FUSION COMPUTING FIELD GUIDE

Claude Cowork for business: a secure-adoption guide for Canadian SMBs

Mike Pearlstein, CISSP, MSc Computer Science (AI)

Published May 26, 2026; Updated August 12, 2026

WHAT'S INSIDE

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

Last month a managing partner at a Toronto law firm forwarded me a one-line note from her office manager. The team had started using Claude Cowork to tidy up case files, and was that a problem? She wasn't asking about productivity. She was asking whether client confidentiality had just walked onto a desktop app nobody had approved.

I get a version of that question weekly now. According to Statistics Canada (https://www150.statcan.gc.ca/n1/pub/11-621-m/11-621-m2026010-eng.htm), 12.2% of Canadian businesses used AI to produce goods or deliver services by the second quarter of 2025, double the 6.1% of a year earlier.

The tools arrived faster than the guardrails, and that gap is the whole reason I wrote this guide. My aim is to help you say yes to Claude Cowork without losing the thing that keeps regulated clients trusting you.

Mike Pearlstein, CISSP, MSc Computer Science (AI), is the founder of Fusion Computing, which has managed IT and security for Canadian SMBs across Toronto, Hamilton, and Metro Vancouver since 2012.

Key takeaways

- Claude Cowork is an agentic assistant that works inside your files and apps, so its risk profile differs from a chat window.
- Plan tier is the first control: only Team and Enterprise carry the “not trained on by default” commitment.
- Where a session runs decides what you can supervise. Local sessions stay on the laptop and cannot be exported by an admin. Cloud sessions on web and mobile are captured in the Compliance API.
- Cloud sessions are on by default for Team and off by default for Enterprise. Check which one your organization is running before you write the policy.
- Scope folder and connector access tightly before widening it. Over-broad access is the failure we flag most.
- Regulated firms in law, accounting, wealth, and healthcare need a written rule set before the first pilot, not after.

What is Claude Cowork, explained, and the difference from Claude chat

Claude Cowork is Anthropic’s agentic AI for knowledge work. A chat window answers questions. Cowork completes tasks: it reads and writes files, runs steps in an isolated environment on Anthropic’s servers for remote sessions, or in an isolated virtual machine on the member’s device for local sessions, calls the connectors you approve, and can run on a schedule. It reached general availability on the macOS and Windows desktop apps on April 9, 2026. It now also runs on the web and on mobile, in beta for Team and Enterprise.

Book a Consultation

That distinction matters more than it sounds. When your bookkeeper pastes a question into Claude chat, the model writes back text. When the same person opens Cowork and says “reconcile last month and flag anything odd,” Claude starts opening files, running calculations, and producing a finished spreadsheet. According to Anthropic’s release notes (https://support.claude.com/en/articles/12138966-release-notes), that capability runs on every paid plan.

Where a Cowork session actually runs

Two answers, and the difference decides everything downstream. A **local session** executes on the user’s own computer, with code in an isolated environment on Anthropic’s servers for remote sessions, or in an isolated virtual machine on the member’s device for local sessions. A **cloud session** runs on Anthropic’s infrastructure, so work continues across desktop, web, and mobile and a scheduled task fires when the laptop is shut.

Cloud sessions are the newer half, and their default differs by plan. Anthropic turns them **on by default for Team** and **off by default for Enterprise**. That single line decides whether your firm is already running work on Anthropic infrastructure without anyone deciding to.

I tell clients to picture a sharp new hire who can touch every file you point them at and never asks twice. That framing moves the conversation to access and supervision. It is why every AI services rollout we run starts with scoping rather than a licence purchase.

Cowork also connects to the tools small businesses already run. Anthropic’s Claude for Small Business (https://www.anthropic.com/news/claude-for-small-business) package ships native links to QuickBooks, Microsoft 365, Google Workspace, HubSpot, and Docusign. That reach is the upside, and the reason scoping is step one, exactly as with a custom business AI platform.

Why this matters for Canadian SMBs right now

> AI adoption among Canadian businesses doubled in a single year, and the firms using it are measurably more productive. The pressure to adopt is real, but adoption is uneven by sector, and the slow movers are the regulated ones with the most to lose from a careless rollout. The window to set rules before staff set their own habits is closing.

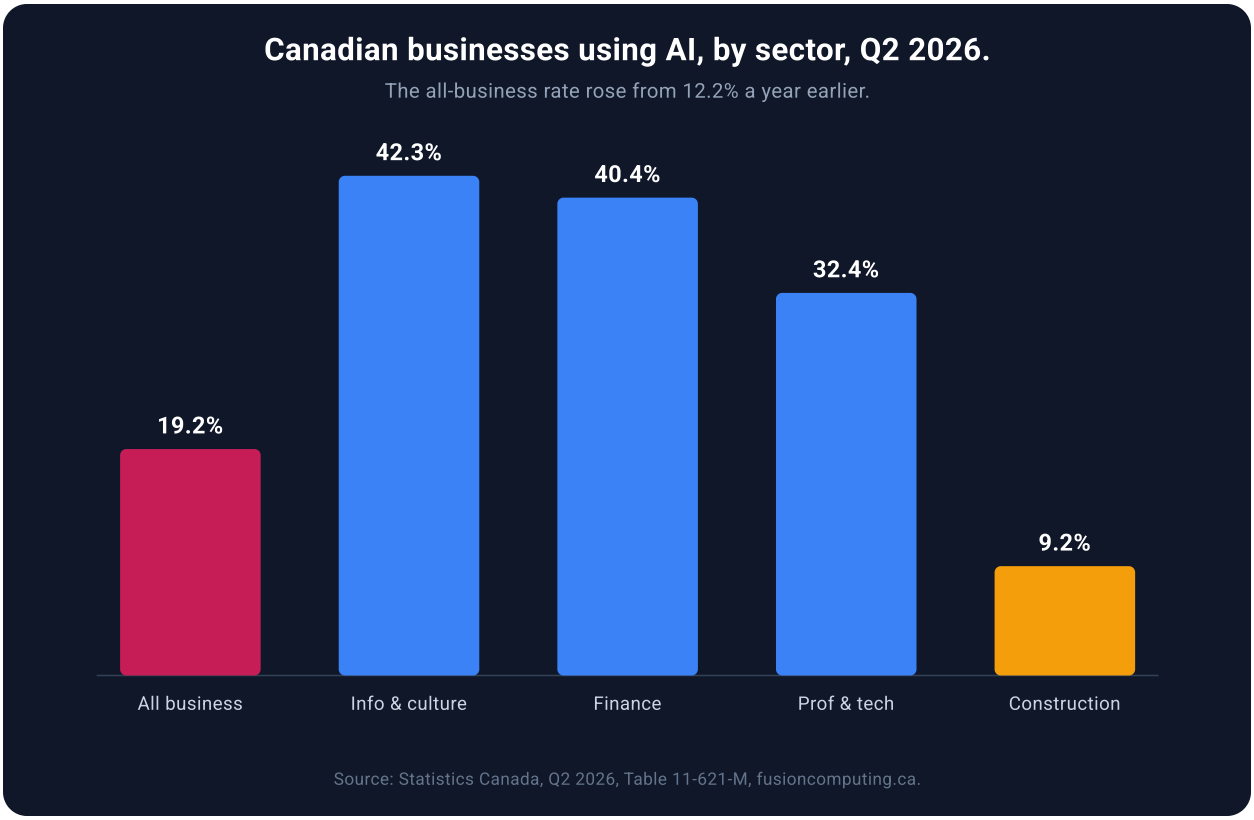

Canadian businesses using AI by sector, Q2 2026. Finance and insurance (40.4%) now leads professional, scientific and technical services (32.4%); the all-business rate was 19.2%, up from 12.2% in Q2 2025. Source: Statistics Canada, Table 11-621-M.

The sector split owners miss

The sector split is the part owners miss. Statistics Canada put AI use at 32.4% in professional, scientific and technical services and 40.4% in finance and insurance, against 1.8% in transportation and warehousing. Roughly

14.5% more firms (https://www150.statcan.gc.ca/n1/pub/11-621-m/11-621-m2025011-eng.htm) said they planned to adopt within a year. Your competitors and your own staff are already in the tools.

Ottawa is pushing the same way. The Business Development Bank of Canada (https://www.bdc.ca/en/about/mediaroom/news-releases/bdc-launches-lift-getting-canadian-smes-off-the-ai-sidelines) launched a CA$500 million loan program in 2026 to get smaller firms off the AI sidelines. When financing shows up, so does the shadow version, where people run company data on personal accounts.

The shadow-adoption pattern

I see that shadow pattern constantly. Someone wires up Cowork on a personal plan because the firm hasn't offered one, and client work now sits on an unmanaged account. Set the official path before the unofficial one hardens. We treat it as part of the same AI strategy work we do across every industry we serve.

Book a free 30-minute call to scope your Claude Cowork rollout safely.

The first question to ask: where does your data actually go?

> Cowork runs its steps in an isolated environment on Anthropic's servers for remote sessions, or in an isolated virtual machine on the member's device for local sessions, and it only touches the files, folders, and connectors the user grants. The reasoning still runs through Anthropic's service, so local execution does not mean nothing leaves the building. The control that matters is scope: decide what Cowork can reach before you turn it loose.

Here's where I slow clients down. Local execution is genuinely useful: drafts and working files stay on the machine. The prompts and the content Claude reasons over still travel to Anthropic, the way any cloud AI works. Both are true at once.

> *Running locally is not the same as staying private. The model still reasons in the cloud, so the one control a firm fully owns is the scope of what the agent can open.*
>
> *Mike Pearlstein, CISSP, Fusion Computing*

What can this install actually open?

So the practical question is never whether Cowork is safe. It's what this install can actually open. If someone points it at a synced OneDrive root, the answer is the entire firm. If they point it at one project folder, the answer is one project. We size that access the way we'd size any permission in a cybersecurity review: least privilege first, widen only on evidence.

When I walk an owner through this, the lightbulb is the folder map. We list what lives where, mark what is client or regulated data, and only then decide what an agent should ever see. That ten-minute exercise prevents most of the incidents I get called about later.

Your plan tier is a compliance decision, not a billing one

> Claude Cowork runs on Pro, Max, Team, and Enterprise plans. The line that matters for any firm touching client data is the training default. On Team and Enterprise plans your content is not used to train Anthropic's models by default. Those plans also add the owner and admin controls that personal plans lack. For regulated work, the plan tier is the first thing to get right.

Personal plans versus business plans, side by side

Here's the side by side I show clients.

Question	Pro and Max (personal)	Team and Enterprise (business)
Content used to train models by default?	Governed by personal privacy settings, not the business default.	No, not by default.
Central admin controls	None.	Owner and admin controls for access and features.
Per-team enablement	Not applicable.	Enterprise only, via groups and custom roles. Team is all-or-nothing.
Cloud sessions default	Not applicable.	On by default for Team. Off by default for Enterprise.
Audit logs	Not available.	Available on Enterprise, capturing metadata rather than chat content.
Compliance API	Not available.	Enterprise. Captures Cowork sessions run via web and mobile.
Seat minimum	One.	Enterprise starts at 20 seats self-serve, 50 sales-assisted.
Right fit for	Personal productivity, non-client data.	Any firm handling client or regulated data.

The first change we make on a regulated pilot

The first change we make for a regulated client piloting Cowork is almost always the plan tier. A law or accounting firm running client matters on a personal Max account is the risk I remediate before anything else. According to Anthropic's privacy centre (https://privacy.claude.com/), the "not trained on by default" promise lives on the business plans, which is exactly where regulated firms belong.

> **Field note from Mike**
>
> Across our Canadian SMB client base, the plan tier is almost never the thing the owner asks about first. They ask whether the tool is “secure.” When we open the billing page together and they see a personal Max subscription paying for work on client matters, the conversation changes in about ten seconds. Our engineers found that swap is usually a same-afternoon fix, and it is the cheapest control on the whole list.

One catch is worth flagging, and it lands squarely on smaller firms. The master Cowork toggle is organization-wide, so on a **Team** plan it is genuinely all-or-nothing: you cannot allow it for marketing and block it for the litigation team. **Enterprise** is where that changes. Groups and custom roles let an admin enable Cowork, or just the “Run Cowork in the cloud” capability, for specific teams.

The catch behind the catch is the seat count. Anthropic sells Enterprise (https://support.claude.com/en/articles/9797531-what-is-the-enterprise-plan) self-serve at a **20-seat minimum** and sales-assisted at **50**. A 12-person Toronto firm that wants per-team control is therefore choosing between buying 20 seats for 12 people or accepting Team’s all-or-nothing switch. We work that trade-off explicitly into the rollout, and into the AI governance work we do for advisory practices.

The oversight gap nobody mentions: local Cowork history sits outside your central audit logs

> It depends where the session ran. Local sessions store conversation history on the user’s own computer, outside Anthropic’s retention policies, and an admin cannot centrally manage or export them. Cloud sessions are different. Anthropic states that Cowork via mobile and web is captured in the Compliance API. Enterprise audit logs record metadata rather than content, so for local work the supervision record is the one you build yourself.

The split that decides your supervision plan

Anthropic's guidance on using Cowork on Team and Enterprise plans (https://support.claude.com/en/articles/13455879-use-claude-cowork-on-team-and-enterprise-plans) is precise about it. Of local sessions it says the history "cannot be centrally managed or exported by admins." Of the other half it is just as plain. "Cowork via mobile and web is captured in the Compliance API."

Think about what the local half breaks. A Law Society audit, a CIRO record-keeping request, a PHIPA access request, or a litigation hold all assume the firm can produce what happened. If a paralegal drafted a memo in a local session, that session lives on a laptop. It does not live anywhere your compliance officer can search.

So the instruction is narrower than "Cowork is unauditable." Decide which surface your regulated work runs on. If you need retrievable records, push that work to cloud sessions. If it must stay local, you are the system of record.

> *Treat every folder a desktop AI agent can reach as already disclosed. Scope the access first, monitor what you can through OpenTelemetry, and write the supervision rule before the first pilot, because the platform will not keep that record for you.*
>
> *Mike Pearlstein, CISSP, Fusion Computing*

Build the visibility yourself. If the pilot needs centralized monitoring, define the destination, filtering, access and retention before enabling OpenTelemetry. Security teams then see tool calls and approvals even when a transcript stays local.

Get a CISSP-led review of which AI tools can reach your client data.

What about Canadian data residency?

Ask it precisely, because the usual framing is wrong. PIPEDA does not require personal information to stay in Canada. The Office of the Privacy Commissioner (https://www.priv.gc.ca/en/privacy-topics/airports-and-borders/gl_dab_090127/) is blunt about it. PIPEDA "does not prohibit organizations in Canada from transferring personal information to an organization in another jurisdiction for processing." Schedule 1, Principle 4.1.3 sets the real duty: comparable protection, by contract.

What Anthropic sells on Enterprise is **US-only inference**. That keeps processing inside the United States, which is a real control and not a Canadian one. Document the transfer, hold the vendor to comparable protection, and tell clients where their information is processed.

A secure-adoption framework we use with clients: the criteria and steps

> A safe Cowork rollout comes down to seven controls. Pick a business plan and scope folder and connector access to least privilege. Keep Claude in "Manually approve" mode for sensitive work and write an acceptable use policy. Stream events to a SIEM through OpenTelemetry, keep a human approving consequential actions, and review Anthropic's data handling against your regulator. Set these before the pilot.

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

The rollout checklist we run

Here's the checklist we actually run. None of it is exotic. The technical setup can be quick; the privacy, logging and approval work is what decides how long the pilot takes.

1. **Choose Team or Enterprise.** The training default and admin controls come with the business plans.

2. **Scope access tightly.** Point Cowork at one working folder, not a synced drive root. Widen only when there's a reason.
3. **Default to "Manually approve."** Reserve the faster "act without asking" mode for low-stakes, non-client folders. Claude always asks before deleting files.
4. **Write an acceptable use policy.** Name what data is allowed in, who may run it, and what's off limits. Our guide on what belongs in an AI acceptable use policy is the template I hand clients.
5. **Turn on OpenTelemetry monitoring only after deciding what may be logged.** For local sessions it is the one visibility path you have. Use it.
6. **Decide your cloud-session posture.** Check "Run Cowork in the cloud" in Organization settings. Team is on by default, Enterprise is off, and the answer changes what the Compliance API can retrieve later.
7. **Leave connector write-approvals per task.** The "Always allow" setting for write-capable connector tools ships off by default. Leaving it off is the right call for client work.
8. **Keep a human on consequential actions.** Anything that sends, deletes, files, or pays gets a person in the loop.
9. **Check the vendor terms against your regulator.** Map Anthropic's data handling to PIPEDA, PHIPA, or your professional body before go-live.

I also reset one expectation in every kickoff. Cowork makes a team faster on the easy 30% of the work: file cleanup, the first-draft memo, the reconciliation pass. AI triages, humans own. It does not sign the compliance record or own the client relationship.

The firms that treat it as a co-worker with a supervisor do well. The ones treating it as a replacement for judgment are the ones I worry about. Fusion Computing builds the supervised version, often inside a co-managed IT engagement.

What this looks like in your industry

> The secure-adoption pattern is the same everywhere, but the sensitive data and the regulator change by vertical. A law firm guards privilege, a clinic guards health information under PHIPA, an advisor guards records under CIRO rules, and a manufacturer guards trade secrets. The control set stays constant; what you scope and supervise depends on what you hold.

Pick your vertical guide

We've published a vertical-by-vertical Cowork guide for each of these, each with the regulator and the workflows that fit. Pick yours:

- **Law firms.** Privilege and Law Society confidentiality. See Claude Cowork for law firms.
- **Accounting firms.** CPA confidentiality and QuickBooks workflows. See Claude Cowork for accounting firms.
- **Wealth management.** CIRO record-keeping and client PII. See Claude Cowork for wealth management.
- **Healthcare clinics.** PHIPA and personal health information. See Claude Cowork for healthcare clinics.
- **Manufacturers.** Trade secrets and supplier data. See Claude Cowork for manufacturers.
- **Architecture and engineering.** Project files and design IP. See Claude Cowork for A&E firms.
- **Construction.** Bid confidentiality and contracts. See Claude Cowork for construction firms.
- **Transport and logistics.** Customer contracts and rate data. See Claude Cowork for transport and logistics.
- **Ontario municipalities.** MFIPPA and public records. See Claude Cowork for municipalities.
- **Non-profits.** Donor data and grant reporting. See Claude Cowork for non-profits.

- **Financial services.** Client data and PIPEDA. See Claude Cowork for financial services.
- **Real estate brokerages.** Deal files, TRESA conduct, and FINTRAC records. See Claude Cowork for real estate brokerages.
- **Insurance brokerages.** Client files, RIBO conduct, and E&O exposure. See Claude Cowork for insurance brokerages.
- **Dental practices.** PHIPA custodianship and admin-only scope. See Claude Cowork for dental practices.

Common mistakes Canadian SMBs make rolling out Cowork

> Most Cowork problems trace to five mistakes. Pointing it at too broad a folder, running client data on a personal plan, assuming local storage means nothing leaves, skipping a written policy, and letting it act without a human on sensitive steps. None of these are exotic, and all of them are cheap to prevent before the first task runs.

> **Field note from Mike**
>
> A sixth mistake belongs on that list and rarely makes anyone's: banning the tool without checking what people already run. In our experience the ban never holds. It just moves the work to a personal account we cannot see, which is the worst of both outcomes. I would rather write a two-page policy in an afternoon and know where the work happens.

Why size does not get you out of it

The "we're too small for this to matter" reflex is the one I push back on hardest. A 12-person firm holds the same privileged files a 200-person firm does, and the regulator doesn't scale its expectations to your headcount. If anything, the small shop is more exposed, because it rarely has a compliance officer watching. Compliance doesn't skip you because you're small.

The second trap is the over-broad folder, the most frequent thing Fusion Computing flags in an AI readiness review. Staff point Cowork at an entire synced cloud drive instead of one project folder, and a single task can read the whole firm. Scope it down and most of the risk evaporates.

Want an AI acceptable use policy your regulator will actually accept?

The last one is quiet but serious: assuming the audit log has you covered. For local sessions it does not. If your supervision plan depends on reconstructing what an AI did, build that visibility on purpose. Want a second set of eyes on your setup? Talk to us, or read how we approach managed IT for regulated firms.

Start with one low-risk workflow. If the controls hold and the numbers are right, expand from there. Pick a business plan, scope the access, write the policy, and keep a human in the loop. That's the whole game.

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 business data? –

Claude Cowork can be safe for business data when it runs on a Team or Enterprise plan with access scoped to specific folders. The work runs in an isolated environment on Anthropic's servers for remote sessions, or in an isolated virtual machine on the member's device for local sessions, and prompts reach Anthropic for reasoning. Safety depends on setup: limit what the agent can open, keep a human approving sensitive actions, and write a usage policy before staff begin.

Q. What plan do I need for Claude Cowork? –

Claude Cowork runs on the Pro, Max, Team, and Enterprise plans. For any firm handling client or regulated data, Team or Enterprise is the right choice. Those business plans do not use your content to train Anthropic's models by default and add owner and admin controls. Personal Pro and Max plans follow individual privacy settings and lack central administration.

Q. Does Anthropic train its models on Cowork data? –

On Team and Enterprise plans, Anthropic does not use your content to train its models by default. Personal Pro and Max plans follow the individual privacy settings on the account, which differ from the business default. For a regulated firm, that difference is the main reason to standardise on a business plan before running any client work through Cowork.

Q. Can my administrator see what Claude Cowork did?

It depends on where the session ran. Local sessions keep their conversation history on the user's own computer, outside Anthropic's retention policies, and admins cannot centrally manage or export them. Cowork run via web and mobile is captured in the Compliance API, so those sessions are retrievable.

Enterprise audit logs record metadata rather than content. Team and Enterprise owners can also stream Cowork events to a SIEM through OpenTelemetry, which Anthropic notes does not replace audit logging for compliance. 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.

Q. Is Claude Cowork allowed under PHIPA, CIRO, or Law Society rules?

No regulator names Claude Cowork directly, so the answer depends on how a firm sets it up. PHIPA, CIRO record-keeping, and Law Society confidentiality duties all expect a firm to control and account for client information. Because local Cowork sessions keep their history on the user's computer, where admins cannot centrally manage or export it, a regulated firm should scope access tightly, keep its own supervision record, and document the controls before use.

Q. How is Claude Cowork different from ChatGPT agents for business? –

Both aim to complete multi-step tasks rather than answer single questions. Cowork works inside local files on macOS and Windows and keeps local session history on the user's machine, while its web and mobile sessions run in Anthropic's cloud. The practical differences for a business are the data-handling defaults, the admin controls, and where each tool stores its activity.

Book a Consultation

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

Both, and more than both. Cowork reached general availability on the macOS and Windows desktop apps on April 9, 2026. It also runs on the web at claude.ai and in the Claude iOS and Android apps, in beta for Team and Enterprise plans, where sessions execute on Anthropic's infrastructure rather than the user's machine. Confirm the current feature list for your platform inside the app, because the surfaces are still moving.

Q. Is Claude Cowork subject to Canadian data-residency rules? –

PIPEDA does not require personal information to stay in Canada. The Office of the Privacy Commissioner is explicit. PIPEDA “does not prohibit organizations in Canada from transferring personal information to an organization in another jurisdiction for processing.” Schedule 1, Principle 4.1.3 requires comparable protection by contract instead. Anthropic sells US-only inference on Enterprise, which keeps processing in the United States rather than Canada, so treat residency as a contractual and disclosure question.

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