GenAI Policy Exposure Map

Version 8 — Public Edition · April 2026

For organizational policy and procurement decisions. Not legal advice. Verify with vendor documentation before final decisions.

LEGEND: **GREEN — Low risk / Strong protection** **YELLOW — Medium risk / Use with care** **RED — High risk / No org controls / Requires action**

Sensitive Document Use

These ratings reflect vendor data-handling policies, not a compliance determination for your organization. Before approving any service for documents containing personal data, client records, HR information, financial data, or legal communications, confirm that your specific plan tier includes the contractual protections shown in Part 2.

Do not assume AI interactions are protected by attorney-client privilege or work product doctrine. Do not process legally sensitive or privileged materials through any AI service without guidance from your legal or compliance team. When in doubt, treat the use case as requiring review.

For individual staff guidance on which tools to use for specific tasks, see the *GenAI Safety Quick Cards*.

Policy Risk Summary: Priority Tools

Use this to decide: Which high-frequency tools have acceptable policy protections and which need restrictions, staff guidance, or prohibition.

Key assumption: Ratings reflect vendor policy, not your org's actual configuration. An admin enforcement rating of GREEN means controls exist — not that your admin has applied them.

Policy Dimension	ChatGPT Free/Plus	ChatGPT Business	Claude Team	Gemini in Workspace	M365 Copilot
Training & Model Use	**RED** ON by default Opt-out in Settings	**GREEN** OFF by default Commercial plan	**GREEN** OFF — commercial Verify in Admin	**GREEN** OFF (contractual) All Workspace tiers	**GREEN** OFF (contractual) Anthropic subprocessor added Jan 2026
Data Retention	**YELLOW** Indefinite by default Delete + 30d purge	**YELLOW** 30d abuse-monitoring baseline Additional controls for qualifying orgs/API	**YELLOW** 30 days standard +30d purge; no ZDR on product UI	**YELLOW** Admin-configurable Default 18 months; min 90 days	**YELLOW** Purview policy No universal default; admin-set
Human Review	**RED** Possible Safety, quality, annotation snippets	**YELLOW** Policy violations Highly limited scope only	**YELLOW** Policy violations Not for training	**YELLOW** Flagged content only Automated trigger; not for training	**GREEN** No abuse-monitoring review Explicitly documented; distinct from Azure OAI
Admin Enforcement	**RED** None Personal account	**GREEN** Workspace admin platform.openai.com / admin.openai.com	**GREEN** Admin console claude.ai/admin; SSO, project controls	**GREEN** Admin Console admin.google.com; OU-level controls	**GREEN** M365 Admin + Purview admin.microsoft.com; Entra ID; DLP
Connectors & Agent Scope	**YELLOW** Custom GPTs: review each Per-action policy; no admin restriction	**YELLOW** Admin can restrict GPT governance; approve/block external GPTs	**YELLOW** Admin-manageable Drive/GitHub/M365; no documented disable control	**YELLOW** Drive/Gmail scope Permissions hygiene critical; DLP can restrict	**YELLOW** Reads all M365 user can see Sensitivity labels + DLP essential; tighten perms

Policy Risk Summary: Additional Tools (US)

Use this to decide: Whether to permit or restrict tools that may already be in use outside your core policy agenda — including personal-tier Claude, NotebookLM, and Perplexity.

Key assumption: The two NotebookLM entries have meaningfully different risk profiles — see the callout below this table before drawing conclusions about either.

Policy Dimension	Claude Personal	NotebookLM Free	NotebookLM Business	Perplexity Personal	Perplexity Ent. Pro
Training & Model Use	**RED** ON since Sep 2025 Must opt out!	**GREEN** OFF by default Feedback triggers use; avoid thumbs up/down	**GREEN** OFF (contractual) No review/training even with feedback	**RED** ON by default Opt-out available	**GREEN** OFF (contractual) Explicit prohibition §1.3.3
Data Retention	**RED** 5 yrs if training on; 30d if opted out; Sep 2025 change	**YELLOW** Indefinite; feedback 3yr Feedback survives acct deletion	**YELLOW** Workspace policies Admin-configurable; Workspace DPA	**YELLOW** "Reasonably necessary" No specific period stated	**YELLOW** 30 days on request Zero-ret. must be explicitly configured
Human Review	**RED** Possible When flagged or feedback submitted	**RED** If feedback given Uploaded content included; avoid thumbs up/down	**GREEN** No human review Even with feedback — explicit Workspace commitment	**YELLOW** Not documented Not addressed in official docs	**YELLOW** Not documented No explicit restriction in contracts
Admin Enforcement	**RED** None Personal account	**RED** None Personal Google account	**GREEN** Workspace Admin admin.google.com; Workspace controls	**RED** None Personal only	**YELLOW** Enterprise console Admin noted; features not fully documented
Connectors & Agent Scope	**YELLOW** Connectors in Pro/Max Personal-scope; no admin control	**YELLOW** Personal Google only No Workspace scope; no admin control	**YELLOW** Workspace scope Notebook sharing; admin-managed access	**RED** Multi-vendor routing Queries may route to multiple providers; plan/query-dependent	**YELLOW** Multi-vendor routing Training prohibition; inference routing may continue

NotebookLM Free — Conditional GREEN Rating

The GREEN Training rating is conditional on user behavior. Thumbs up/down feedback triggers training use AND human review of uploaded content, including source documents. Feedback data is retained for 3 years, independently of account deletion. If staff use NotebookLM Free, instruct them not to click feedback buttons on any sensitive content. Organizations regularly processing sensitive documents should evaluate NotebookLM for Business instead.

Policy Risk Summary: Additional Tools (Europe)

Use this to decide: Whether Mistral (French AI company) or Lumo by Proton are viable alternatives for org use — particularly where EU/Swiss data jurisdiction, privacy architecture, or training guarantees matter.

Key assumption: Lumo's zero-access encryption means training and human review are architecturally impossible, not just contractually prohibited — a materially stronger guarantee than policy commitments alone. Le Chat Team's L3 rating reflects unconfirmed training status at that tier; confirm directly with Mistral before approving for client-data workflows. Certifications for Mistral commercial plans are not yet independently verified.

Policy Dimension	Le Chat Free/Pro	Le Chat Team	Lumo Free	Lumo Professional	Lumo + Workspace
Training & Model Use	**RED** ON by default Free/Pro not excluded per privacy policy; opt-out in account settings	**YELLOW** Unconfirmed at Team tier Privacy policy exempts only Enterprise + paid API; confirm with Mistral; DPA applies	**GREEN** Architecturally off Zero-access encryption; no server logs; open-source verifiable	**GREEN** Architecturally off Zero-access encryption; SOC 2 Type II validated (Jul 2025)	**GREEN** Architecturally off Same zero-access architecture across full Proton suite
Data Retention	**YELLOW** Until user deletes 30-day abuse-monitoring logs baseline	**YELLOW** Until admin/user deletes No automatic purge; admin data export + audit logs	**GREEN** Deleted after response Optional encrypted history (zero-access); no server-side logs	**GREEN** No server-side logs Zero-access encrypted history; Ghost Mode available	**GREEN** No server-side logs Zero-access or E2E encryption across all Proton services
Human Review	**YELLOW** Abuse monitoring only Not for training or quality review — unlike OpenAI Free/Plus	**YELLOW** Policy violations only Commercial terms; not for training	**GREEN** Architecturally impossible Zero-access encryption — Proton cannot read conversation content	**GREEN** Architecturally impossible Zero-access encryption; SOC 2 Type II audit validates controls	**GREEN** Architecturally impossible Zero-access encryption across all Proton services
Admin Enforcement	**RED** None Personal account; no org controls	**GREEN** Admin API + audit logs Domain verification; data export; DPA available; no SAML SSO (Enterprise only)	**RED** None No account required; not for org deployment	**GREEN** Proton for Business admin Team access management; ISO 27001 + SOC 2 Type II validated	**GREEN** Full Proton for Business Complete admin suite; ISO 27001 + SOC 2 Type II; HIPAA support available
Connectors & Agent Scope	**YELLOW** Web search + MCP beta Each integration independently scoped; no admin restriction on Free/Pro	**YELLOW** Web search; MCP connectors Commercial DPA applies to connector data; admin can manage	**GREEN** Web search only No third-party integrations; rate limited; no logs retained	**YELLOW** Proton Drive + file uploads Encrypted ecosystem; EU/Swiss data residency; web search	**YELLOW** Full Proton suite Drive, Mail, Calendar, VPN, Pass — all zero-access or E2E encrypted

Legal confidentiality note: No vendor on this map makes attorney-client privilege representations for AI interactions. Do not assume AI interactions are protected by attorney-client privilege or work product doctrine. Legal professionals handling privileged communications or documents subject to confidentiality obligations should consult their legal team before using any service for such material.

Data Handling & Retention

Use this to decide: Whether each tool's default retention window is compatible with your data minimization requirements, and where your admin needs to actively configure deletion.

Key assumption: These are vendor defaults. Shorter retention and zero-retention options exist for some tools but typically require higher plan tiers or explicit configuration — check what your specific plan tier actually provides, not just what the vendor advertises. "Undocumented" means not verifiable from official sources — verify with the vendor before high-stakes decisions.

Service / Plan	Retention & Storage	Training / Model Use	Human Review Conditions
ChatGPT Free/Plus	• Indefinite by default; user-deleted + 30d purge • Temporary Chat: not stored	• ON by default • Opt-out: Settings > Data Controls • Temporary Chat excluded from training	• Yes — staff may review flagged content and annotate snippets
ChatGPT	• 30-day abuse-monitoring baseline	• OFF by default (contractual)	• Highly limited — eligible orgs can request modified

Business	• Configurable retention for qualifying orgs • ZDR available — documented primarily for API; verify scope for your plan	• Limited to security, abuse monitoring, legal compliance only	abuse monitoring (no human review)
Claude Personal	• Up to 5 years if training enabled (default since Sep 28, 2025) • 30 days if training opted out • Existing pipeline data not retroactively removed	• ON by default since Sep 2025 • Opt-out: Settings > Privacy > "Help improve Claude"	• Yes — when flagged by safety classifiers or feedback submitted
Claude Team	• 30 days standard retention • Policy violations: up to 2 years • No ZDR on product UI	• OFF by default (commercial terms) • Training prohibition is primary default commitment	• Policy violations / safety flagging only; not for training
Gemini in Workspace	• 18 months default; admin-configurable (90 days to indefinite) • History disabled: 72-hour temp retention only	• NOT used for training outside org without permission • Explicit contractual commitment in CDPA	• Automated safety flagging only; not for quality or product improvement
M365 Copilot	• No automatic deletion; stored in tenant • Governed by Purview retention policies (admin-configured)	• NOT used to train LLMs (explicitly stated) • No benchmarking or human annotation use	• No abuse-monitoring-based human review — explicitly documented
NotebookLM Free	• Indefinite until manually deleted • Feedback data: 3 years, survives account deletion	• OFF for ordinary use • Feedback (thumbs up/down) triggers training use	• YES when feedback given — scope includes uploaded sources and queries
NotebookLM for Business	• Admin-configurable (Workspace policies) • History disabled: 72-hour temp retention only	• NOT used for training — explicit Workspace commitment • Applies even when feedback is given	• NO — explicit Workspace commitment, even with feedback
Perplexity Personal	• Retained "as long as reasonably necessary" • Account deletion: personal data deleted within 30 days	• ON by default; opt-out available • Queries may route to multiple third-party providers depending on plan and settings	• Not documented in official materials
Perplexity Enterprise Pro	• 30-day deletion on written request; otherwise "reasonably necessary" • Zero-retention must be explicitly requested in writing	• OFF — explicit contractual prohibition (§1.3.3) • Prohibition extends to all third-party subprocessors	• Not documented in official materials
Le Chat Free/Pro	• Conversations saved until user deletes • 30-day abuse-monitoring logs retained after deletion • User account deletion triggers full purge	• ON by default for Free and Pro (consumer) plans • Opt-out: account settings > Privacy controls • Privacy policy excludes only Enterprise + paid API from training	• Abuse monitoring and policy enforcement only • Not for model training or quality improvement — unlike OpenAI consumer plans
Le Chat Team	• Retained until admin or user deletes; no automatic purge • Admin data export and audit logs available	• Training status NOT explicitly confirmed as OFF at Team tier • Privacy policy exempts only "Enterprise and paid API" • DPA applies — Mistral acts as data processor • Verify with Mistral before use on client data	• Policy violations only; not for training per commercial terms
Lumo Free	• No server-side logs — conversations deleted after response generated • Optional encrypted chat history (zero-access encryption; only user can decrypt)	• Architecturally off — zero-access encryption means Proton cannot read conversations • Explicit policy + open-source code publicly verifiable	• Architecturally impossible • Zero-access encryption prevents any human access to conversation content • Not a policy promise — a technical architectural

	• Ghost Mode: conversation not saved in any form		guarantee
Lumo Professional	• No server-side logs; zero-access encrypted saved history • Ghost Mode available for maximum privacy • Proton Drive integration also zero-access encrypted	• Architecturally off — same zero-access architecture as Free • Open-source code verifiable; SOC 2 Type II audit (Jul 2025) validates controls	• Architecturally impossible at all plan tiers • SOC 2 Type II audit (Schellman, Jul 2025) independently validates zero-access architecture
Lumo + Workspace Premium	• No server-side logs for AI queries • Zero-access or E2E encryption across all Proton services (Mail, Drive, Calendar, VPN, Pass) • Ghost Mode available	• Same zero-access architecture — training architecturally off across full suite • ISO 27001 (May 2024) + SOC 2 Type II (Jul 2025) cover full Proton infrastructure	• Architecturally impossible across all Proton services • Zero-access encryption is uniform across the full suite

Tool Suitability & Certifications

Use this to decide: Which tools to formally approve or restrict, and what rationale to document in your AI policy.

Key assumption: The tier labels (L1–L4) are baseline recommendations for professional or nonprofit organizations not subject to sector-specific regulations. Healthcare, legal, financial, and other regulated orgs typically need stricter criteria — what this map labels L2 may require L1-equivalent controls for regulated or privileged data in those contexts. Organizations with no client or regulated data may find more flexibility at L3. If your org has negotiated a custom DPA or BAA with a specific vendor, that may warrant reassigning that tool's tier.

Policy tiers: L1 = Approved (routine internal use, org plan) · L2 = Approved (org commercial plan required; review for regulated use) · L3 = Review required before sensitive-document use · L4 = Not approved for privileged or regulated use without explicit authorization. Tier labels reflect tool-level protections only. Org policy should separately define approved task types and data categories — tool approval does not authorize use for regulated data, clinical or legal judgment, individual case decisions, or final determinations affecting people.

Service / Plan	Key Certifications / Compliance Signals	Policy Tier & Suitability
ChatGPT Free/Plus	• No certification scope confirmed for Free/Plus accounts • SOC 2, ISO 27001 certifications apply to OpenAI commercial products — Free/Plus accounts are not in scope	**POLICY TIER:** L4 — Not approved for privileged or regulated use without explicit authorization • Training on by default; no org controls; no DPA • Do not use for any org-sensitive content • Custom GPT actions can route content to external services without restriction
ChatGPT Business	• SOC 2 Type II, ISO 27001/27701/27017/27018 • No FedRAMP • HIPAA: requires dedicated product or API + ZDR + BAA	**POLICY TIER:** L3 — Review required before sensitive-document use • Training off contractually; admin controls available • Retention and ZDR controls vary by plan tier — confirm before approving for regulated content • Custom GPT/agent actions route externally; admin must explicitly restrict
Claude Personal	• SOC 2, ISO 27001, ISO 42001 • Consumer account — certifications apply at company level	**POLICY TIER:** L4 — Not approved for org use • Training on by default since Sep 2025; up to 5-year retention in pipelines • No admin controls; no DPA; personal account only
Claude Team	• SOC 2, ISO 27001, ISO 42001 • HIPAA-ready (API + BAA required)	**POLICY TIER:** L2 — Approved on org commercial plan • Training off contractually; 30-day retention; DPA automatic • No audit logs or EU data residency on Team plan — verify against your requirements • Suitable for most org use; EU-regulated orgs should confirm residency needs

Gemini in Workspace	• SOC 2, ISO 27001/17/18/27701/42001 • FedRAMP High • HIPAA BAA available	**POLICY TIER:** L1 — Approved for routine internal use (org Workspace plan) • Training off; CDPA governs; strongest certification set in this comparison • Primary risk: broadly shared Drive/Gmail files are immediately accessible to Gemini • Run permissions audit before broad rollout; apply IRM/DLP to sensitive files
M365 Copilot	• ISO 27001/18/42001; HIPAA (properly configured) • GDPR compliant • SOC 2 scope for Copilot not separately confirmed	**POLICY TIER:** L2 — Approved on org commercial plan; review Anthropic subprocessor • Training off; no abuse-monitoring human review documented • Anthropic added as subprocessor Jan 2026 — on by default in US/global; disable for sensitive workflows if needed • Copilot surfaces all M365 content user can view — permissions audit required before rollout
NotebookLM Free	• Inherits Google infra certifications • Not separately scoped for consumer tier	**POLICY TIER:** L4 — Not approved for org sensitive-document use • Feedback triggers training use and human review; 3-year feedback retention • No org controls; personal account only • See conditional-risk callout above Part 2
NotebookLM for Business	• Full Workspace: SOC 2, ISO 27001/17/18/27701/42001 • FedRAMP High, HIPAA BAA, BSI C5	**POLICY TIER:** L1 — Approved for routine internal use (Workspace plan) • Strongest training/human-review commitment: explicit contractual no-training, no-review even with feedback • Full Workspace certifications; CDPA governs • Verify admin controls are configured, not just available
Perplexity Personal	• SOC 2 Type 2, FedRAMP 20x Low • ISO 27001: not documented • HIPAA: gap assessment only	**POLICY TIER:** L4 — Not approved for org use • Training on by default; queries may route to multiple AI providers • No admin controls; no DPA; encryption not documented
Perplexity Enterprise Pro	• SOC 2 Type 2, FedRAMP 20x Low • ISO 27001: not documented • HIPAA: gap assessment only	**POLICY TIER:** L3 — Review required before sensitive-document use • Training prohibition contractual; DPA executed • Significant gaps: encryption standards, data residency, and admin controls not documented • Multi-provider inference routing continues — content may be sent to multiple providers for inference
Le Chat Free/Pro	• French company (Paris); GDPR-subject; EU law applies • ISO 27001 / SOC 2: not separately confirmed for consumer tier	**POLICY TIER:** L4 — Not approved for org use • Training on by default; no org controls; no DPA • EU/French jurisdiction is favorable but does not substitute for admin controls or a commercial plan
Le Chat Team	• French company (Paris); GDPR-subject • DPA available • ISO 27001 / SOC 2: not yet confirmed — verify with Mistral	**POLICY TIER:** L3 — Review required before sensitive-document use • Training status at Team tier unconfirmed — confirm before approving for client data • Admin controls present (API, audit logs, domain verification, data export) • EU/French jurisdiction; GDPR applies; suitable for general internal use pending training confirmation
Lumo Free	• Swiss/EU jurisdiction (Proton AG, Geneva) • ISO 27001 certified May 2024 • SOC 2 Type II attested Jul 2025 (Schellman) • GDPR + Swiss DPA compliant	**POLICY TIER:** L4 — Not for org deployment • No account required; no admin controls; no org audit capability • Strong privacy posture does not compensate for absence of org oversight • Individual staff use does not create shared data exposure, but cannot be managed or monitored
Lumo Professional	• Swiss/EU jurisdiction (Proton AG, Geneva) • ISO 27001 certified May 2024	**POLICY TIER:** L1 — Approved for routine org use • Strongest privacy architecture of any tool in this map: zero-access encryption makes training and human review architecturally impossible — not just policy-prohibited

	• SOC 2 Type II attested Jul 2025 (Schellman) • GDPR + Swiss DPA compliant; HIPAA support available; DPA available	• ISO 27001 + SOC 2 Type II independently validate admin controls • EU/Swiss jurisdiction; open-source code verifiable; HIPAA support for qualifying workflows
Lumo + Workspace Premium	• Same Proton AG certifications: ISO 27001 (May 2024) + SOC 2 Type II (Jul 2025) • Covers full Proton infrastructure — Mail, Drive, Calendar, VPN, Pass, Lumo • GDPR + Swiss DPA; HIPAA support available	**POLICY TIER:** L1 — Approved for routine org use • Full Proton suite under same zero-access architecture as Lumo Professional • Broader integration scope — adds encrypted Mail, Drive, Calendar, VPN, Pass • EU/Swiss jurisdiction; ISO 27001 + SOC 2 Type II cover full infrastructure

Sources

All ratings are based on official vendor documentation accessed April–May 2026. Sources listed below are sufficient to search and verify the key claims in this document.

ChatGPT Free/Plus · ChatGPT Business — OpenAI Privacy Policy; OpenAI Terms of Use; OpenAI Help Center (data controls; chat and file retention; temporary chat); OpenAI Security & Privacy documentation; OpenAI Subprocessor List; OpenAI Business Terms; OpenAI Enterprise Privacy

Claude Personal · Claude Team — Anthropic Privacy Policy; Anthropic Commercial Terms of Service; Anthropic Data Processing Addendum; Anthropic Trust Center (trust.anthropic.com); Anthropic Help Center (usage policies; privacy settings; data retention)

Gemini in Workspace — Google Workspace Privacy Notice; Google Cloud Data Processing Addendum; Google Cloud Compliance and Certifications; Google Admin Help (Gemini for Google Workspace settings); Google Cloud Terms of Service

M365 Copilot — Microsoft Privacy Statement; Microsoft Product Terms; Microsoft Trust Center (microsoft.com/trust-center); Microsoft Online Services DPA; Microsoft Copilot for Microsoft 365 documentation; Microsoft Purview Compliance documentation; Microsoft Security documentation (Copilot)

NotebookLM Free · NotebookLM for Business — Google Privacy Policy; Google Workspace Terms of Service; Google Cloud CDPA; NotebookLM Help Center; NotebookLM FAQ (training and data use); Google Blog (NotebookLM for Business announcement)

Perplexity Personal · Perplexity Enterprise Pro — Perplexity Privacy Policy; Perplexity Terms of Service; Perplexity Enterprise Terms; Perplexity Enterprise Security Page; Perplexity Help Center; Perplexity Blog (FedRAMP 20x, enterprise announcements)

Le Chat (all plans) — Mistral AI Privacy Policy (effective April 8, 2026); Mistral AI Commercial Terms of Service; Mistral AI Data Processing Addendum; Mistral AI Pricing (mistral.ai/pricing); Mistral AI Legal Center (legal.mistral.ai)

Lumo (all plans) — Proton Lumo for Business page (proton.me/business/lumo); Proton SOC 2 Type II blog post (proton.me/blog/soc-2); Proton Privacy Policy; Proton for Business overview; Lumo security model documentation