FUSION COMPUTING FIELD GUIDE

Claude Cowork for wealth management firms: secure client and portfolio work for Canadian advisors

Mike Pearlstein, CISSP, MSc Computer Science (AI)

Published May 26, 2026; Updated August 12, 2026

WHAT'S INSIDE

---

fusioncomputing.ca · 416-566-2845 · CISSP-led managed IT, cybersecurity, and AI for Canadian SMBs · Toronto · Hamilton · Metro Vancouver

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

Advisors want to know whether Claude Cowork can speed up client and portfolio work without putting account data or a CIRO record obligation at risk. 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 keep client information confidential stays with the firm.

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

> *A wealth-management firm does not need a different tool than everyone else. It needs the agent scoped to 1 household, a decision on record about whether sessions run in Anthropic's cloud, and a reviewer on anything a client or a CIRO examiner will ever read.*
>
> **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 3-person planning team can pilot without a firm-wide purchase.
- Scope Cowork to 1 household's folder, never the whole CRM or client-document store.
- Local sessions sit outside your audit trail. Cowork run through web and mobile is captured in the Compliance API, and on Enterprise that cloud mode is off until an owner turns it on.
- CIRO bulletin 26-0034, published February 17, 2026, says FinOps examiners will ask about AI use and review the operational controls behind it.
- An advisor or compliance reviews anything client-facing. Never put account numbers on a personal account.

Can wealth management firms use Claude Cowork with client and account data?

> Yes, on a Team or Enterprise plan, with access scoped to 1 household's folder and a written policy. According to Anthropic's commercial terms (2026), "Anthropic may not train models on Customer Content from Services". Your CIRO record-keeping and supervision duties, and PIPEDA, survive that promise.

Confidentiality and supervision are the firm's obligations under CIRO rules, and no vendor setting removes them. What a practice controls is the scope: which files the agent opens, which plan governs the data, and who reviews the output.

This spoke is the tactical companion to our deeper work on AI governance for wealth management firms. It applies the secure-adoption logic from the pillar guide on using Claude Cowork securely in your business, and it sits alongside our broader IT for wealth management firms work.

If your firm sits outside registered dealer work, the broader Claude Cowork guide for financial services firms covers brokerages, MGAs, lenders, and fintech operations under FSRAO and PIPEDA rather than CIRO.

What does Claude Cowork actually do for an advisory firm?

> According to Anthropic's release notes (2026), Cowork reached general availability on macOS and Windows on April 9, 2026, and gained web and mobile access on July 7, 2026. It completes multi-step document work instead of answering 1 question. In a wealth practice that means KYC synthesis, portfolio reporting, meeting notes, suitability prep, and compliance filing.

Here's how those 5 jobs map to the work, with the guardrail that keeps each one safe. Fusion Computing walks firms through this before any pilot.

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

Task	What to test	The guardrail that keeps it safe.
KYC and onboarding synthesis	Reads onboarding documents and drafts a client summary	Scope to 1 household; the advisor verifies.
Account and portfolio reporting	Assembles a draft report from statements and holdings	Figures checked before anything is sent.
Client-meeting note summaries	Turns raw notes into a structured summary and action list	A draft for the file, not the official record.
Suitability-review prep	Pulls KYC and holdings into a review-ready package	Compliance review before any recommendation.
Compliance-document organization	Sorts and renames disclosures and records	1 household folder, never the whole CRM.

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 supervision guardrails.

> Least privilege is the whole control: scope Cowork to 1 household's folder, not the CRM. According to Anthropic's safety guidance (2026), prompt-injection risk stays "non-zero" even with classifiers running. The folder you point the agent at is the real boundary, not the model.

The first thing I check is scope. When a Toronto practice connects the agent to its whole CRM, 1 task can read every client's holdings. Scope it to the active file and you have cut most of the exposure.

> **FIELD NOTE FROM MIKE.** The first thing I change is access. I've watched an advisor point an agent at a CRM export holding every household's positions. We scoped it to 1 folder, and the workflow that felt reckless became routine. The work's identical; the exposure isn't.

What may go in, and what must not.

- **In:** working documents for the 1 active household, meeting notes, and draft disclosures.
- **Out:** account numbers and identifiers for any household outside that folder.
- **Out:** the CRM export, the full client-document store, and shared network drives.
- **Always:** an advisor or compliance reviewer on anything a client will read.

Anthropic documents 3 approval modes, and the one a registered firm wants is Manually approve. Deletion is protected in every mode, since Cowork “requires your explicit permission before permanently deleting any files”. Treat that as a floor rather than a scoping control.

The policy is the other half. A short rule set names the approved tool, the data that may go in, and who may run it, and I keep it to 1 page for a CIRO dealer. We cover the shape of it in our guide on what belongs in an AI acceptable use policy, paired with a cybersecurity review so the firm has a defensible position with its regulator.

What does CIRO’s 2026 compliance report require of dealers using AI?

> According to the CIRO Compliance Report for 2026 (2026), s.1.5 says examiners “will be enquiring about the use of AI in dealers’ operations”. CIRO published it as bulletin 26-0034 on February 17, 2026. Examiners then review the operational controls behind that use, and the same clause raises a registration question most firms miss.

> **REGULATOR QUOTE.** CIRO, Compliance Report for 2026, s.1.5, says it will be reviewing “the operational controls they implemented to ensure AI is working as designed” wherever a dealer uses AI. The report also

> tells each member firm to assess whether automating a regulatory function is a material business change needing advance written notice and a Form 33-109F5 filing.

Read that second sentence twice. In my reading, automating a supervisory step at a CIRO dealer is not purely a tooling decision, because it can become a registration filing. Most advisors I speak to have never connected those 2 ideas, and the check costs nothing before a pilot.

Incident reporting and outsourcing.

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. Section 1.1 also points investment firms at IDPC Rule 3703, which mandates reporting of cybersecurity incidents meeting set criteria, and at Guidance Note GN-2300-21-003 on outsourcing arrangements.

That outsourcing note matters here. Pointing an agent at client files is a third-party engagement, and CIRO asks a registered firm to assess risk before, during, and after it. In our practice that framing turns an AI question into a vendor-management question the firm already knows how to answer.

> **CONTRARIAN THESIS.** A CIPF membership does not backstop an AI mistake. CIPF protects client property held by an insolvent member firm and states plainly that it “does not provide protection against any other type of risk or loss”. Advisors who assume the fund covers an operational or data incident are reading a solvency backstop as insurance.

The oversight gap that matters for CIRO 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 stays on the user’s computer and “cannot be centrally managed or exported by admins”. Cowork via mobile and web is captured in the Compliance API instead.

That split is the setting I check first on any advisory engagement. Cloud sessions are on by default for Team plans and off by default for Enterprise, where an owner switches them on and grants the capability through custom roles.

Where the record actually lives.

- **Local session:** history sits on the laptop, outside admin export and outside the Compliance API.
- **Web or mobile session:** runs on Anthropic infrastructure and is captured in the Compliance API.
- **OpenTelemetry stream:** tool calls and file access, in your own SIEM, for either mode.

The Enterprise audit logs capture metadata rather than the work itself. Owners can also stream Cowork events to a SIEM through OpenTelemetry, which Anthropic notes “doesn’t replace audit logging for compliance purposes” on its own. 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 wealthtech platform.

Wealthtech AI is embedded in a planning, CRM, or portfolio system and is scoped to that system's data. Claude Cowork reaches across your own files instead. The practical difference is blast radius: a rebalancing model sees 1 dataset, while a desktop agent sees whatever folder and connectors you granted it 5 minutes ago.

That difference cuts both ways, and the question I ask first is which system already hands you an audit trail. A wealthtech vendor usually carries part of that burden for you. 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.

How much does Claude Cowork cost an advisory 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 licence with a larger usage allowance is US$100 annually. Enterprise is US$20 per seat plus usage charged at API rates.

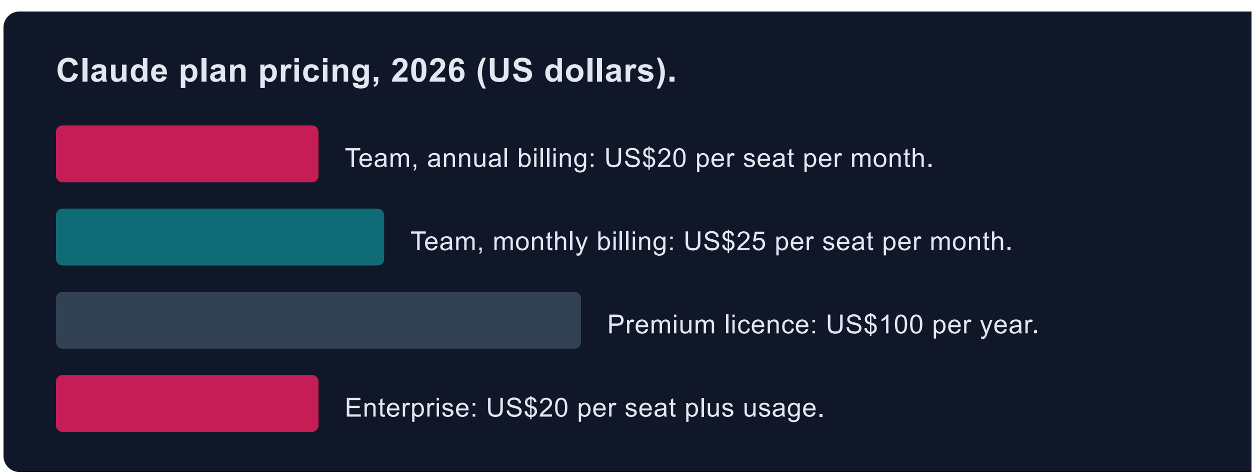

Two numbers matter when I size this for a practice. The floor is 2 licences, so a chief compliance officer and 1 senior advisor can pilot without a firm-wide purchase. The ceiling is 150 people, above which the plan is Enterprise.

In our experience 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 advisory firm rather than a per-seat charge.

Plan tier and a setup checklist for an advisory firm.

> According to Anthropic's privacy centre (2026), an individual plan with the improvement setting enabled may keep 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 5-year gap is the whole argument for buying 2 seats.

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

So plan tier is the first decision, because only Team and Enterprise carry the contractual no-training commitment plus the admin controls a firm needs.

> **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.

Get a CISSP-led review of how AI tools reach your client accounts →

The 8-step rollout.

1. **Choose Team or Enterprise.** An advisor running client work on a personal account is the first risk to fix.
2. **Scope to 1 household folder.** Never the whole CRM or client-document store. Widen only with a reason.

3. **Set the approval mode to Manually approve.** Deletion always prompts, and client data 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 an advisor or compliance signing off.** Nothing client-facing ships without review.
8. **Check the material-change question.** Automating a regulatory function may owe CIRO a Form 33-109F5 filing.

> **FIELD NOTE FROM MIKE.** I ask 1 question before any of the 8 steps above: who reviews this document today, when a human writes it? If nobody at the firm can name that person, the AI pilot is not the problem I want to solve first. Cowork simply makes an informal review step visible sooner than a regulator would.

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 across Toronto and the GTA. The same pattern carries to accounting firms and insurance brokerages under their own regulators.

Where an advisory firm should start this week.

Buy 2 Team seats, scope them to 1 live household folder, set approvals to Manually approve, and put the cloud-session decision in writing before anyone opens a second folder. If you want a CISSP-led second opinion first, talk to us or read more about how we work.

Frequently Asked Questions

Q. Is Claude Cowork safe for client accounts? –

Claude Cowork can be safe for client account data on a Team or Enterprise plan, with access scoped to 1 household's folder and an advisor 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. Anthropic's commercial terms say it 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. Does Claude Cowork meet CIRO 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. Cowork via mobile and web is captured in the Compliance API. A firm meets the rule by choosing the cloud mode deliberately and streaming Cowork events to a SIEM through OpenTelemetry.

Q. What plan does an advisory firm need for Claude Cowork? –

Team or Enterprise, never a personal Pro or Max account. Claude Team runs from 2 to 150 members at US$20 per seat per month billed annually. Above that ceiling, or where you need groups and custom roles to enable Cowork for 1 department only, the answer is Enterprise. An advisor running client work on a consumer account is the first risk to remediate.

Q. Is client PII 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 registered advisory firm, that 5-year gap is the reason to standardise on business seats.

Want an AI use policy that holds up to CIRO supervision? →

Q. What does CIRO expect from a dealer using AI? –

The CIRO Compliance Report for 2026, bulletin 26-0034 of February 17, 2026, says FinOps examiners will enquire about AI use and review the operational controls behind it. The report also tells each firm to assess whether automating a regulatory function is a material business change requiring advance written notice and a Form 33-109F5 filing. Mapping those controls is the first half of the CISSP-led review we run before a pilot.

Q. Can Claude Cowork summarize client meetings? –

Yes. Cowork can turn raw notes from a client meeting into a structured summary and an action list, which saves an advisor time on file notes. Treat the result as a draft for the file, not the official record under CIRO books-and-records rules, and keep client identifiers inside the 1 scoped household folder. An advisor reviews the summary before it informs any recommendation.

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

Both, and on more surfaces since launch. 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 who understand the client relationship and the supervision duty, not the whole firm at once. The Cowork toggle 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.

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