ChatGPT Free

Safety Quick Card · **AlignIQ**

Policy Tier: **L4** — Not approved for org use without explicit authorization

YOUR DEFAULT STATE

Training: ON — conversations may be used to train OpenAI models by default

Retention: Conversations are saved until you delete them; once deleted, purged from OpenAI systems within 30 days (legal holds may extend this in some cases)

Human review: Possible — OpenAI staff may review conversations for safety and quality

Admin controls: None — personal account only

DO THESE NOW (about 2 minutes)

- ☐ chatgpt.com > click your profile (top right) > Settings > Data controls > turn OFF "Improve the model for everyone"
- ☐ Optional: Settings > Data controls > turn OFF "Chat History & Training" to stop saving new conversations (also disables history — existing chats purged within 30 days)
- ☐ For reduced-risk sessions: use Temporary Chat (pencil icon in sidebar > "Temporary Chat") — excluded from training and not saved in history; note that OpenAI still retains these for up to 30 days before system deletion

WHAT CHANGES AFTER YOU DO THIS

Your conversations will no longer be used to train OpenAI models. Turning off Chat History stops new conversations from being saved and they are purged within 30 days. Temporary Chat is excluded from training and not visible in history — meaningfully lower risk, but not zero-retention.

WHAT STAYS RISKY REGARDLESS

- ! No admin enforcement — settings are per-person and can be re-enabled at any time
- ! Human review remains possible regardless of training opt-out
- ! No audit log, DPA (Data Processing Agreement — legal contract governing how your data is handled), or compliance documentation at this tier
- ! Do not assume conversations are protected by attorney-client privilege or work product doctrine — no court has recognized such protection for AI chat logs

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Client or partner data (contracts, case details, identifying info), internal work product (strategy, unreleased research, pre-publication drafts), sensitive personal or regulated data (SSNs, health records, financials). Training is ON by default.

YELLOW — USE WITH GUARDRAILS: General work (internal comms, meeting notes, non-confidential drafts) after opting out of training OR using Temporary Chat

GREEN — GENERALLY FINE: Public or already-published content (published articles, public reports, general research questions)

Sources:

OpenAI: How your data is used

OpenAI: Data Controls FAQ

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: aligniq.ai/resources

ChatGPT Plus

Safety Quick Card · **AlignIQ**

Policy Tier: **L4** — Not approved for org use without explicit authorization

YOUR DEFAULT STATE

Training: ON by default — same as Free; opt-out available in Data Controls

Retention: Conversations saved until you delete them; deleted conversations purged within 30 days. History ON by default.

Temporary Chat: Available — excluded from training and not saved in history; OpenAI still retains temporarily for up to 30 days before system deletion (not zero-retention)

Admin controls: None — personal subscription

DO THESE NOW (about 2 minutes)

- ☐ chatgpt.com > profile (top right) > Settings > Data controls > turn OFF "Improve the model for everyone"
- ☐ For reduced-risk sessions: use Temporary Chat (pencil icon > "Temporary Chat") — excluded from training and history; OpenAI still retains for up to 30 days before system deletion
- ☐ Review Memory: Settings > Personalization > Memory > disable or review what has been retained across sessions

WHAT CHANGES AFTER YOU DO THIS

Training opted out. Temporary Chat sessions are excluded from training and not saved in history; OpenAI retains them for up to 30 days before system deletion — reduced risk but not zero-retention.

WHAT STAYS RISKY REGARDLESS

- ! Still a personal plan — no organizational admin oversight or policy enforcement
- ! Memory feature (if enabled) stores info across all sessions — review what has been retained
- ! No compliance documentation or DPA (Data Processing Agreement — legal contract governing how your data is handled) at this tier
- ! Do not assume conversations are protected by attorney-client privilege or work product doctrine — no court has recognized such protection for AI chat logs

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Client or partner data (contracts, case details, identifying info), internal work product (strategy, unreleased research, pre-publication drafts), sensitive personal or regulated data (SSNs, health records, financials). Training ON; Memory stores content across sessions.

YELLOW — USE WITH GUARDRAILS: General work (internal comms, meeting notes, non-confidential drafts) after opting out of training OR using Temporary Chat. Review Memory settings separately.

GREEN — GENERALLY FINE: Public or already-published content (published articles, public reports, general research questions)

Sources:

OpenAI: Data Controls FAQ

OpenAI: New ways to manage your data

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: aligniq.ai/resources

ChatGPT Business

Safety Quick Card · **AlignIQ**

Policy Tier: **L3** — Review required before sensitive-document use

Applies to: ChatGPT Business workspace (previously called ChatGPT Team — rename standardized by OpenAI). For ChatGPT Enterprise or Edu plans, see your contract terms — those tiers have additional data controls, and apps are disabled by default (the reverse of Business).

YOUR DEFAULT STATE

Training: OFF by default — Business workspace conversations are not used to train OpenAI models

Retention: Conversations, files, and canvas documents are retained indefinitely unless explicitly removed by a user or admin. Admin-configurable retention policies are an Enterprise/Edu capability — not available at Business tier.

Human review: Not for training; OpenAI may review for policy or safety violations

Admin controls: YES — workspace admin dashboard at platform.openai.com; apps (plugins/GPTs) are enabled by default at Business tier — review and restrict these immediately

DO THESE NOW (about 2 minutes)

- ☐ Admin: platform.openai.com > Business settings > confirm "Model training" is OFF (verify after any plan changes)
- ☐ Admin: review and restrict which apps (plugins, GPTs) are enabled — Business tier has apps ON by default, unlike Enterprise/Edu where apps default to disabled
- ☐ Admin: establish a manual deletion practice — data is retained indefinitely with no auto-purge at this tier; delete sensitive threads at project close and conduct a quarterly history review
- ☐ All users: confirm you are signed into the Business workspace (not personal ChatGPT) before starting work
- ☐ Admin: audit any custom GPTs or connected apps — each has its own data policy separate from the Business agreement

WHAT CHANGES AFTER YOU DO THIS

Training is already off by default. Admin can verify settings and enforce policy across all workspace members.

WHAT STAYS RISKY REGARDLESS

- ! Conversations are retained indefinitely unless explicitly removed — training off does not mean data is deleted; Business tier has no admin-configured auto-purge; deletion is manual
- ! Apps (plugins/GPTs) are enabled by default at Business tier — unlike Enterprise where they default to disabled; review these before staff start using them
- ! Custom GPTs and connected apps may have separate data policies not covered by the Business agreement
- ! Staff using personal chatgpt.com accounts outside the Business workspace have no organizational protection
- ! Do not assume conversations are protected by attorney-client privilege or work product doctrine — no court has recognized such protection for AI chat logs

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Sensitive personal or regulated data (SSNs, health records, financial records). Data retained indefinitely with no automatic purge.

YELLOW — USE WITH GUARDRAILS: Client or partner data (contracts, case details, identifying info) with care. Training off and admin controls active, but data is retained indefinitely. Verify DPA before using for client work.

GREEN — GENERALLY FINE: Internal work product (strategy, research, pre-publication drafts), general work, public content. Commercial plan with training off by default.

Sources:

OpenAI: Data Controls FAQ ·
OpenAI: ChatGPT data & privacy ·
OpenAI: New ways to manage your data

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: alignq.ai/resources

Claude Personal (Free / Pro / Max)

Safety Quick Card · **AlignIQ**

Policy Tier: **L4** — Not approved for org use without explicit authorization

Note: Anthropic updated their training policy on September 28, 2025 — training now defaults to ON for all personal accounts. Opt-out is required. Claude Pro and Max now include connector/integration capabilities that extend data scope beyond the base chat.

YOUR DEFAULT STATE

Training: ON by default since September 28, 2025 — all personal accounts; must opt out explicitly
Retention if opted in: Up to 5 years (data retained for potential training use)
Retention if opted out: Conversations deleted within 30 days once opt-out is active
Human review: Possible — Anthropic may review flagged conversations
Admin controls: None — personal account only; note that Claude Pro and Max include connector capabilities (Google Drive, GitHub, etc.) that operate under personal account terms without org oversight

DO THESE NOW (about 2 minutes)

- ☐ claude.ai > click initials/name (bottom left) > Settings > Privacy > turn OFF "Help improve Claude"
- ☐ Verify the setting saved — the toggle should show as off
- ☐ Review Projects: check that no sensitive documents are stored in any Project's memory or context
- ☐ If using Claude Pro/Max: Settings > Integrations — review any connected apps; each connector extends data reach to third-party services under personal-account terms

WHAT CHANGES AFTER YOU DO THIS

Your conversations will be excluded from training data. Retention drops from up to 5 years to 30 days once opted out.

WHAT STAYS RISKY REGARDLESS

- ! If you haven't opted out yet — data shared since September 2025 may already be in the 5-year retention window
- ! No admin enforcement — individual opt-out can be re-enabled; no org-level enforcement at this tier
- ! Claude Pro/Max connectors (Google Drive, GitHub, etc.) extend data scope to third-party services and run under personal account terms without organizational oversight
- ! No SOC 2 (independent security audit certification) or DPA (Data Processing Agreement — legal data-handling contract) at personal tier
- ! Do not assume conversations are protected by attorney-client privilege or work product doctrine — no court has recognized such protection for AI chat logs

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Client or partner data (contracts, case details, identifying info), internal work product (strategy, unreleased research, pre-publication drafts), sensitive personal or regulated data (SSNs, health records, financials). Training has been ON by default since September 2025.

YELLOW — USE WITH GUARDRAILS: General work (internal comms, meeting notes, non-confidential drafts) after opting out in Settings → Privacy

GREEN — GENERALLY FINE: Public or already-published content (published articles, public reports, general research questions)

Sources:
Anthropic: September 2025 policy update
Anthropic: Is my data used for training?
Anthropic: Privacy Policy

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.
Latest updates: aligniq.ai/resources

Claude Team

Safety Quick Card · **AlignIQ**

Policy Tier: **L2** — Approved on org commercial plan

YOUR DEFAULT STATE

Training: OFF — Team is a commercial plan; Anthropic's Sep 2025 consumer policy changes do not apply

Retention: Conversations persist until deleted by the user or admin; deleted conversations are removed from backend servers within 30 days. Admin-configurable custom retention is an Enterprise capability — not available at Team tier.

Human review: Not for training; possible for safety or policy violations

Admin controls: YES — admin console with usage policies and member management

DO THESE NOW (about 2 minutes)

- ☐ Admin: claude.ai > click org name (top left) > Admin settings > Privacy & Data > confirm "Model training" is OFF
- ☐ Admin: establish a regular deletion practice — conversations are not auto-purged; they persist until explicitly removed
- ☐ Admin: define and communicate approved use cases and data classification rules to all team members
- ☐ All users: sign into the Team workspace — verify your org name appears top left, not personal claude.ai

WHAT CHANGES AFTER YOU DO THIS

Training is already off by default for commercial plans. Admin console provides oversight and policy enforcement across members.

WHAT STAYS RISKY REGARDLESS

- ! Conversations remain on Anthropic's infrastructure until explicitly deleted — this is a longer storage window than commonly assumed; plan a deletion cadence
- ! Staff using personal claude.ai accounts (Free/Pro/Max) are subject to the opt-in training policy — the most common gap to close
- ! No zero-retention option or custom retention configuration at Team tier (those are Enterprise features)
- ! Do not assume conversations are protected by attorney-client privilege or work product doctrine — no court has recognized such protection for AI chat logs

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Sensitive personal or regulated data (SSNs, health records, financial records). Best practice caution regardless of plan protections.

YELLOW — USE WITH GUARDRAILS: Client or partner data (contracts, case details, identifying info). DPA available, training off, admin-enforced. You're still transmitting third-party data outside your org.

GREEN — GENERALLY FINE: Internal work product (strategy, research, pre-publication drafts), general work, public content. Commercial plan protections apply.

Sources:

Anthropic: September 2025 policy update · Anthropic: Claude for Teams · Anthropic: Privacy Policy

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: aligniq.ai/resources

Gemini in Google Workspace

Safety Quick Card · **AlignIQ**

Policy Tier: **L1** — Approved for routine org use

Applies to: Google Workspace Business Starter, Standard, or Plus — including Google for Nonprofits. Workspace account protections now cover both Workspace apps (Docs, Gmail, Meet) and the Gemini app when signed in with a qualifying Workspace account.

YOUR DEFAULT STATE

Training: OFF — Google does not use Workspace customer data to train AI models (contractual guarantee for all tiers)

Retention — Workspace apps: Prompts and responses within Workspace apps (Docs, Gmail, Meet) are not retained after the session ends

Retention — Gemini app: When signed in with a Workspace account, conversation history is admin-controlled: configure 3, 18, or 36 months, or disable with up to 72-hour retention for service delivery

Human review: Not for AI training; possible for Terms of Service violations only

Admin controls: YES — Google Workspace Admin Console at admin.google.com

Key distinction: Workspace account protections apply to both Workspace apps AND the Gemini app on a Workspace account — NOT to personal Google accounts or personal Gmail through gemini.google.com

DO THESE NOW (about 2 minutes)

- ☐ Admin: admin.google.com > Apps > Google Workspace > Gemini for Workspace > Settings > verify Gemini is scoped to your org domain
- ☐ Admin: configure Gemini app conversation history retention (3 / 18 / 36 months, or off) to match your compliance requirements
- ☐ Admin + staff: access Gemini while signed into the org Workspace account — this protection now covers both Workspace apps and the Gemini app; NOT covered when using personal Gmail
- ☐ Admin: review Google Drive sharing permissions — Gemini surfaces files users can already access, including broadly shared ones

WHAT CHANGES AFTER YOU DO THIS

You are already protected by default when using Gemini through a qualifying Workspace account — this covers both Workspace apps and the Gemini app. Key action: ensure staff use their org account, not personal Gmail, and configure Gemini app retention appropriately.

WHAT STAYS RISKY REGARDLESS

- ! No. 1 risk: staff using personal Google accounts or personal Gmail on gemini.google.com — Workspace protections do not apply; those conversations may train Google's models
- ! Gemini surfaces broadly shared Drive/Gmail content — permissions hygiene across the org matters
- ! Google for Nonprofits eligibility requires annual verification — confirm your org's status is current

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Sensitive files with broad org access (SSNs, health records, client data). Gemini can surface any file your account can access, including ones you didn't intend to include. Restrict permissions before enabling.

YELLOW — USE WITH GUARDRAILS: Client or partner data (contracts, case details, identifying info). Workspace DPA applies and training is off. Confirm file-level permissions and retention settings.

GREEN — GENERALLY FINE: Internal work product (strategy, research, drafts), general productivity (meeting notes, email, documents). This is what the tool is built for.

Sources:

Google Workspace: AI privacy & security ·
Google Support: Generative AI privacy hub ·
Google: Gemini for Nonprofits

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: alignq.ai/resources

Microsoft 365 Copilot (Business)

Safety Quick Card · **AlignIQ**

Policy Tier: **L2** — Approved on org commercial plan

Applies to: M365 Business Basic, Standard, or Premium with the Microsoft 365 Copilot add-on license.

YOUR DEFAULT STATE

Training: OFF — Microsoft does not use M365 customer data to train AI foundation models (contractual guarantee)

Retention: Governed by your Microsoft Purview retention policies — there is no single universal default across all workloads; configure per workload and compliance requirement before assuming any specific number

Human review: Not for AI training; possible for compliance and security violations

Data boundary: Prompts and responses remain within your M365 service boundary; web search queries may involve Bing. Note: As of January 7, 2026, Anthropic became a subprocessor for some M365 Copilot AI capabilities; Anthropic models are currently out of scope for the EU Data Boundary.

Admin controls: YES — M365 Admin Center + Microsoft Purview for sensitivity labels and audit

DO THESE NOW (about 2 minutes)

- ☐ Admin: admin.microsoft.com > Settings > Microsoft 365 Copilot > review enabled features and user/group assignments
- ☐ Admin: BEFORE broad rollout, tighten SharePoint/OneDrive/Teams permissions — Copilot reads whatever users can access
- ☐ Admin: configure Microsoft Purview sensitivity labels (e.g., Confidential, Internal, Public) for key file types — Copilot respects these labels
- ☐ Admin: review third-party connectors (Settings > Integrations) — each connector extends data reach

WHAT CHANGES AFTER YOU DO THIS

Already protected by contract at this tier. The primary work is data classification and permissions hygiene to define what Copilot can surface.

WHAT STAYS RISKY REGARDLESS

- ! Permissions sprawl: Copilot reads any file the user can access — broad permissions = broad AI exposure
- ! Copilot summarizes meetings, emails, and documents — staff need to understand what it is reading
- ! Third-party connectors (if enabled) extend data reach; web search queries may route through Bing

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Sensitive files with broad permissions (SSNs, client health records, financial data). Copilot reads what users can access. Apply sensitivity labels and restrict permissions before enabling.

YELLOW — USE WITH GUARDRAILS: Client or partner data (contracts, case details, identifying info). M365 DPA applies and training is off. Verify file permissions and EU data boundary applicability.

GREEN — GENERALLY FINE: Internal work product (strategy, research, documents), general productivity (email, meetings, drafting). This is what the tool is built for.

Sources:

Microsoft: M365 Copilot privacy ·
Microsoft: Enterprise data protection ·
Microsoft: Copilot data protection architecture

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: alignq.ai/resources

NotebookLM Free

Safety Quick Card · **AlignIQ**

Policy Tier: **L4** — Not approved for org use without explicit authorization

Applies to: notebooklm.google.com with a personal Google account. Designed for research and source analysis. Does not train on your data — but contains one critical confidentiality risk to understand before uploading non-public content.

YOUR DEFAULT STATE

Training: OFF — Google does not use your sources, queries, or responses to train NotebookLM
Retention: Notebooks and sources stored until you delete them; associated with your personal Google account
Human review: Possible only if you click a feedback button (thumbs up/down) — the critical risk at this tier
Admin controls: None — personal Google account only

DO THESE NOW (about 2 minutes)

- ☐ If uploading any non-public content: never click the thumbs up or thumbs down feedback buttons — this is the sole trigger that opens your content to human review at Google
- ☐ Use only publicly available sources and documents you would be comfortable sharing if feedback were accidentally triggered
- ☐ For confidential work: consider Google Workspace with NotebookLM for Business — that tier removes the human review risk entirely
- ☐ Check for policy changes: Google's terms can update; review the NotebookLM privacy page when notified of changes

WHAT CHANGES AFTER YOU DO THIS

Avoiding feedback buttons ensures no human reviewer at Google can access your uploaded sources, queries, or responses. The free tier already protects you from AI training — confidentiality depends on not triggering that one review pathway.

WHAT STAYS RISKY REGARDLESS

- ! Human review is one accidental click away — a thumbs up or thumbs down instantly opens your full session to Google reviewers; there is no undo
- ! No admin controls — any protection depends entirely on individual user discipline; cannot be enforced across a team
- ! No DPA (Data Processing Agreement) or organizational compliance documentation at this tier
- ! Personal Google account terms apply — not the stronger Workspace protections

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER UPLOAD: Client or partner data (contracts, case details, identifying info), sensitive personal or regulated data (SSNs, health records, financials). No DPA. Human review is possible if feedback features are used.

YELLOW — USE WITH GUARDRAILS: Internal work product (strategy, research, pre-publication drafts). No training on uploads. Avoid feedback buttons (thumbs up/down) to prevent human review exposure.

GREEN — GENERALLY FINE: Public or already-published content (published articles, reports, source documents). This is what the tool is built for.

Sources:
NotebookLM: Privacy and data protection · NotebookLM: Help Center · Google: Privacy Policy · Latest updates: aligniq.ai/resources

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.
Latest updates: aligniq.ai/resources

NotebookLM for Business

Safety Quick Card · **AlignIQ**

Policy Tier: **L1** — Approved for routine org use

Applies to: NotebookLM for Business (requires a qualifying Google Workspace account). Provides significantly stronger protections than the free tier — confidentiality is contractual, not just a matter of avoiding feedback buttons.

YOUR DEFAULT STATE

Training: OFF — Google contractually does not use your organization's data to train NotebookLM

Retention: Governed by your Google Workspace account policies; admin-configurable; stays within your org

Human review: OFF — even when providing feedback, Google does not expose content to human reviewers at this tier

Admin controls: YES — managed through Google Workspace Admin Console at admin.google.com

Data boundary: Data stays private to you and whoever you share the notebook with — Google's explicit contractual commitment

DO THESE NOW (about 2 minutes)

- ☐ Admin: confirm all staff access NotebookLM while signed into the org Workspace account — personal Google accounts bypass all these protections
- ☐ Admin: admin.google.com → verify NotebookLM access is scoped to your organization domain
- ☐ Admin: review notebook sharing settings — treat shared notebooks like any sensitive Workspace document; sharing extends access to all recipients
- ☐ Admin: run a permissions audit on Google Drive — broadly shared files accessible to users may also surface within their NotebookLM context

WHAT CHANGES AFTER YOU DO THIS

Confidentiality protection is now contractual, not dependent on individual behavior. Human reviewers cannot access content at this tier regardless of what staff click. Admin controls let you manage access scope and verify org-wide compliance.

WHAT STAYS RISKY REGARDLESS

- ! Staff using free personal NotebookLM (personal Gmail) are not covered by these protections — require and communicate org Workspace account use
- ! Sharing settings determine who can see a notebook — treat shared notebooks with the same care as any confidential Workspace document
- ! Workspace policy terms can change; verify protections when Google announces updates to Workspace or NotebookLM agreements

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER UPLOAD: Sensitive personal or regulated data (SSNs, health records, financial records). Best practice caution regardless of contractual protections.

YELLOW — USE WITH GUARDRAILS: Client or partner data (contracts, case details, identifying info). DPA in place, no training, no human review. You're still transmitting third-party data outside your org.

GREEN — GENERALLY FINE: Internal work product (strategy, research, pre-publication drafts), general work, public content. Contractual protections and admin controls apply.

Sources:

NotebookLM for Business: Privacy ·
Google Workspace: AI privacy & security ·
NotebookLM: Help Center

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: aligniq.ai/resources

Perplexity Free

Safety Quick Card · **AlignIQ**

Policy Tier: **L4** — Not approved for org use without explicit authorization

YOUR DEFAULT STATE

Training: Perplexity's own data use ON by default — queries may be used to improve Perplexity's service; opt-out available in Settings

Retention: Queries logged; linked to account if signed in; session-only if browsing without an account

Third-party models: Routes queries to OpenAI and Anthropic models; Perplexity's agreements with these providers prohibit them from using your data for their own model training

Dual risk surface: Search engine + AI assistant — both search logs and AI conversation logs apply simultaneously

Admin controls: None — personal account only

DO THESE NOW (about 2 minutes)

- ☐ perplexity.ai > profile icon (top right) > Settings > AI Data > turn OFF "Allow my data to be used to improve Perplexity"
- ☐ For sensitive research: consider browsing without signing in — reduces data linkage to your profile
- ☐ Know the dual risk: your query goes to Perplexity's search infrastructure AND to a third-party AI model simultaneously

WHAT CHANGES AFTER YOU DO THIS

Your data will no longer be used to improve Perplexity's service. Unauthenticated use reduces data-to-profile linkage.

WHAT STAYS RISKY REGARDLESS

- ! Queries route to OpenAI and Anthropic models as subprocessors — Perplexity's contracts prohibit those providers from training on your data, but you remain subject to Perplexity's own data practices
- ! Search history creates a richer data profile than AI-only tools (research intent + AI queries combined)
- ! No DPA (Data Processing Agreement — legal contract for organizational data handling) or admin controls at any personal tier

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Client or partner data (client names, case details, partner info), internal work product (strategy, unreleased research), sensitive personal or regulated data (SSNs, health, financials). Training is ON. Search queries are logged separately from chat.

YELLOW — USE WITH GUARDRAILS: General research (background topics, public industry info) without org-specific identifiers. Prefer unauthenticated. Exclude client names and confidential context from queries.

GREEN — GENERALLY FINE: Public content, general research with no org-specific data (topic research, public information lookup). This is what the tool is built for.

Sources:

Perplexity: Privacy Policy ·
Perplexity Help: Data collection ·
Perplexity: Enterprise Pro data practices

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: aligniq.ai/resources

Perplexity Pro / Max

Safety Quick Card · **AlignIQ**

Policy Tier: **L4** — Not approved for org use without explicit authorization

YOUR DEFAULT STATE

Training: Perplexity's own data use ON by default — same opt-out as Free; Pro/Max adds more AI model access but does not change privacy defaults

Third-party models: Routes to OpenAI, Anthropic, or others; model selection available in Pro/Max. Perplexity's contracts prohibit these providers from training on your Perplexity data — the primary risk is Perplexity's own practices, not the model providers.

Dual risk surface: Same search + AI dual exposure as Free tier

Admin controls: None — personal subscription. Perplexity Enterprise Pro has org controls; this tier does not.

DO THESE NOW (about 2 minutes)

- ☐ perplexity.ai > Settings > AI Data > turn OFF "Allow my data to be used to improve Perplexity"
- ☐ Model selection is available in Pro/Max — Perplexity's contracts already prohibit all model providers from training on your queries; the key risk is Perplexity's own data practices regardless of model choice
- ☐ Evaluate: for organizational use, Perplexity Enterprise Pro provides a DPA and configurable zero-retention — personal Pro/Max does not

WHAT CHANGES AFTER YOU DO THIS

Data use for improvement opted out. Pro/Max model selection gives flexibility in AI model choice, though all models are already contractually prohibited from training on your Perplexity data.

WHAT STAYS RISKY REGARDLESS

- ! Still a personal plan — no DPA (Data Processing Agreement), no admin enforcement
- ! Dual search + AI risk surface persists regardless of subscription level
- ! Personal Pro/Max does not provide the contractual protections that Perplexity Enterprise Pro gives for organizational use

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Client or partner data (client names, case details, partner info), internal work product (strategy, unreleased research), sensitive personal or regulated data (SSNs, health, financials). Training is ON by default.

YELLOW — USE WITH GUARDRAILS: General research (background topics, public industry info) without org-specific identifiers. Opt out of training in Settings first.

GREEN — GENERALLY FINE: Public content, general research with no org-specific data. Well-suited for personal professional use.

Sources:

Perplexity: Privacy Policy ·
Perplexity Help: Enterprise data retention ·
Perplexity: Enterprise Pro security

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: aligniq.ai/resources

Perplexity Enterprise Pro

Safety Quick Card · **AlignIQ**

Policy Tier: **L3** — Review required before sensitive-document use

Applies to: Perplexity Enterprise Pro (custom pricing). Requires a signed contract and DPA with Perplexity. Meaningfully different data posture from personal Free/Pro/Max accounts.

YOUR DEFAULT STATE

Training: OFF — Perplexity contractually prohibits using Enterprise data to train AI models. Third-party model providers (OpenAI, Anthropic, others) are bound by annual agreements prohibiting them from training on Enterprise customer data.

Retention — Conversations: Configurable: zero data retention is available as an option and must be explicitly enabled via the admin console. Not zero by default.

Retention — Files: 7-day default retention; organizations with 50+ seats can configure custom retention periods and force-delete at any time

Admin controls: YES — enterprise admin console with SSO/SAML, usage policies, and member management

Compliance: SOC 2 Type II certified; DPA (Data Processing Agreement — legal data-handling contract) executed with your organization

DO THESE NOW (about 2 minutes)

- ☐ Admin: execute DPA with your legal team before rolling out to staff — this is the key document establishing data protections
- ☐ Admin: set up SSO/SAML at enterprise.perplexity.ai to centralize access control and prevent personal account mixing
- ☐ Admin: explicitly configure conversation data retention in the admin console — zero data retention is available but must be actively enabled, it is not the default
- ☐ Admin + staff: communicate and enforce that all org-related use of Perplexity must be through Enterprise Pro accounts — personal accounts bypass all these protections

WHAT CHANGES AFTER YOU DO THIS

Training is off and contractually protected. DPA creates legal accountability. Admin controls let you enforce policy, monitor usage, and configure zero-retention for conversations — but zero-retention must be actively enabled.

WHAT STAYS RISKY REGARDLESS

- ! Zero-retention for conversations must be explicitly configured — do not assume it is active by default
- ! Multi-vendor architecture persists — queries still route through third-party AI providers, though contractual protections apply to all of them
- ! Staff using personal Perplexity accounts (Free/Pro/Max) get zero enterprise protection — require org policy and communicate it explicitly

DATA CLASSIFICATION — THESE RULES STILL APPLY

RED — NEVER PASTE: Sensitive personal or regulated data (SSNs, health records, financial records, legal matters). Best practice caution even with zero retention configured.

YELLOW — USE WITH GUARDRAILS: Client or partner data (contracts, case details, partner info). DPA in place, zero retention if configured. Search queries still leave your org boundary.

GREEN — GENERALLY FINE: General research (competitive intelligence, topic research), public content. Works well here when zero retention is active.

Sources:

Perplexity: Enterprise Pro security ·
Perplexity Help: Enterprise data retention ·
Perplexity: Data Processing Addendum

AlignIQ

AI-assisted research starting point — not final advice. Policies change; always verify with official vendor documentation before acting. Best practice depends on your org's context.

Latest updates: alignq.ai/resources