FUSION COMPUTING FIELD GUIDE

Claude Cowork for healthcare clinics: secure administrative work under PHIPA

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

Clinic managers want to know whether Claude Cowork can take administrative work off their plate without putting patient information at risk under PHIPA. According to Statistics Canada, 12.2% of Canadian businesses now use AI, up from 6.1% a year earlier, so clinic staff are already reaching for these tools. The duty to protect personal health information stays with the clinic.

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

Key takeaways

Book a Consultation

- A clinic can use Claude Cowork for administrative work on a Team or Enterprise plan with scoped access and a written policy.
- Keep personal health information out of the tool, de-identify it first, or run a documented PHIPA assessment. Prompts reach Anthropic.
- Scope it to an admin folder, never the EMR or clinical drive.
- Cowork keeps local-session history on the user's own computer, so the firm builds its own record of AI-assisted work. PHIPA records of use and disclosure need a record you build yourself.

Can healthcare clinics use Claude Cowork under PHIPA?

Yes, a clinic can use Claude Cowork for administrative work on a Team or Enterprise plan, with access scoped to a folder that holds no patient identifiers and a written policy. According to Ontario's Personal Health Information Protection Act, 2004, the duty to protect PHI stays with the clinic, so patient information stays out of the tool, gets de-identified first, or goes through a documented PHIPA assessment with the right consent.

The duty to protect health information is the clinic's, and no vendor setting removes it. What a clinic controls is the scope: which files the agent opens, which plan governs the data, and whether any patient information is involved at all.

It’s the same secure-adoption logic from the pillar guide on using Claude Cowork securely in your business, applied to a PHIPA custodian, and it sits alongside our broader IT for healthcare clinics work. The parallel guides cover law firms, accounting firms, and dental practices under their own regulators.

What does Claude Cowork safely do for a clinic?

> Claude Cowork is strongest at clinic administration, not clinical records. Per Anthropic’s Cowork documentation, the agent works across local files and apps, so the safe jobs use no patient identifiers: drafting administrative letters and forms, writing scheduling and reminder messages from a de-identified list, maintaining policy documents, preparing non-PHI billing references, and building intake-form templates. Each output is a draft a human reviews.

Here’s how those administrative jobs map to the work, with the guardrail that keeps each one inside PHIPA. Fusion Computing walks clinics through this before any pilot, the same way we scope any AI services engagement, and the scoping call is run by a CISSP-led team at a Microsoft Solutions Partner.

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

Task	What to test	The guardrail
Administrative letters and forms	Drafts referral cover letters, recall notices, and templates	No patient identifiers in the source folder
Scheduling and reminders	Drafts reminder and rebooking messages	De-identified or admin data only
Policy and procedure documents	Writes and updates clinic policies and SOPs	Internal documents, no patient data
Billing and claim prep	Organizes non-PHI billing references and code lists	Patient identifiers stay out of the folder
Intake-form templates	Builds and refines intake and consent templates	Templates only, never completed forms

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.

How do the PHIPA and PHI guardrails work?

> According to Anthropic's deployment guidance, Cowork 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, but prompts still reach Anthropic, so the core guardrail is keeping personal health information out of the tool by default. PHI is information about an identifiable patient. Scope the agent to an administrative folder, de-identify anything that must be processed, and run a documented PHIPA assessment before any PHI use. Keep the EMR off limits.

The first thing we check is scope. When a clinic points the agent at an EMR export or a clinical drive, every task touches PHI. Section 12(1) of PHIPA requires a custodian to take reasonable steps to protect health information against theft, loss, and unauthorized use or disclosure. Scope Cowork to an administrative folder with no identifiers and the exposure drops to near zero. That's the single change that helps most.

> **Field note.** The first thing I do is draw a hard line around the EMR. I've seen staff want to summarize charts on day one. We start with letters and scheduling instead, prove the workflow, and only revisit PHI after a privacy officer has signed off on an assessment.

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, states that PHI stays out unless assessed, and says who may run it. Fusion Computing pairs that with a cybersecurity review so the clinic has a defensible position.

The oversight gap for PHIPA records and audits

> According to Anthropic, Claude Cowork stores its conversation history locally on each user's computer, and that activity is not captured by audit logs, the Compliance API, or data exports. For a clinic this matters: PHIPA expects a custodian to account for how health information is used and disclosed. Team and Enterprise owners can 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.

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." The Enterprise audit logs that do exist capture metadata, not the content of the work.

> **Field note.** When I walk a clinic manager through the oversight gap, I open the local Cowork history on the demo machine and ask who else can see it. Nobody can. Few privacy programs contemplate a transcript that lives on one coordinator's laptop, and that is a record a privacy officer may have to produce.

> *"Clinic managers ask me whether Cowork can help with charts. It shouldn't go near them. The EMR stays off limits, the admin folder does the work, and the clinics that hold that line get the hours back without a PHIPA problem."*
>
> ***Mike Pearlstein**, CISSP, CEO, Fusion Computing*

That gap is a strong reason to keep PHI out of Cowork. If the pilot needs centralized monitoring, define the destination, filtering, access and retention before enabling OpenTelemetry. If health information is ever in play, the clinic keeps its own record on purpose.

Plan tier and a setup checklist for a clinic

> The plan tier is the first decision: per Anthropic's plan lineup, only Team ($25 USD monthly, $20 USD annual, 2-seat minimum) and Enterprise carry the "not trained on by default" commitment plus the admin controls a clinic needs. Then: scope to an administrative folder, keep "Manually approve" on, put a PHI rule in the policy, turn on OpenTelemetry monitoring, and run a PHIPA assessment before patient information goes near the tool.

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

Cowork runs on Pro, Max, Team, and Enterprise plans per Anthropic's release notes, and on the business tiers your content is not used to train models by default, as Anthropic's privacy commitments set out. Here's the

7-step checklist Fusion Computing runs with a clinic. 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.

Get a CISSP-led review of where AI tools touch patient information →

1. **Choose Team or Enterprise.** A staff member running clinic work on a personal account is the first risk to fix.
2. **Scope to an administrative folder.** Never the EMR or clinical drive.
3. **Default to "Manually approve."** Cowork always asks before deleting files; keep approvals on.
4. **Put a PHI rule in the policy.** PHI stays out unless de-identified or assessed.
5. **Turn on OpenTelemetry monitoring only after deciding what may be logged.** It gives the firm a structured visibility stream into what the agent did.
6. **Keep a privacy officer signing off.** Nothing patient-facing ships without review.
7. **Run a PHIPA assessment first.** Complete it before any patient information goes near the tool.

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 clinics, and the same pattern carries to wealth management firms under their own regulators. If you want a second set of eyes before your clinic pilots Cowork, **talk to us**. The review is CISSP-led at a Microsoft Solutions Partner, securing Canadian clinics since 2012.

Clinics operating inside a public health unit will also want the Claude Cowork guide for Ontario municipalities, which covers MFIPPA rather than PHIPA.

Claude Cowork is worth adopting for the administrative load that fills a clinic's day.Start with one low-risk workflow. If the controls hold and the numbers are right, expand from there.

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. Can a clinic put patient information into Claude Cowork? –

Treat that as a decision, never a default. Personal health information sent to Cowork reaches Anthropic, so it leaves the clinic's direct custody. A clinic should keep PHI out of the tool, de-identify data before processing, or complete a documented PHIPA assessment with appropriate consent first. Administrative work with no patient identifiers is the safe starting point.

Q. Is Claude Cowork PHIPA-compliant? –

PHIPA compliance depends on how a clinic uses a tool, never on the tool alone. Cowork can be used in a PHIPA-aligned way for administrative tasks with no patient identifiers, on a business plan, with a policy and a privacy officer involved. Using it with personal health information requires a documented assessment before you start.

Q. What can a clinic safely use Claude Cowork for? –

The safe uses involve no patient identifiers: drafting administrative letters and forms, writing scheduling and reminder messages from a de-identified list, maintaining clinic policies and procedures, preparing non-PHI billing references, and refining intake-form templates. Start there, prove the workflow, and keep clinical records and the EMR out of scope until a privacy officer has assessed any PHI use.

Q. What plan does a clinic need for Claude Cowork? –

A clinic should use the Team or Enterprise plan, never 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 a clinic needs. A staff member running clinic work on a personal account is the first risk to remediate.

Want a PHIPA-aware AI use policy before your clinic pilots Cowork? →

Q. Is patient data used to train the model? –

On Team and Enterprise plans, your content is not used to train Anthropic's models by default. The bigger point for a clinic is custody. Even without training, PHI sent to the tool reaches Anthropic, so the safe default is to keep patient information out unless a PHIPA assessment says otherwise.

Q. How is Claude Cowork different from health-specific AI? –

Health-specific AI tools are built into clinical systems and scoped to clinical use. Claude Cowork is a general desktop agent that works across your own files and apps, which makes it a strong fit for clinic administration and a poor fit for raw clinical records. The practical differences are where the data lives and how broadly the agent can reach.

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. Cowork is available on paid plans across desktop, web and mobile. Web and mobile remote sessions are in beta and are rolling out gradually across plans, so confirm availability for your plan before launch. 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 clinic should run Claude Cowork? –

Start with administrative staff and a privacy officer, never the whole clinic and never clinical users handling charts. Cowork has an organization-wide enablement switch, and on Enterprise, groups and custom roles can restrict it to selected users or teams while Team remains organization-wide, so a deliberate pilot with named users beats a broad rollout. Pair it with training and a written PHI rule first.

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