FUSION COMPUTING FIELD GUIDE

Claude Cowork for financial services firms: secure client and document work

Mike Pearlstein, CISSP, MSc Computer Science (AI)

Published May 27, 2026; Updated August 12, 2026

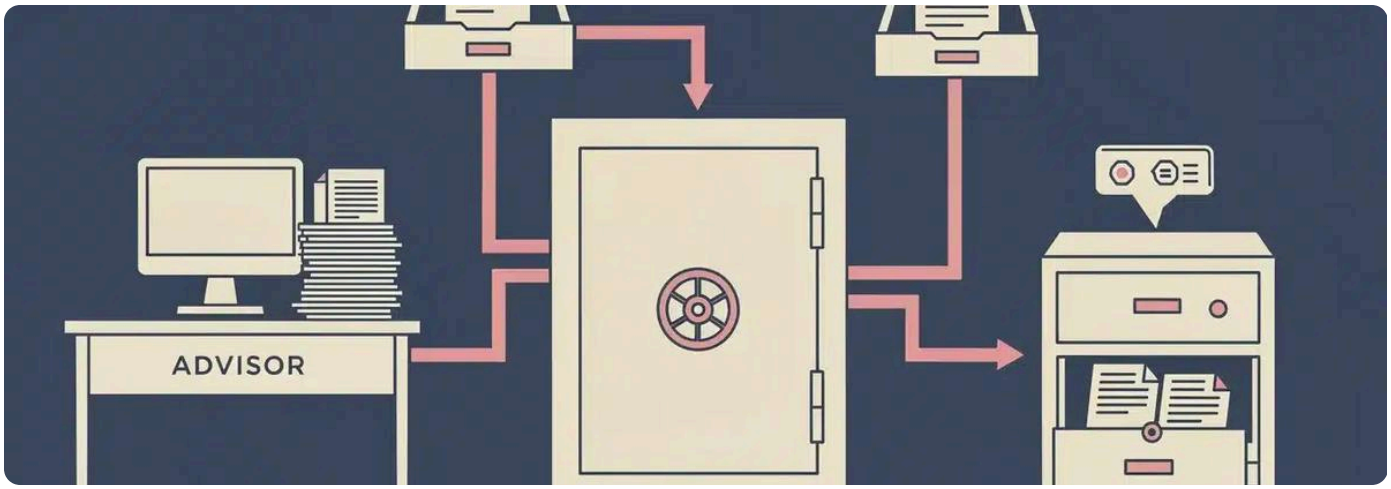

WHAT'S INSIDE

0. Frequently Asked Questions

TRUSTED BY **Toronto law firms** **Hamilton manufacturers** **Vancouver clinics** **GTA accounting firms** **Ontario non-profits** **British Columbia professional services**

Financial services firms want to know whether Claude Cowork can take on application and document work without exposing client financial data. According to Statistics Canada (2026), 40.4% of finance and insurance businesses used AI in the 12 months to the second quarter of 2026, against 19.2% across all industries. The duty to protect client information stays with the firm.

Mike Pearlstein, CISSP, MSc Computer Science (AI), founder of Fusion Computing, which has secured IT for Canadian financial services firms across Toronto, Hamilton, and Metro Vancouver since 2012.

> *The firms that get this right do not buy a different tool than everyone else. They scope it to one client file, decide in advance whether sessions may run in the vendor cloud, and keep compliance on anything a client or a regulator will ever see.*
>
> **Mike Pearlstein, CISSP, founder of Fusion Computing**

Key takeaways

Book a Consultation

- Claude Team runs from 2 to 150 members at US$20 per seat per month billed annually, so a 6-person compliance team is a real starting point.
- Scope it to one client file, never the whole CRM or policy system.
- Local Cowork sessions sit outside your audit trail. Sessions run in the vendor cloud are captured in the Compliance API, and on Enterprise that cloud mode is off until an owner turns it on.
- CIRO bulletin 26-0034 says examiners will ask about AI use and review the controls behind it.

Can financial services firms use Claude Cowork with client data?

Yes, on a Team or Enterprise plan with access scoped to one client file and a written policy. According to Anthropic's commercial terms (2026), "Anthropic may not train models on Customer Content from Services." Your PIPEDA duties and your sector record obligations survive that promise, and no vendor setting retires them.

This guide covers the broader financial firm: brokerages, MGAs, lenders, and fintech operations. A registered investment advisory firm should also work through our wealth-management guide for CIRO duties, and a CPA practice our accounting guide for CPA duties.

It's the same secure-adoption logic from the pillar guide on using Claude Cowork securely in your business, applied to a financial firm, and it sits alongside our broader IT for financial services work.

What does Claude Cowork actually do in a financial back office?

> According to Anthropic's release notes (2026), Cowork reached general availability on macOS and Windows on April 9, 2026. It completes multi-step document work instead of answering one question. In a back office that means onboarding synthesis, client correspondence, disclosure upkeep, claim summaries, and compliance packages. Every output is a draft a person verifies.

Here's how those jobs map to the work, with the guardrail that protects the client. Fusion Computing walks firms through this before any pilot, the same way we scope any AI services engagement.

Book a 30-minute call to scope Claude Cowork for your firm safely →

Task	What to test	The guardrail
Application and onboarding synthesis	Reads onboarding documents and drafts a client summary	Scope to one client file; staff verify
Client communications	Drafts letters and follow-ups from records	Client identifiers stay in a scoped folder
Policy and disclosure organization	Sorts and updates policies and disclosures	Internal documents, reviewed before issue
Claims or file summaries	Summarizes a claim or file from source documents	A draft for review, not the official record
Compliance-document prep	Assembles compliance and reporting documents	One client file, reviewed by compliance

Treat these as pilot hypotheses to validate in your own tenant, not vendor-guaranteed capabilities. Anthropic documents a subset of them directly; the rest are workflows to test before a firm relies on them.

The client-data and compliance guardrails

> Least privilege is the whole control: scope Cowork to one client file, not the CRM. Classify what may go in, which is the working documents for the active file, and what stays out, which is account numbers beyond that folder. According to Anthropic's safety guidance (2026), prompt-injection risk is "non-zero" even with classifiers running.

The first thing I check is scope. When an Ontario brokerage connects the agent to its whole CRM, 1 task can read every client's file. Scope it to the active client and you've cut most of the exposure without changing the work at all.

> **FIELD NOTE FROM MIKE.** The first thing I change is access. I've watched an operations lead point an agent at a CRM holding every client's financial profile. We scoped it to one client file, and the workflow that felt reckless became routine. The work's identical; the exposure isn't.

Anthropic documents three approval modes, and the one a regulated firm wants is Manually approve. Deletion is protected in every mode: Cowork "requires your explicit permission before permanently deleting any files." That is a useful floor, though it is not a substitute for scoping.

The policy is the other half. A short rule set, the kind we cover in our guide on what belongs in an AI acceptable use policy, names the approved tool, the data that may go in, and who may run it. Fusion Computing pairs that with a cybersecurity review so the firm protects client trust and its regulatory standing.

What CIRO's 2026 report requires of dealers using AI

According to the CIRO Compliance Report for 2026 (2026), published as bulletin 26-0034 on February 17, 2026, s.1.5 states: "As part of our FinOps examination approach, we will be enquiring about the use of AI in dealers' operations." Examiners then review the operational controls behind that use.

REGULATOR QUOTE. CIRO, Compliance Report for 2026, section 1.5: "To the extent dealers are using AI, we will be reviewing the operational controls they implemented to ensure AI is working as designed." That passage adds that a dealer should assess whether AI or automation of a regulatory function is a material business change requiring advance written notice and a Form 33-109F5 filing.

Read that second sentence twice. In my reading, automating a supervisory or regulatory function at a CIRO dealer is not only a tooling decision, it can be a registration filing. Most firms I speak to have never connected those two ideas, and the fix costs nothing if you check before the pilot rather than after.

Clause 1.1 of the same report says CIRO will run another cybersecurity table-top exercise in 2026 aimed at small and mid-sized dealers. If your firm is not CIRO-regulated, the equivalent question comes from your own regulator. A firm supervised by the Financial Services Regulatory Authority of Ontario faces the same evidence burden under a different name.

The oversight gap for record-keeping and audits

> Where a Cowork session runs decides whether you can produce a record of it. According to Anthropic's admin guidance (2026), local session history is stored on the user's computer and "cannot be centrally managed or exported by admins." Sessions run in the cloud are captured in the Compliance API instead.

That split is the setting I check first on any engagement. Cloud sessions are on by default for Team plans and off by default for Enterprise, where an owner must enable them and then grant the capability through custom roles. A dealer that wants a retrievable record of AI-assisted work should be running in the cloud on purpose.

The Enterprise audit logs capture metadata rather than the work itself. Team and Enterprise owners can also stream Cowork events to a SIEM through OpenTelemetry, which Anthropic notes "doesn't replace audit logging for compliance purposes. Cowork exports the full text of user prompts by default, along with tool parameters, file paths and user email addresses, so configure filtering or redaction at the collector and set SIEM access and retention before enabling export." Our engineers found that stream is where tool calls and file access become visible.

If the pilot needs centralized monitoring, define the destination, filtering, access and retention before enabling OpenTelemetry.

Cowork versus the AI already inside your lending platform

> Fintech AI is embedded in a lending, advisory, or policy system and is scoped to that system's data. A general desktop agent like Cowork reaches across your own files instead. The practical difference is blast radius: a scoring model sees 1 dataset, while an agent sees whatever folder and connectors you granted it 5 minutes ago.

That difference cuts both ways. A platform vendor already carries part of your control burden and can usually hand you an audit trail. With Cowork the firm owns the scoping decision, so the upside is that the agent works on the odd, cross-system documents no platform ever covered.

> **CONTRARIAN THESIS.** Canadian data residency is a choice, not a legal requirement. PIPEDA Schedule 1, clause 4.1.3 says an organization stays responsible for information "transferred to a third party for processing" and "shall use contractual or other means to provide a comparable level of protection". Choose residency because it makes accountability cheap to evidence, not because a statute demands it.

How much does Claude Cowork cost for a financial firm?

> According to Anthropic's published pricing (2026), Claude Team is built for 2 to 150 members at US$20 per seat per month billed annually, or US$25 billed monthly. A premium seat with a larger usage allowance is US$100 annually. Enterprise is US$20 per seat plus usage charged at API rates.

Two numbers matter when I size this for a firm. The floor is 2 seats, so a compliance officer and one operations lead can pilot without a company-wide purchase. The ceiling is 150 members, above which the plan is Enterprise, which is also where groups and custom roles live.

In our practice the licence is rarely the expensive part. The real cost sits in scoping, policy drafting, and the monitoring wiring, which we scope at 3 to 6 hours for a Toronto brokerage rather than a per-seat charge. That is why I tell firms to price the control work first and the seats second.

Plan tier and a setup checklist for a financial firm

> Plan tier is the first decision, because only Team and Enterprise carry the contractual no-training commitment plus admin controls. From there a safe rollout is short: scope to one client file, keep Manually approve on, decide the cloud-session question, write a usage policy, turn on OpenTelemetry, and keep compliance signing off.

Need the policy first? Use our AI acceptable use policy template.

The contrast with a personal account is stark. According to Anthropic's privacy centre, an individual plan with the improvement setting enabled may retain chats "in a de-identified format for up to 5 years in our model training pipelines". Business inputs and outputs are deleted within 30 days by default. That gap is the whole argument for buying seats.

Get a CISSP-led review of where AI tools touch client financial data →

> **Why Canadian firms bring this work to Fusion Computing**
>
> CISSP-led, a Microsoft Solutions Partner and a CompTIA Managed Services Trustmark holder, securing IT for Canadian SMBs across Toronto, Hamilton, and Metro Vancouver since 2012.

1. **Choose Team or Enterprise.** Client financial data on a personal account is the first risk to fix.
2. **Scope to one client file.** Never the whole CRM or policy system. Widen only with a reason.
3. **Set the approval mode to Manually approve.** Deletion always prompts, but the other actions should too.
4. **Decide the cloud-session question.** Team has it on by default; Enterprise owners must switch it on deliberately.
5. **Write an acceptable use policy.** Name the approved tool, the data that may go in, and who may run it.

6. **Turn on OpenTelemetry monitoring only after deciding what may be logged.** On local sessions it's an additional visibility stream into what the agent did.
7. **Keep compliance signing off.** Nothing client-facing or regulator-facing ships without review.
8. **Check the material-change question.** A CIRO dealer automating a regulatory function may owe a Form 33-109F5 filing.

> **FIELD NOTE FROM MIKE.** I ask one question before any of the eight steps above: who currently reviews this document when a human writes it? If nobody can name that person, the AI pilot is not the problem I want to solve first. That question surfaces an informal review step in a good share of our clients, and Cowork simply makes the gap visible sooner.

None of it's exotic. The technical setup can be quick; the privacy, logging and approval work is what decides how long the pilot takes. Fusion Computing sets it up as part of the managed IT work we already do for firms, and the same pattern carries to wealth-management guide readers and to accounting guide readers under their own regulators.

Fusion Computing helps Canadian businesses across Toronto and the GTA, Hamilton, and Metro Vancouver with managed IT, cybersecurity, and Microsoft 365.

Where a financial firm should start this week

Buy 2 Team seats, scope them to one live client file, set approvals to Manually approve, and answer the cloud-session question in writing before anyone opens a second folder. If you want a CISSP-led second opinion before that first file goes near an agent, talk to us or read more about how we work.

Frequently Asked Questions

Q. Is Claude Cowork safe for client financial data?

Claude Cowork can be safe for client financial data on a Team or Enterprise plan, with access scoped to one client file and compliance reviewing the output. Team and Enterprise do not train on organization content by default, but that alone is not a confidentiality determination: verify the execution mode, Anthropic's current terms and DPA, retention, the applicable privacy law and your professional obligations before client data is used. The contract states that Anthropic may not train models on Customer Content, and business inputs and outputs are deleted within 30 days by default. Expose only the documents a task needs.

Q. What plan does a financial services firm need for Claude Cowork?

A financial firm should use Team or Enterprise, never an individual Pro or Max account. Team runs from 2 to 150 members at US$20 per seat per month billed annually. Above 150 members, or where you need groups and custom roles to enable Cowork for one department, the answer is Enterprise. Client data on a consumer account is the first risk to fix.

Q. Is our data used to train the model?

On Team and Enterprise plans, Anthropic's commercial terms say it may not train models on Customer Content from the Services. A personal Pro or Max account is different: with the improvement setting enabled, chats may be kept in de-identified form for up to 5 years in Anthropic's training pipelines. For a firm holding client financial data, that gap is the reason to buy business seats.

Q. How does this differ from the wealth-management and accounting guides? –

This guide covers the broader financial-services firm: brokerages, MGAs, lenders, and fintech operations. A registered investment advisory firm should read the wealth-management guide for CIRO duties, and a CPA or bookkeeping practice the accounting guide for its own rules. The Cowork setup is the same across all three; the regulator and the records differ.

Q. How is Claude Cowork different from fintech AI? –

Fintech AI is usually built into a lending, advisory, or policy platform and scoped to that system. Claude Cowork is a general desktop agent that works across your own files and apps, which suits application, communication, and compliance documents more than scoring or transactions. The practical differences are where the data lives and how broadly the agent can reach, which is the first thing our CISSP-led review pins down.

Want an AI use policy that fits PIPEDA and your sector rules? →

Q. Does Claude Cowork meet record-keeping rules? –

On its own it does not, and the answer depends on where the session runs. Local session history stays on the user's computer and cannot be centrally managed or exported by admins. Sessions run in Anthropic's cloud are captured in the Compliance API. A firm meets its record-keeping rule by choosing the cloud mode deliberately and streaming Cowork events to a SIEM through OpenTelemetry.

Q. Does Claude Cowork work on Windows, or only Mac? –

Both, and more since it launched. Cowork reached general availability on macOS and Windows through the Claude desktop app on April 9, 2026. On July 7, 2026 Anthropic added web at claude.ai and the Claude mobile apps, in beta for Team and Enterprise, with those sessions running on Anthropic's infrastructure rather than the laptop. Confirm the current feature list inside the app.

Q. Who at the firm should run Claude Cowork? –

Start with a small group in operations or compliance who understand client confidentiality. The Cowork toggle itself is organization-wide, so a Team plan is all-or-nothing. On Enterprise, groups and custom roles let an admin enable Cowork, or the cloud-session capability, for named teams only. Pair whichever you choose with training and a written policy, which is the CISSP-led review Fusion Computing runs before a pilot starts.

Talk to Fusion

Put this guide to work

Fusion Computing helps Canadian SMBs adopt AI safely. The review work is CISSP-led at a Microsoft Solutions Partner and CompTIA Managed Services Trustmark holder, securing Canadian businesses since 2012.

- Managed IT support with a 1-hour priority response
- Cybersecurity, managed detection and response, and compliance reviews
- AI and Copilot consulting, including safe Claude Cowork pilots
- Microsoft 365 management and licensing

Book a 30-minute scoping call

fusioncomputing.ca/contact-us · 416-566-2845
Toronto · Hamilton · Metro Vancouver