FUSION COMPUTING FIELD GUIDE

Claude Cowork for insurance brokerages: secure client files and E&O-safe drafting for Canadian brokers

Mike Pearlstein, CISSP, MSc AI

August 04, 2026

WHAT'S INSIDE

TRUSTED BY **Toronto law firms** **Hamilton manufacturers**

Vancouver clinics **GTA accounting firms** **Ontario non-profits**

British Columbia professional services

Principal brokers are asking whether Claude Cowork can take the submission and renewal grind off their producers without exposing client files. According to Statistics Canada, finance and insurance firms already use AI at 30.6%, one of the highest rates of any industry, so staff at most brokerages are testing these tools. The duty to protect client information stays with the brokerage.

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

Key takeaways

- A brokerage can use Claude Cowork on a Team or Enterprise plan with access scoped to one client folder and a written policy.
- Submission prep, renewal summaries, and claims correspondence drafts are strong fits. A licensed broker reviews anything that states coverage.
- Unreviewed AI-drafted coverage advice is an errors and omissions exposure. The review step is the control, not a courtesy.
- Cowork stores its work locally, so it sits outside your audit logs. For RIBO oversight and E&O defence, you build that record yourself.

Plan a Safe Cowork Pilot

Can insurance brokerages use Claude Cowork with client files?

Yes, an insurance brokerage can use Claude Cowork on a Team or Enterprise plan, with access scoped to one client's working folder and a written usage policy. RIBO holds Ontario brokers to standards of professional conduct and competence, and PIPEDA binds the brokerage on client data, so the duty never moves to the software. On the business plans, your content is not used to train Anthropic's models by default.

QUICK SELF-CHECK

How AI-ready is your business, really?

Score your AI readiness across data, use cases, governance, people and ROI in about 2 minutes, with your ranked gaps and a 30/60/90-day plan.

Take the AI readiness assessment →

or book a free 30-min call →

Free · no email until you see your score, or talk to a senior engineer.

What a brokerage controls is the scope: which client folders the agent opens, which plan governs the data, and who reviews the work before it reaches a client or an insurer. On the business tiers, Anthropic's privacy commitments keep that data out of model training.

It's the same secure-adoption logic from the pillar guide on using Claude Cowork securely in your business, applied to a brokerage, and it sits beside the parallel guides for wealth management firms and financial services under their own regulators.

What does Claude Cowork actually do for an insurance brokerage?

> Claude Cowork completes multi-step office work rather than answering a single question. Per Anthropic's Cowork documentation, the agent works across your own files and apps on the desktop, so the practical jobs are submission packages for insurers, renewal comparison summaries, first drafts of claims correspondence, certificate request handling, and cleaning up exported client lists. Each output is a draft for a licensed broker to review.

Here's how those jobs map to the work, with the guardrail that keeps each one safe. Fusion Computing walks brokerages 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 brokerage safely →

Task	What Cowork does	The guardrail
Submission prep	Assembles applications and loss runs into a clean package for markets	Scope to one client folder; a broker reviews
Renewal summaries	Compares expiring and renewal terms into a plain-English table	A licensed broker verifies every coverage statement
Claims correspondence	Drafts acknowledgment and status letters from the file	Draft only; nothing goes out unreviewed
Certificate requests	Drafts certificates and cover notes from policy details	Issued through your broker management system, not the agent
Client-list cleanup	Deduplicates and categorizes exported contact and policy lists	Work on exports in the scoped folder, never the live system

How do the confidentiality and E&O guardrails work?

> According to Anthropic's deployment guidance, the core guardrail is least privilege: scope Cowork to one client's working folder, never the whole book of business. A renewal summary that misstates a limit or an exclusion is an errors and omissions claim waiting for a loss, so a licensed broker verifies anything that states coverage. Cowork runs in an isolated virtual machine, but prompts still reach Anthropic, so scope limits what can be exposed.

The mistake I flag most often is scope. When a brokerage points the agent at a full book export, a single task can read every insured's file. Scope it to the active client and most of the risk disappears.

> **Field note.** A first-person field observation: in the brokerage pilots I've run, the first thing I change is access. I've seen a CSR hand an agent a full book-of-business export to answer one renewal question. We scoped it to the single client folder, and the workflow that felt reckless became routine. The work's identical; the exposure isn't.

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. I pair the policy with a technical review so the scope survives renewal season.

Fusion Computing pairs that policy work with a cybersecurity review so the brokerage has a defensible position.

The oversight gap for RIBO oversight and E&O defence

> According to Anthropic, Claude Cowork stores its conversation history locally on each user's computer, outside audit logs, the Compliance API, and data exports. For a brokerage it matters twice: RIBO's conduct oversight assumes the office can show how client work was done, and an E&O defence turns on documentation of the advice. Team and Enterprise owners can stream Cowork events to a SIEM through OpenTelemetry, which Anthropic notes does not replace audit logging for compliance.

According to Anthropic's guidance on using Cowork on Team and Enterprise plans, the local history "is not subject to Anthropic's standard data retention policies and cannot be centrally managed or exported by admins."

> **Field note.** When I walk a principal broker through the oversight gap, I open the local Cowork history on the demo machine and ask who else can see it. Nobody can. If a client disputes what was recommended at renewal, the file the defence lawyer wants is the one that shows how the summary was produced, and that transcript lives on one CSR's laptop.

> *"Principal brokers ask me whether Cowork can run renewals. It can draft them. The coverage advice still belongs to the licensed broker who gives it, and the offices that hold that line get the speed without the E&O problem."*
>
> ***Mike Pearlstein**, CISSP, CEO, Fusion Computing*

That doesn't rule Cowork out. It means the brokerage designs its own record of AI-assisted work.

Fusion Computing wires the OpenTelemetry stream into the same monitoring we run for managed detection and response, so a brokerage sees tool calls and file access even though the transcript stays on the device. If a draft could support a coverage decision or an E&O defence, the office keeps that record on purpose.

Plan tier and a setup checklist for an insurance brokerage

> The plan tier is the first decision: per Anthropic's plan lineup, only Team ($25 USD monthly, $20 USD annual, 5-seat minimum) and Enterprise carry the "not trained on by default" commitment a brokerage needs. A safe rollout: scope to one client folder, keep "ask before acting" on for client files, write a usage policy, turn on OpenTelemetry monitoring, and keep a licensed broker reviewing anything that states coverage.

Cowork runs on Pro, Max, Team, and Enterprise plans per Anthropic's release notes, but only the two business tiers fit client work. Here's the checklist Fusion Computing runs with a brokerage before the first client file goes near the tool. The steps come from our 2026 professional-firm pilots; in our practice the plan-tier fix is the first change we make.

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.** A producer running client files on a personal account is the first risk to fix.
2. **Scope to one client folder.** Never the book of business or the shared drive. Widen only with a reason.
3. **Default to "ask before acting."** Cowork always asks before deleting files; keep approvals on for client data.
4. **Write an acceptable use policy.** Name the approved tool, the data that may go in, and who may run it.
5. **Turn on OpenTelemetry monitoring.** It's the only visibility you have into what the agent did.
6. **Keep a licensed broker signing off.** Nothing that states coverage, from summaries to certificates, ships without review.
7. **Map the terms to your duties.** Check Anthropic's data handling against RIBO conduct standards and PIPEDA before go-live.

None of it's exotic, and most of it takes an afternoon.

Fusion Computing sets it up as part of the managed IT work we already do for financial firms, and the same pattern carries to accounting firms and law firms under their own regulators. The review work is CISSP-led at a Microsoft Solutions Partner, the same team that has secured Canadian practices since 2012.

Claude Cowork is worth adopting for the submission and renewal work that fills a brokerage. The offices that set the plan, the scope, and the policy first are the ones that'll use it calmly through renewal season while their competitors are still arguing about whether it's allowed.

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

Frequently Asked Questions

Q. Is Claude Cowork safe for insurance client files? –

Claude Cowork can be safe for client files on a Team or Enterprise plan, with access scoped to one client folder and a licensed broker reviewing the output. The work runs locally and, on the business plans, is not used to train Anthropic's models by default. The duty to protect client information stays with the brokerage, so the controls around the tool are what make it safe.

Q. Can Claude Cowork read my broker management system? –

Treat the answer as no by design. Cowork works across local files, so the safe pattern is exporting what a task needs into a scoped folder rather than connecting the agent to the live broker management system. The BMS stays the system of record, certificates and policy changes issue from it, and the agent only ever sees the export.

Q. Does using Claude Cowork create an E&O exposure? –

The exposure comes from unreviewed output, not from the tool. A renewal summary that misstates a limit or an exclusion can ground an errors and omissions claim, so the control is simple: a licensed broker verifies every coverage statement before a client sees it. Documented review closes most of the gap, and a written policy makes the practice provable.

Q. What plan does an insurance brokerage need for Claude Cowork?

A brokerage should use the Team or Enterprise plan, not a personal Pro or Max account. Only the business tiers carry Anthropic's commitment not to train on your content by default, plus the owner and admin controls an office needs. A CSR running client files on a personal account is the first risk to remediate.

Q. Is client policy data used to train the model?

On Team and Enterprise plans, your content is not used to train Anthropic's models by default, so client and policy data processed under a business plan stays out of training. Personal Pro and Max plans follow individual privacy settings, which differ from the business default. For a brokerage, that difference is the reason to standardise on a business plan.

Q. How is Claude Cowork different from AI built into my BMS?

AI features inside a broker management system work only on the data that platform holds. Claude Cowork is a general agent that works across your own files and apps on the desktop, so it can draft a submission from a folder of applications, loss runs, and emails the BMS never sees. Most offices use both, with the BMS as the system of record.

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

Claude Cowork works on both macOS and Windows through the Claude desktop app, and it reached general availability on both on April 9, 2026. It is not available on the web or on mobile. Some capabilities, such as computer use, arrived first as research previews, so confirm the current feature list for your platform inside the app.

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

Start with a small group who understand the client files and the conduct duty, not the whole office at once. Cowork is an organization-wide setting that owners can switch on or off, and granular per-user controls are limited, so a deliberate pilot with named users beats a broad rollout. Pair it with training and a written policy before wider use.

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