third-party AI subprocessors (Anthropic and OpenAI) introduces data transfer considerations [8].

Security Certifications:

Notion holds SOC 2 Type II and ISO 27001 certifications [8]. Notion does not hold ISO 42001 certification [8]. HIPAA compliance is not available [8].

Meeting Transcription Risks:

Notion does not provide native meeting transcription services. However, users may paste meeting transcripts into Notion AI prompts, creating data exposure risks [8]. Organizations must implement policies to prevent employees from entering confidential information into Notion AI [8].

Integration Risks:

Notion integrates with various third-party services and provides API access for custom integrations [8]. The primary integration risk lies in the use of Anthropic and OpenAI as AI subprocessors, which means customer data is shared with these vendors [8]. Organizations must evaluate the security and privacy practices of these subprocessors as part of their risk assessment [8].

Prompt/Document Leakage Risks:

Notion's documentation does not provide detailed information about prompt injection protection or adversarial attack mitigation [8]. The reliance on third-party AI providers introduces potential data exposure risks, though Notion's agreements with these providers include no-training commitments [8].

Enterprise Governance:

Notion Enterprise provides admin controls for user management, permissions, and access [8]. The platform supports Single Sign-On (SSO) via SAML 2.0 and user provisioning through SCIM [8]. Audit log features allow admins to track activity [8]. However, granular controls over AI feature usage and data sharing with subprocessors are not extensively documented [8].

Known Incidents:

No major security incidents or data breaches specific to Notion AI have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Medium	SOC 2 and ISO 27001 certified; but uses Anthropic/OpenAI as subprocessors
R&D	Medium	No training on customer content; but third-party subprocessors; EU data residency not documented
HR Data	Medium	Strong certifications but third-party subprocessors introduce complexity; no HIPAA
Customer Data	Medium	DPA available; but third-party subprocessors; EU data residency not documented
Legal/Compliance Meetings	Medium-High	Third-party subprocessors; prompt leakage risks; limited AI-specific governance documentation
Public Sector/University (EU)	Medium	ISO 27001 certified; GDPR compliant; but third-party subprocessors; no ISO 42001; EU data residency not documented

3.12 Fathom

Data Processing & Storage:

Fathom processes meeting audio, generates transcripts, and provides AI-powered summaries [12]. Data is encrypted at rest and in transit [12]. However, specific details about data residency options, regional storage, and data retention policies are not comprehensively documented in public materials [12].

Model Training Policies:

Fathom's public documentation does not provide explicit statements about whether customer data is used to train AI models [12]. The absence of a clear no-training policy is a notable gap compared to enterprise-grade competitors [12].

GDPR/EU Compliance:

Fathom claims GDPR compliance [12]. However, the company does not provide a publicly available Data Processing Addendum (DPA) or comprehensive documentation of GDPR compliance measures [12]. EU data residency options are not documented [12].

Security Certifications:

Fathom holds SOC 2 Type II certification [12]. The company does not hold ISO 27001 or ISO 42001 certifications [12]. HIPAA compliance is not claimed or documented [12].

Meeting Transcription Risks:

Fathom automatically joins meetings to provide transcription services. The primary risk is third-party consent: participants may not be aware that an AI transcription service is recording and processing the meeting [12]. Organizations must ensure that meeting hosts obtain proper consent from all participants before enabling Fathom [12].

Integration Risks:

Fathom integrates with Zoom, Microsoft Teams, Google Meet, and other video conferencing platforms [12]. The platform also integrates with CRM systems and collaboration tools [12]. These integrations access meeting metadata, participant information, and potentially meeting content [12]. The security and privacy practices of these integrations are not extensively documented [12].

Prompt/Document Leakage Risks:

Fathom's documentation does not provide detailed information about data isolation, access controls, or leakage prevention mechanisms [12]. The lack of comprehensive security documentation raises concerns about inadvertent data exposure [12].

Enterprise Governance:

Fathom provides admin controls for user management and access [12]. However, advanced governance features such as comprehensive audit logs, data lineage tracking, and granular permission management are not extensively documented [12].

Known Incidents:

No major security incidents or data breaches specific to Fathom have been publicly reported as of June 2026.

Risk Classification:

Use Case	Risk Level	Rationale
General Office Use	Medium	SOC 2 certified but limited public documentation of data practices
R&D	Medium-High	No clear no-training policy; EU data residency not documented; limited certifications
HR Data	Medium-High	No HIPAA; third-party consent risks; limited governance documentation
Customer Data	Medium-High	No public DPA; GDPR claimed but not detailed; limited security documentation
Legal/Compliance Meetings	High	Limited documentation of security controls; third-party consent risks; no ISO certifications
Public Sector/University (EU)	High	No ISO 27001/42001; EU data residency not documented; limited compliance transparency; no clear no-training policy

4. Cross-Tool Risk Comparison

4.1 Model Training on Customer Data

Lowest Risk (Explicit No-Training Policies):

- Microsoft 365 Copilot [1]
- Google Gemini for Workspace (with appropriate settings) [2]
- Anthropic Claude Team/Enterprise [3]
- ChatGPT Business/Enterprise [7]
- Fireflies.ai (0-day retention with vendors) [5]
- Read AI (opt-in only) [10]
- Perplexity Enterprise Pro [6]
- Notion AI (via subprocessor agreements) [8]

Moderate Risk (Unclear or Conditional Policies):

- Zoom AI Companion (policy clarified after 2023 controversy) [14]
- Fellow (no explicit statement) [16]
- Fathom (no explicit statement) [12]

Highest Risk (Alleged Training on Customer Data):

- Otter.ai (class action lawsuit alleges indefinite retention and use for training) [11]

4.2 EU Data Residency

Full EU Data Boundary:

- Microsoft 365 Copilot (except Anthropic subprocessor) [1]
- Google Gemini for Workspace [2]

Partial EU Infrastructure:

- Zoom AI Companion (EU infrastructure available but not guaranteed) [4]
- Fireflies.ai (Private Storage option on Enterprise) [5]

No EU Data Residency Guarantee:

- ChatGPT Business/Enterprise (US-based) [7]
- Anthropic Claude Team/Enterprise (AWS/GCP worldwide) [3]
- Perplexity Enterprise Pro (not documented) [6]
- Notion AI (not documented) [8]

- Otter.ai (not documented) [15]
- Fellow (not documented) [16]
- Read AI (not documented) [10]
- Fathom (not documented) [12]

4.3 Security Certifications

Most Comprehensive (ISO 42001 AI Management Systems):

- Microsoft 365 Copilot (SOC 2, ISO 27001, ISO 42001, HIPAA) [1]
- Anthropic Claude Enterprise (SOC 2, ISO 27001, ISO 42001, CSA STAR, HIPAA, NIST 800-171, FedRAMP High) [3]

Strong Enterprise Certifications (SOC 2 + ISO 27001):

- Google Gemini for Workspace (SOC 2, ISO 27001) [2]
- Notion AI (SOC 2, ISO 27001) [8]

SOC 2 Type II Only:

- Zoom AI Companion (SOC 2, ISO 27701) [4]
- ChatGPT Business/Enterprise (SOC 2) [7]
- Fireflies.ai (SOC 2, HIPAA on Enterprise) [5]
- Perplexity Enterprise Pro (SOC 2, HIPAA claimed) [6]
- Fellow (SOC 2, HIPAA claimed) [16]
- Read AI (SOC 2) [10]
- Fathom (SOC 2) [12]

Limited Certifications:

- Otter.ai (no major certifications documented) [15]

4.4 Known Regulatory Incidents

Major Incidents:

- ChatGPT: Italy ban (March 2023), €15M fine (January 2025) [13]
- Zoom AI Companion: ToS controversy (August 2023, resolved) [14]
- Otter.ai: Class action lawsuit (August 2025, active) [11]

No Known Incidents:

- Microsoft 365 Copilot, Google Gemini for Workspace, Anthropic Claude, Fireflies.ai, Perplexity, Notion AI, Fellow, Read AI, Fathom

4.5 Third-Party Consent Risks (Meeting Transcription)

Highest Risk (Automatic Joining, Active Litigation):

- Otter.ai (class action lawsuit for unauthorized recording) [11]

High Risk (Automatic Joining, Limited Consent Documentation):

- Fireflies.ai [5]
- Fathom [12]

Moderate Risk (Automatic Joining, Some Consent Mechanisms):

- Zoom AI Companion [4]
- Read AI [10]

Lower Risk (Native Integration, Existing Consent Flows):

- Microsoft 365 Copilot (Teams native) [1]
- Google Gemini for Workspace (Meet native) [2]

Not Applicable (No Native Transcription):

- ChatGPT, Claude, Perplexity, Notion AI, Fellow

5. GDPR-Specific Concerns for EU Organizations

5.1 Data Localization Requirements

EU organizations subject to strict data localization requirements (e.g., public sector, healthcare, financial services) face significant constraints when deploying AI tools. Only Microsoft 365 Copilot and Google Gemini for Workspace offer comprehensive EU Data Boundary commitments with in-region processing guarantees [1], [2]. All other tools either process data outside the EU or rely on subprocessors not covered by EU data residency commitments.

Critical Exception: Microsoft 365 Copilot's use of Anthropic models is explicitly out of scope for the EU Data Boundary [1]. Organizations must evaluate whether this exception is acceptable for their use cases.

5.2 Legal Basis for Processing

Under GDPR Article 6, organizations must establish a lawful basis for processing personal data. For AI tools that process employee and customer data, the most common legal bases are:

- Legitimate interests (Article 6(1)(f)): Requires a balancing test demonstrating that the organization's interests do not override data subjects' rights.
- Contractual necessity (Article 6(1)(b)): Applicable when AI processing is necessary to fulfill contractual obligations.
- Consent (Article 6(1)(a)): Required for processing special categories of personal data (Article 9) or when other legal bases do not apply.

Meeting transcription tools present particular challenges because they often process data of third parties (meeting participants from other organizations) who have not consented to AI processing. The Italian DPA's actions against ChatGPT emphasized the importance of transparency and legal basis for AI training [13]. Organizations must ensure that:

1. Meeting participants are informed that AI transcription is active.
2. Participants have the opportunity to object or request exclusion.
3. The legal basis for processing is clearly documented and communicated.

5.3 Data Subject Rights

GDPR Articles 15-22 grant data subjects rights including access, rectification, erasure, restriction, portability, and objection. AI tools must support these rights, which requires:

- Data access mechanisms: Users must be able to retrieve their data processed by AI systems.
- Deletion capabilities: Organizations must be able to delete data upon request (subject to legal retention requirements).
- Rectification processes: Inaccurate data must be correctable.

Challenges:

- ChatGPT faced criticism from the Italian DPA for inaccuracy in outputs, which the DPA argued violated data subjects' rights to accurate processing [13].
- Otter.ai allegedly retains conversational data indefinitely, which may conflict with data minimization and storage limitation principles [11].

5.4 Data Processing Agreements (DPAs)

Under GDPR Article 28, organizations must execute Data Processing Agreements with AI vendors acting as data processors. DPAs must specify:

- The subject matter, duration, nature, and purpose of processing.
- The types of personal data and categories of data subjects.
- The processor's obligations, including security measures and subprocessor management.
- Data subject rights support and breach notification procedures.

DPA Availability:

- Available: Microsoft 365 Copilot [1], Zoom AI Companion [4], ChatGPT Enterprise [7], Google Gemini for Workspace [2], Anthropic Claude [3], Notion AI [8]
- Not Publicly Available: Otter.ai [15], Fellow [16], Read AI [10], Fireflies.ai [5], Perplexity [6], Fathom [12]

5.5 Subprocessor Management

GDPR Article 28(4) requires that processors obtain prior authorization before engaging subprocessors. Organizations must:

- Maintain a list of subprocessors.
- Receive notification of subprocessor changes.
- Have the right to object to new subprocessors.

Subprocessor Transparency:

- High Transparency: Microsoft 365 Copilot (Anthropic disclosed) [1], Anthropic Claude (subprocessor list available) [3], Notion AI (Anthropic/OpenAI disclosed) [8]
- Limited Transparency: Most other tools do not provide comprehensive subprocessor lists.

5.6 Regulatory Precedents

The Italian DPA's actions against ChatGPT establish important precedents for EU organizations:

1. Transparency is mandatory: AI tools must clearly explain how they process personal data [13].
2. Legal basis must be documented: Organizations cannot rely on vague or post-hoc justifications for AI processing [13].
3. Age verification is required: Tools accessible to minors must implement age verification mechanisms [13].
4. Accuracy matters: AI outputs that contain personal data must be accurate, or the tool must clearly indicate uncertainty [13].

Implications for Organizations:

- Conduct Data Protection Impact Assessments (DPIAs) before deploying AI tools, especially for high-risk processing.
- Document legal basis for AI processing in privacy notices and internal policies.
- Implement technical and organizational measures to ensure accuracy and transparency.
- Monitor regulatory developments and adjust practices accordingly.

6. Recommended Enterprise Policy Framework

6.1 Pre-Deployment Assessment

Before deploying any AI tool, organizations should conduct a structured risk assessment:

Step 1: Data Classification

- Identify what types of data will be processed (personal data, sensitive data, confidential business information).
- Classify data according to organizational data classification policies (e.g., Public, Internal, Confidential, Restricted).

Step 2: Use Case Evaluation

- Define specific use cases for the AI tool (e.g., meeting transcription, document summarization, customer support).
- Assess whether the use case involves high-risk processing under GDPR (e.g., automated decision-making, special categories of data).

Step 3: Vendor Due Diligence

- Review vendor security certifications (SOC 2, ISO 27001, ISO 42001).
- Verify GDPR compliance and DPA availability.
- Assess data residency options and subprocessor transparency.
- Review model training policies and data retention practices.
- Check for known incidents or regulatory actions.

Step 4: Data Protection Impact Assessment (DPIA)

- Conduct a DPIA for high-risk processing activities (required under GDPR Article 35).
- Document risks, mitigation measures, and residual risks.
- Consult with the Data Protection Officer (DPO) and, if necessary, the supervisory authority.
-

6.2 Acceptable Use Policy

Organizations should establish clear acceptable use policies for AI tools:

Prohibited Uses:

- Entering confidential customer data, trade secrets, or legally privileged information into AI tools without proper safeguards.
- Using AI tools to process special categories of personal data (health, biometric, genetic data) without explicit authorization.
- Recording meetings with AI transcription tools without obtaining consent from all participants.
- Bypassing organizational security controls or data loss prevention (DLP) mechanisms.

Required Practices:

- Obtain consent from meeting participants before enabling AI transcription.
- Use organizational accounts (not personal accounts) for work-related AI tool usage.
- Review AI-generated outputs for accuracy before sharing externally.
- Report suspected data breaches or security incidents immediately.

6.3 Technical Controls

Data Loss Prevention (DLP):

- Implement DLP policies to detect and block sensitive data from being entered into AI tools.
- Use content inspection to identify confidential information in prompts and documents.

Access Controls:

- Enforce Single Sign-On (SSO) for all AI tools.
- Implement role-based access control (RBAC) to limit AI tool access based on job function.
- Require multi-factor authentication (MFA) for administrative accounts.

Monitoring and Auditing:

- Enable audit logging for AI tool usage.
- Monitor for anomalous activity (e.g., bulk data uploads, unusual access patterns).
- Conduct periodic access reviews to ensure least privilege.

Data Retention:

- Configure AI tools to minimize data retention periods.
- Implement automated deletion policies for AI-generated content.

- Ensure that deleted data is removed from backups and archives.

6.4 Training and Awareness

Employee Training:

- Provide mandatory training on AI tool acceptable use policies.
- Educate employees about data privacy risks and best practices.
- Conduct scenario-based training (e.g., "What should you do if a customer asks you not to record a meeting?").

Leadership Awareness:

- Brief executives and managers on AI governance requirements.
- Ensure that leadership understands the legal and reputational risks of AI misuse.

6.5 Incident Response

AI-Specific Incident Response Plan:

- Define procedures for responding to AI-related data breaches (e.g., prompt injection attacks, unauthorized data access).
- Establish escalation paths for reporting incidents to the DPO, legal team, and supervisory authorities.
- Conduct post-incident reviews to identify root causes and implement corrective actions.

6.6 Continuous Monitoring

Regulatory Monitoring:

- Track regulatory developments (e.g., EU AI Act, US state privacy laws).
- Monitor vendor compliance status and incident disclosures.
- Update policies and controls in response to regulatory changes.

Vendor Performance:

- Conduct annual vendor reviews to assess ongoing compliance.
- Request updated security certifications and audit reports.
- Evaluate vendor incident response and breach notification practices.

7. Final Ranking: Most to Least Suitable for Data-Sensitive Organizations

Tier 1: Lowest Risk (Highly Suitable)

1. Microsoft 365 Copilot

- Strengths: EU Data Boundary, ISO 42001, comprehensive certifications, no model training, mature governance, native integration with Microsoft 365.
- Limitations: Anthropic subprocessor outside EU Data Boundary.
- Best For: EU organizations, public sector, healthcare, financial services, legal/compliance.

2. Google Gemini for Workspace

- Strengths: EU data residency, ISO 27001, SOC 2, no model training (with appropriate settings), native integration with Google Workspace.
- Limitations: No ISO 42001, conditional no-training policy requires configuration verification.
- Best For: EU organizations, education, general office use, customer data processing.

3. Anthropic Claude Enterprise

- Strengths: ISO 42001, comprehensive certifications (FedRAMP High, NIST 800-171), no model training, HIPAA available.
- Limitations: No EU data residency guarantee, AWS/GCP worldwide infrastructure.
- Best For: US organizations, government (FedRAMP), healthcare (HIPAA), R&D, general office use.

Tier 2: Moderate-Low Risk (Suitable with Conditions)

4. Zoom AI Companion

- Strengths: ISO 27701, SOC 2, GDPR compliant, EU infrastructure available, no model training (after 2023 policy clarification).
- Limitations: 2023 ToS controversy, no ISO 27001/42001, third-party consent risks.
- Best For: General office use, external meetings (with consent), organizations with existing Zoom deployments.

5. Fireflies.ai Enterprise

- Strengths: SOC 2, 0-day retention with vendors, no model training, Private Storage option, HIPAA available (Enterprise).
- Limitations: No ISO certifications, third-party consent risks, EU data residency requires Enterprise plan.

- Best For: Organizations prioritizing data retention controls, sales/customer success teams (with consent protocols).

6. Perplexity Enterprise Pro

- Strengths: SOC 2, no model training, HIPAA claimed, configurable file retention.
- Limitations: No ISO certifications, limited public compliance documentation, EU data residency not documented.
- Best For: Research teams, knowledge work, organizations prioritizing search and citation features.

Tier 3: Moderate Risk (Requires Enhanced Due Diligence)

7. ChatGPT Business/Enterprise

- Strengths: SOC 2, no model training, DPA available, HIPAA available (API).
- Limitations: €15M GDPR fine (2025), Italy ban (2023), no EU data residency, no ISO certifications.
- Best For: US organizations, non-sensitive use cases, organizations with robust prompt governance.

8. Notion AI Enterprise

- Strengths: SOC 2, ISO 27001, no model training (via subprocessor agreements), DPA available.
- Limitations: Uses Anthropic/OpenAI as subprocessors, EU data residency not documented, no ISO 42001.
- Best For: Knowledge management, documentation, organizations with existing Notion deployments.

9. Fellow

- Strengths: SOC 2, HIPAA claimed, SSO support.
- Limitations: No ISO certifications, limited public documentation, no clear no-training policy, EU data residency not documented.
- Best For: Meeting management, 1-on-1s, organizations prioritizing meeting structure over transcription.

10. Read AI

- Strengths: SOC 2, strong privacy architecture (user-level isolation), no model training (opt-in only).
- Limitations: No ISO certifications, limited public compliance documentation, EU data residency not documented.
- Best For: Organizations prioritizing user-level data control, knowledge workers requiring privacy-by-design.

Tier 4: Elevated Risk (Not Recommended for Sensitive Data)

11. Otter.ai

- Strengths: Popular platform, ease of use.
- Limitations: Active class action lawsuit (unauthorized recording), no major certifications, unclear model training policy, no EU data residency, limited GDPR documentation.
- Best For: Non-sensitive, low-risk use cases only; not recommended for enterprise deployment.

12. Fathom

- Strengths: SOC 2, ease of use.
- Limitations: No ISO certifications, limited public documentation, no clear no-training policy, EU data residency not documented, third-party consent risks.
- Best For: Non-sensitive, low-risk use cases; not recommended for data-sensitive organizations.

References

[1] Microsoft Learn. (2026). Data, Privacy, and Security for Microsoft 365 Copilot. Retrieved from https://learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-privacy

[2] Google Workspace Help. (2026). Generative AI in Google Workspace Privacy Hub. Retrieved from https://knowledge.workspace.google.com/admin/generative-ai/generative-ai-in-google-workspace-privacy-hub

[3] Anthropic Trust Center. (2026). Welcome to the Anthropic Trust Center. Retrieved from https://trust.anthropic.com/

[4] Zoom. (2026). Privacy at Zoom. Retrieved from https://www.zoom.com/en/trust/privacy/privacy-statement/

[5] Fireflies.ai. (2026). Fireflies Data Security & Privacy for Meeting Notes. Retrieved from https://fireflies.ai/security

[6] Perplexity. (2026). Perplexity Enterprise. Retrieved from https://www.perplexity.ai/enterprise

[7] OpenAI. (2026). Enterprise privacy at OpenAI. Retrieved from https://openai.com/enterprise-privacy

[8] Notion. (2026). Security & Compliance. Retrieved from https://www.notion.com/security

[9] Notion. (2026). Privacy Policy. Retrieved from https://www.notion.com/privacy

[10] Read AI. (2026). AI and Privacy: What Your Company Should Know. Retrieved from https://www.read.ai/articles/ai-and-privacy-what-your-company-should-know

[11] Fisher Phillips LLP. (2025). New Lawsuit Highlights Concerns About AI Notetakers: 7 Steps Businesses Should Take. Retrieved from https://www.fisherphillips.com/en/insights/insights/new-lawsuit-highlights-concerns-about-ai-notetakers

[12] Fathom. (2026). Is Fathom secure? Retrieved from https://help.fathom.video/en/articles/296512

[13] Lewis Silkin. (2025). OpenAI faces €15 million fine as the Italian Garante strikes again. Retrieved from https://www.lewissilkin.com/insights/2025/01/14/openai-faces-15-million-fine-as-the-italian-garante-strikes-again-102jtqc

[14] Data Protection Report. (2023). Italian Garante bans Chat GPT from processing personal data of Italian data subjects. Retrieved from https://www.dataprotectionreport.com/2023/04/italian-garante-bans-chat-gpt-from-processing-personal-data-of-italian-data-subjects/

[15] Otter.ai. (2026). Privacy Policy. Retrieved from https://otter.ai/privacy-policy

[16] Fellow. (2026). Security and Compliance. Retrieved from https://help.fellow.ai/en/articles/4302231-security-and-compliance

Document Control:

Version: 1.0

Date: June 23, 2026

Classification: Internal Use – Management Review

Next Review: December 2026 or upon significant regulatory/vendor changes

Disclaimer: This review is based on publicly available information as of June 23, 2026. Organizations should conduct independent due diligence, including vendor questionnaires, security audits, and legal review, before deploying AI tools for sensitive data processing. Regulatory requirements and vendor practices may change; continuous monitoring is essential.