

MEMORANDUM

Vendor Risk Assessment: Anthropic Claude (Team plan) for a Solo or Small Law Firm

TO: Solo and small-firm practitioners evaluating Claude for client work
FROM: Mike Bittner, General Counsel, MB Legal, PLLC
DATE: July 4, 2026
RE: Risk assessment of Anthropic's Claude Team plan (Team chat, Claude Code, and Claude Cowork) for a firm handling confidential and privileged client information

I. Service overview and scope

Anthropic, PBC is a San Francisco AI company that builds Claude, a large language model available through web, desktop, and mobile applications and a developer API. As of this assessment, Anthropic is the most highly valued AI company in the market. Its May 2026 [Series H round](#) raised roughly \$65 billion at a \$965 billion post-money valuation, it has raised on the order of \$125 billion in total, and it filed confidentially for an initial public offering in June 2026. Amazon and Google are its largest strategic investors. The company publishes its safety research, its Responsible Scaling Policy, and its security posture at trust.anthropic.com.

This assessment evaluates the Claude Team plan, the paid subscription built for groups, as it would be used by a solo practitioner or a small firm. Three Anthropic products run under a single Team subscription and a single contractual framework: (1) Claude Team, the web and desktop chat interface, for drafting, contract review, research, and general work; (2) Claude Code, a command-line tool for developers, less common in a law office but available; and (3) Claude Cowork, a desktop tool that works with files in a folder the user selects on their own computer. Cowork reached general availability across all paid plans, on macOS and Windows, on April 9, 2026. All three are governed by the same Commercial Terms of Service and the same Data Processing Addendum. There is no separate contract for Code or Cowork.

All three run under the same base commercial framework: the Commercial Terms of Service, the incorporated Data Processing Addendum, and the Service Specific Terms, which carry supplemental terms for particular services. There is no separately negotiated contract for Code or Cowork. What differs across the products is the data they can reach. Team chat processes only what the user types into a prompt. Claude Code processes the contents of the working directory it runs in. Cowork is an agent that can read, create, and modify files in a folder the user mounts, and it can use connectors to reach services like email, calendar, and cloud drives. In a law practice a mounted folder can include client files, drafts, and privileged analysis, and a connector can expose far more than the document the lawyer meant to share. The contractual protections are identical across all three, but the exposure runs low to high from chat to Code to Cowork, and that gradient is managed by the firm's own use policy, not by the contract.

Data footprint for a law firm. A typical small firm will put confidential client information into Claude: client and adverse-party names, contact details, contract terms, matter facts, and attorney work product. This information is subject to the duty of confidentiality (ABA Model Rule 1.6 and its state analogs), the duty of competence including technology competence (Model Rule 1.1, Comment 8), and the duty of communication (Model Rule 1.4). Rule 1.6(c) requires a lawyer to make reasonable efforts to prevent unauthorized disclosure of client information, which is the standard this assessment measures the vendor against. This is a confidentiality-and-privilege analysis, not a children's-data or sector-specific one.

Scope and use. This memo assesses the Team tier under Anthropic's Commercial Terms of Service (effective June 17, 2025), the incorporated Data Processing Addendum (effective February 24, 2025), and the incorporated Service Specific Terms (effective June 8, 2026), together with the Usage Policy and Supported Regions Policy that the Commercial Terms bring in by reference. It is written as a general reference for the small-firm market and is current as of the date above. It is not legal advice for any particular firm, and it does not create an attorney-client relationship. Terms and plan features change; confirm the current versions at trust.anthropic.com and claude.com before relying on any provision. Any firm with a mandatory client or regulatory requirement (see Section VII) should read the Team-versus-Enterprise discussion closely.

II. Documents assessed

This assessment rests on Anthropic's current public contractual and security documents for commercial (paid) customers.

Document	What it governs	Risk	Notes
Commercial Terms of Service (eff. June 17, 2025)	Confidentiality, input/output ownership, the no-training commitment, IP indemnity, liability cap, suspension, term and termination.	Low	Click-through; strong customer-side terms.
Data Processing Addendum (eff. Feb. 24, 2025)	Processor obligations, sub-processors, cross-border transfers (SCCs, UK and Swiss addenda), security schedule, 48-hour breach notice, deletion.	Low	Incorporated into the Commercial Terms; no separate signature.
Service Specific Terms (eff. June 8, 2026)	Supplemental terms per service, including the Development Partner Program (opt-in data/training) and Covered Models (retention and safety review overriding ZDR).	Medium	Read alongside the base terms; two settings a firm should understand.
Privacy Policy	Data handling, controller designations for EEA/UK and other regions, data-subject rights.	Low	GDPR-aligned rights framework.
Usage Policy (AUP)	Prohibited uses; human-review and disclosure expectations for high-risk uses, including legal advice.	Low	Aligns with a lawyer's supervision duty.
Trust Center / certifications	SOC 2 Type II, ISO 27001:2022, ISO 42001:2023, CSA STAR, NIST 800-171; reports available under NDA.	Low	trust.anthropic.com
Sub-processor list	Current sub-processors and infrastructure providers.	Low	Published; changes on notice with an objection window.

III. Risk analysis

The analysis below tracks five risk categories. Each provision reference points to the substance of Anthropic's public terms; verify current numbering before relying on a specific citation.

A. Data privacy and confidentiality

Confidentiality. The Commercial Terms (Section E) include mutual confidentiality obligations, and they classify Customer Content as the customer's Confidential Information. The recipient must protect it with no less than reasonable care and use it only to perform under the agreement, and must destroy it on request, subject to routine backup copies that stay under the confidentiality obligation. Two ordinary limits apply:

Anthropic may disclose confidential information where required by law or court order (with notice to the customer where not prohibited), and backup copies persist until purged. This is standard enterprise-SaaS language and is a low contractual concern for client information, subject to the use controls in Section VII.

Data privacy compliance. The DPA makes Anthropic a data processor acting on the customer's documented instructions. Anthropic commits not to sell or share customer personal data, not to use it outside the direct business relationship, and not to combine it with data from other sources. The DPA incorporates the EU Standard Contractual Clauses (Modules Two and Three), a UK addendum, and a Swiss addendum for cross-border transfers, and Anthropic will support a transfer impact assessment on request. The Privacy Policy provides a GDPR-style rights framework. For a firm with clients or matters touching the EU or UK, the transfer package is in place. One drafting note: the DPA's SCC annex lists special categories of personal data as 'None,' meaning the customer has not designated sensitive-category processing; a firm whose matters involve sensitive personal data should be aware of that default.

Retention and deletion. Retention follows a two-layer model on the Team plan. Conversations persist in the product until the user deletes them; on deletion, the content is removed from history and purged from back-end storage within 30 days. On termination of the account, the DPA (Section H) requires Anthropic to delete customer data within 30 days, with a copy returned or exported on request first. Deletion exceptions are narrow (legal requirement, an active dispute, and abuse enforcement).

Four points qualify that clean picture, and a firm should understand them. First, the Team plan does not offer configurable retention or a true zero-data-retention setting for the chat interface; that is an Enterprise and API feature. Second, the Service Specific Terms let Anthropic designate certain models as 'Covered Models,' for which it may retain and run safety reviews on inputs, outputs, and usage data, and that expressly overrides any zero-data-retention commitment that would otherwise apply. Third, if a conversation is flagged by Anthropic's automated trust-and-safety systems as violating the Usage Policy, the inputs and outputs are kept for up to two years and the associated trust-and-safety classification scores for up to seven years, regardless of user deletion. Fourth, data tied to feedback a user submits (a thumbs up or down, or a bug report) is retained for five years. For most small-firm work, user-controlled deletion plus the 30-day purge is acceptable, paired with a habit of deleting conversations that carry sensitive matter facts once the work is done and an awareness that the safety-flagging and feedback tails sit outside that control.

B. AI-specific risks

Training on your data. This is the question most lawyers ask first, and the answer on the commercial plan is clean. The Commercial Terms (Section B) state that Anthropic 'may not train models on Customer Content from Services.' That is a contractual prohibition by default, not an opt-out you have to remember to toggle. Two qualifications, both worth knowing. First, the Service Specific Terms create an opt-in 'Development Partner Mode'; if a firm affirmatively enables that setting, Anthropic may use its submitted content to train models. Leave it off. Second, feedback a user chooses to submit can be used by Anthropic. Neither changes the default no-training posture, but a firm should know the one switch that would. The consumer tiers (Free, Pro, Max) work differently and can use inputs for training unless the user opts out, which is why client work belongs on a commercial plan (Team or Enterprise), not a personal account.

Ownership of prompts and outputs. The customer keeps all rights in inputs, and Anthropic assigns to the customer whatever rights it may have in the outputs. There is no license back to Anthropic to use your outputs commercially. This is the strongest ownership position available among the major AI vendors. The 'if any' phrasing on outputs reflects the open question of whether AI-generated text is copyrightable at all, not a reservation of rights by the vendor.

Indemnification for AI output. Anthropic defends customers against third-party intellectual-property claims arising from paid, authorized use of the Services, which covers the core scenario where an output infringes someone's IP. The defense has real exclusions (Section K.3): it does not apply to claims arising from the customer's own inputs, from modifications the customer makes, from combining outputs with non-Anthropic technology or content, from use the customer knew or should have known was infringing, from the practice of a patented invention in an output, or from trademark use of an output in commerce. For a law firm those exclusions are manageable but worth understanding. Separately, no AI vendor indemnifies for the accuracy of outputs; the terms place responsibility on the user to review outputs, which for a lawyer is the same duty that already applies to a drafting tool or a junior's work.

C. Security and infrastructure

Certifications. Anthropic holds SOC 2 Type II, ISO 27001:2022 (information security), ISO/IEC 42001:2023 (AI management systems), CSA STAR, and NIST 800-171 attestations, with reports available under NDA through the [Trust Center](#). This is a strong, independently audited posture and, notably, matches what Anthropic offers on the Enterprise plan on the certifications that matter for confidential legal work. The one certification gated to Enterprise is HIPAA (see Section III.D).

Technical controls and breach notice. The DPA's security schedule specifies AES-256 encryption at rest, TLS 1.2 or higher in transit, multi-factor authentication on systems that access customer data, role-based least-privilege access, endpoint detection and response, centralized security monitoring, and annual third-party penetration testing (summaries available on request). Anthropic commits to notify the customer in writing within 48 hours of becoming aware of a security breach, faster than the GDPR's 72-hour controller clock and well inside the 30-day outer limit in Florida's breach statute. The notice must describe the categories and approximate volume of affected data and the response taken.

D. Legal and liability

Indemnity and liability cap. Indemnification is mutual (Section K): Anthropic defends third-party IP claims from authorized use, and the customer covers claims arising from its own inputs or policy violations. General liability is capped at fees paid in the prior 12 months and excludes consequential and indirect damages (Section L.3). One point in the customer's favor: the cap does not apply to either party's indemnification obligations (Section L.3.b), so the IP defense is not squeezed by the fee cap. On a per-seat plan the general cap is still small in dollar terms, which is the trade-off for click-through pricing, and there is no separate carve-out from the cap for a data breach or a confidentiality breach on the Team tier. A firm that needs a negotiated cap or those carve-outs is looking at Enterprise. Disputes go to informal resolution, then binding arbitration (JAMS in California for US customers; UNCITRAL in Dublin for EEA, Swiss, and UK customers), with a jury-trial and class-action waiver, under California or Irish law respectively.

Professional and sector rules. For the practice of law, the Usage Policy's expectation of human review and disclosure for high-risk uses lines up with the lawyer's supervision and candor duties, and ABA Formal Opinion 512 (2024) and parallel state bar guidance (including the Florida Bar's) frame the duties directly: understand the tool's terms and data handling, verify outputs, and get the client's informed consent before putting confidential information into a generative-AI tool where disclosure is a risk. Treat client disclosure as a real analysis, not an afterthought. Three sector points matter. First, HIPAA: if a firm handles protected health information (for example, a plaintiff's-side personal-injury or a healthcare-regulatory practice), a Business Associate Agreement is available only for Anthropic's HIPAA-ready services (its first-party API and Enterprise with HIPAA activated), it is sales-assisted rather than self-serve, and it does not cover Cowork. The Team and consumer tiers sit outside the BAA, so PHI should not go into Team-tier Claude. Second, payment-card data (PCI) should not be entered into any tier. Third, a firm should not bill a client

for the time spent learning a general-purpose AI tool, and should disclose in advance any AI cost it intends to pass through (Formal Opinion 512). These are controls for the firm, not gaps in the vendor.

Audit and transparency. The DPA gives customers access to annual third-party audit reports and certifications on request, plus a customer-initiated audit right subject to reasonable scope and frequency limits. Current certifications are published at trust.anthropic.com. This is a workable transparency framework for a small firm, which will rely on the published reports rather than conducting its own audit.

E. Operational resilience

Availability. The Team plan does not carry a contractual uptime SLA with service credits. Anthropic publishes real-time status at status.anthropic.com. The Services are provided 'as is' and 'as available' (Section L.2), and Anthropic reserves broad rights to suspend access for security, policy, legal, cost, or upstream-vendor reasons, with no liability for losses from a suspension (Section I.3). For a law practice, Claude is a productivity tool, not case-management or docketing infrastructure, so an outage is usually an inconvenience rather than a client emergency. The residual risk is real but modest: a firm that uses Claude for time-sensitive drafting should keep a fallback and not build a filing deadline around its availability. A firm that wants a written uptime SLA is looking at Enterprise.

Termination and portability. Either party may terminate for convenience on 30 days' notice, with a 30-day cure period for material breach. On termination the DPA provides for return or self-service export of data and deletion within 30 days. The terms do not specify an export file format, a minor gap for a chat product.

Vendor stability. Anthropic is one of the best-capitalized companies in the sector, with a roughly \$965 billion valuation as of May 2026, reported run-rate revenue in the tens of billions, and a confidential IPO filing on file. Like most companies at this stage of growth it is not yet profitable, but the capital position and strategic backing from Amazon and Google make near-term vendor failure a remote risk. As with the other frontier AI companies, Anthropic is a party to intellectual-property litigation (including copyright suits by rightsholders); none of that is known to bear on the security or confidentiality posture assessed here, and a firm that wants a current litigation picture should run its own search before relying on this point.

IV. How this compares to the cloud tools a firm already uses

Most small firms already keep confidential client files in third-party cloud services: iCloud, Google Workspace, or Microsoft 365. None of those providers train on business-customer data under their commercial terms, all offer DPAs with SCC-based transfer mechanisms and breach-notification duties, and all hold SOC 2 Type II. On the pure contract terms, Claude's Team plan sits at the same bar. The useful question is not whether Claude is safe in the abstract, but whether it is a bigger step than the cloud storage a firm has already accepted. It is a step, and a narrow one, for three reasons.

Processing, not just storage. When a document sits in iCloud or Google Drive it is encrypted at rest, inert, and no one reads it during normal operations. When a document goes to Claude it is loaded into memory and actively processed by the model during inference. The no-training commitment means it is not retained to improve the model, but during the request itself the data is being computed on, not merely held. That is a different operation with a different threat model.

A human-review pathway exists. If Anthropic's automated systems flag a conversation for a policy review, staff may review that content (disassociated from user identity). Passive file storage has no equivalent pathway. The exception is narrow and standard across AI vendors, but it is an access route that

does not exist when a file only sits in a drive, and it is worth understanding before putting the most sensitive matters through the tool.

Output generation is a new category. Drive storage does not generate new content from stored files; Claude does. Conversation isolation and the no-training term address the obvious concern, but generation is a risk category passive storage does not have.

Bottom line. For a firm already comfortable storing confidential files in iCloud, Google Workspace, or Microsoft 365, using Team chat is a modest step up in risk, not a dramatic one: the contract protections are on par, and the incremental risks (processing, the safety-review pathway, and output generation) are manageable through the controls in Section VII. Cowork and connectors are a larger step, because there the model does not just process a pasted prompt, it can read, change, and act on local files and reach into connected accounts. Treat agentic use as its own decision with its own guardrails, not as an extension of file storage.

V. Key issues

Issue	Why it matters for a firm	Source	Assessment
Use paid, not consumer, accounts	The no-training protection lives in the Commercial Terms. Consumer tiers (Free, Pro, Max) can train on inputs unless the user opts out. Client work belongs on Team or Enterprise.	Commercial Terms; Consumer Terms	Review Needed
No BAA on Team (HIPAA)	A BAA covers only HIPAA-ready services (first-party API and Enterprise with HIPAA activated), is sales-assisted, and does not cover Cowork. Keep PHI out of Team-tier Claude.	HIPAA-ready docs; BAA page	Review Needed
Cowork and connectors are the high-risk surface	Cowork can read, change, and act on mounted files and reach connected accounts (email, calendar, drives). Mount only the working set, keep the most sensitive matters out, and enable connectors only when needed.	Service Specific Terms; Cowork docs	Review Needed
Covered Models / retention overrides ZDR	For designated Covered Models, Anthropic may retain and safety-review inputs and outputs, overriding any ZDR commitment. Team also lacks configurable retention (Enterprise/API only). Manage with deletion hygiene.	Service Specific Terms F; DPA	Review Needed
Set the opt-out switches at signup	Leave Development Partner Mode OFF (enabling it lets Anthropic train on your content), and opt out of the publicity right so Anthropic does not list the firm as a customer.	Service Specific Terms E; CT Section G	Review Needed
Liability cap is small	General liability is capped at trailing 12-month fees; the cap does not apply to indemnification (so IP defense is uncapped), but there is no data-breach or confidentiality carve-out on the click-through tier.	Commercial Terms L.3	Acceptable
Thinner admin controls than Enterprise	Team provides SSO and basic admin controls. SAML for any IdP, SCIM, audit logs, a Compliance API, customer-managed keys, and US-only inference are Enterprise. Build offboarding into a checklist.	Team vs. Enterprise plan docs	Acceptable

VI. Decision matrix

Weights are adjusted from the defaults because this is an AI vendor: AI-specific risks are raised to 20 percent, and legal/liability and operational resilience are each reduced to 15 percent. Data privacy and security stay at 25 percent each, reflecting the centrality of the confidentiality duty for a law firm. No children's-data adjustment applies to this generic profile.

Factor	Weight	Score (1-5)	Weighted	Notes
Data privacy and confidentiality	25%	4.5	1.13	Processor DPA, no-sell, deletion; connectors and lawful process temper it.
AI-specific risks	20%	4.0	0.80	No-training default and output ownership, less Covered Models retention and Dev Partner Mode.
Security and infrastructure	25%	4.75	1.19	Full cert stack; 48-hr notice.
Legal and liability	15%	3.5	0.53	Small cap, indemnity exclusions, arbitration, no HIPAA on Team.
Operational resilience	15%	3.5	0.53	No SLA, broad suspension rights; strong vendor stability.
TOTAL	100%		4.16 / 5.00	Approve with conditions.

Overall score: 4.16 / 5.00. The contractual framework (Commercial Terms, DPA, and Service Specific Terms) gives a small firm strong privacy protections, an audited security posture, a 48-hour breach clock, customer audit rights, and a clear post-termination data lifecycle. The score sits in the low-4 range rather than higher because of the small general liability cap, the indemnity exclusions and broad suspension rights, the Covered Models retention override, the agentic exposure of Cowork and connectors, and the governance features reserved for Enterprise (HIPAA/BAA, configurable retention, SAML/SCIM, audit logs, a Compliance API, a written SLA). None of those is disqualifying for ordinary confidential work, and none is something most small firms need on day one, but together they place this at a conditional approval rather than a clean one.

VII. Recommendation

Verdict: approve the Team plan for confidential client work, with conditions. For a solo or small firm handling ordinary confidential and privileged information, Claude Team under the Commercial Terms and DPA is a sound choice. Start on Team, and move to Enterprise the day a specific requirement makes it necessary, rather than paying for Enterprise against a need you cannot yet see.

Move to Enterprise (sales-assisted) when one of these becomes real: a client engagement or regulation that requires EU data residency, US-only inference, or zero data retention; a HIPAA/BAA obligation because the firm handles PHI; a need for SAML SSO across any identity provider, SCIM provisioning, audit logs, a Compliance API, or customer-managed encryption keys; configurable retention; or negotiated contract terms and a written SLA. Absent one of those, Team is the right call.

Conditions and internal safeguards, regardless of what the vendor commits to:

1. **Keep client work on a paid plan.** Use Team (or Enterprise), never a personal Free, Pro, or Max account. The no-training protection depends on it.
2. **Adopt a short AI use policy.** Cover which matters and data types may go into Claude, the rule against entering PHI or payment-card data, mandatory human review of anything client-facing, and deletion of conversations carrying sensitive matter facts once the work is done.
3. **Be deliberate with Cowork and connectors.** Mount only the folder needed for the task at hand, not an entire client archive; keep the most sensitive files (sealed matters, settlement strategy, PHI) out of the mounted set; and enable connectors to email, calendar, or cloud drives only when a task needs them, since a connector can expose far more than the document you meant to share.
4. **Treat outputs as a first draft.** Review and verify every citation, clause, and factual assertion before it goes to a client or a court. This is the existing competence and supervision duty, applied to a new tool.
5. **Get informed consent where it is required.** Under ABA Formal Opinion 512 and parallel state bar guidance, evaluate the disclosure risk before putting client information into Claude, obtain the client's informed consent where the facts call for it, and add engagement-letter language addressing AI use. Do not treat disclosure as optional boilerplate. Also follow Formal Opinion 512 on fees: do not bill clients for time spent learning the tool, and disclose any pass-through AI cost in advance.
6. **Set the account switches at signup.** Leave Development Partner Mode off (turning it on lets Anthropic train on your content), and submit the publicity opt-out so Anthropic does not list the firm as a customer.
7. **Manage access and offboarding.** Use the Team admin controls to add and remove users promptly. Team's admin tooling is thinner than Enterprise (no SCIM or audit logs), so build offboarding into a checklist rather than relying on directory automation.
8. **File the vendor documents.** Pull the current SOC 2 Type II report and the sub-processor list from the Trust Center and keep them with the firm's records, and note the 15-day sub-processor objection window so a change does not slip past.

Negotiation outlook. Low. The Team terms are click-through and are not individually negotiated. That is acceptable here, because the standard terms already deliver the protections that matter (no training, 48-hour breach notice, output ownership, and customer audit rights), which many negotiated enterprise agreements struggle to reach.

VIII. Sources

Source	Use
Anthropic Commercial Terms of Service (eff. June 17, 2025)	Confidentiality, ownership, no-training commitment, indemnity and exclusions, liability cap, suspension, publicity, arbitration, term.
Anthropic Data Processing Addendum (eff. Feb. 24, 2025)	Processor obligations, transfers, security schedule, 48-hour breach notice, deletion.
Anthropic Service Specific Terms (eff. June 8, 2026)	Development Partner Mode (opt-in training) and Covered Models (retention and safety review overriding ZDR).
Anthropic Trust Center (trust.anthropic.com)	Certifications: SOC 2 Type II, ISO 27001:2022, ISO 42001:2023, CSA STAR, NIST 800-171.

Source	Use
Anthropic certifications (Privacy Center); HIPAA-ready Enterprise plans	Team vs. Enterprise features, HIPAA/BAA scope, retention, SSO, certifications.
Anthropic Series H funding announcement (May 2026)	Valuation, funding, and vendor-stability facts.
ABA Model Rules 1.1, 1.4, 1.6(c) and ABA Formal Opinion 512 (2024)	Competence, communication, confidentiality, informed consent, supervision, and fee duties for generative-AI use.

This assessment is general information current as of July 4, 2026, prepared by MB Legal, PLLC. It is not legal advice and does not create an attorney-client relationship. Vendor terms and certifications change; confirm current terms at trust.anthropic.com and claude.com before relying on any provision.