

Claude Cowork for architecture and engineering firms: secure project and document work

Mike Pearlstein, CISSP, MSc Computer Science (AI)

Published May 27, 2026; Updated August 12, 2026



WHAT'S INSIDE

1. Can architecture and engineering firms use Claude Cowork on client projects?
2. What does Claude Cowork actually do for a design practice, explained
3. The client-confidentiality and design-IP guardrails
4. What the Construction Act requires when an architect or engineer certifies payment
5. The oversight gap for professional records and project audits
6. Cowork versus the AI already inside your CAD and BIM software
7. How much does Claude Cowork cost a design firm?

8. Plan tier and a setup checklist for a design firm

9. Where a design firm should start this week

10. Frequently Asked Questions

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.

Architecture and engineering firms want to know whether Claude Cowork can take on the document load of a project without putting client confidentiality or design IP at risk. According to **Statistics Canada** (2026), 32.4% of professional, scientific and technical services businesses used AI in the 12 months to the second quarter of 2026, against 19.2% across all industries. The professional 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 architecture and engineering firms across Toronto, Hamilton, and Metro Vancouver since 2012.

A design firm does not need a different tool than everyone else. It needs the agent scoped to 1 project folder, a decision on record about whether sessions run in Anthropic's cloud, and a licensed practitioner on anything that will carry a seal.

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 1 principal and 1 project lead can pilot without a firm-wide purchase.
- Scope Cowork to 1 project folder, never the whole archive or drawing library.
- 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.
- The Construction Act defines a “payment certifier” as an architect or engineer, so certification work sits inside statutory clocks measured in days.
- A licensed practitioner signs off on anything sealed or client-facing. O. Reg. 941 s.53(2) attaches to responsibility, not authorship.

Can architecture and engineering firms use Claude Cowork on client projects?

Yes, on a Team or Enterprise plan. According to [Anthropic's commercial terms](#) (2026), “Anthropic may not train models on Customer Content from Services”. Your professional and confidentiality duties survive that promise, so scope the agent to 1 project folder and put the rule in writing first.

The control that matters is scope: which files the agent opens, which plan governs the data, and who reviews the output before it carries a seal or reaches a client. Obligations to bodies like [Professional Engineers Ontario](#) sit with the practitioner, so the setup around the tool is what keeps it defensible.

This spoke applies the secure-adoption logic from the pillar guide on [using Claude Cowork securely in your business](#) to a design practice. It sits alongside our broader [IT for architecture and engineering firms](#) work.

Firms that also self-perform or manage construction will want the companion [Claude Cowork guide for construction firms](#), which covers the same prompt-payment clocks from the contractor side.

AI use by Canadian businesses, Q2 2026.

Share reporting AI use in producing goods or delivering services.

Professional, scientific and technical services.

32.4%

All industries.

19.2%

Source: Statistics Canada, table 33-10-1167-01, Canadian Survey on Business Conditions.

What does Claude Cowork actually do for a design practice, explained

It finishes multi-step document work. According to [Anthropic's release notes](#) (2026), Cowork reached general availability on macOS and Windows on April 9, 2026, then gained web and mobile access on July 7, 2026. In a design practice that means specs, proposals, transmittals, code research, and minutes.

Here's how those 5 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 that keeps it safe.
Specification and report drafting	Drafts specs and reports from your templates and notes	Scope to 1 project; a practitioner reviews.
RFP and proposal drafting	Assembles a proposal from past work and the brief	Client data stays in the scoped folder.
Project-document organization	Sorts and renames drawings, transmittals, and records	1 project folder, never the whole archive.
Code and standard research	Summarizes building codes and standards for a project	A practitioner verifies against the source.
Meeting minutes and transmittals	Turns notes into minutes and transmittal drafts	A draft for the file, reviewed before sending.

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-confidentiality and design-IP guardrails

Least privilege is the whole control. 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, so scope it to the active project rather than the drawing library.

The first thing I check is scope. When a firm points the agent at the whole project archive, 1 task can read every client's drawings. Scope it to the active project and you have cut most of the exposure.

FIELD NOTE FROM MIKE. The first thing I change is access. I've watched a project lead point an agent at a server holding every project the firm had ever drawn. We scoped it to 1 job, 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 project, meeting notes, and draft specifications.
- **Out:** other clients' drawings, proprietary details, and anything under a separate confidentiality agreement.
- **Out:** the full project archive, the drawing library, and shared network drives.
- **Always:** a licensed practitioner on anything that will be sealed or issued to a client.

Anthropic documents 3 approval modes, and the one a design 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 design practice. 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.

What the Construction Act requires when an architect or engineer certifies payment

The clocks are short. According to the Ontario [Construction Act](#) (2026), s.6.4(1) obliges an owner to pay a proper invoice within 28 days. Section 6.4(2) leaves the owner only 14 days to give a notice of non-payment, and s.6.5(1) gives a contractor 7 days to pay subcontractors.

REGULATOR QUOTE. The Construction Act defines “payment certifier” as “an architect, engineer or any other person upon whose certificate payments are made under a contract or subcontract”. That is

your practice, named in the statute. Section 6.9 then accrues interest automatically on anything paid late, at the prejudgment rate under s.127(2) of the Courts of Justice Act.

Put those 2 facts together and the AI question changes shape. An agent that touches certification or invoicing sits inside a clock measured in days, with interest running and adjudication under Part II.1 available to the other side. The 2024, c. 20, Sched. 4 amendments came into force on January 1, 2026.

Why the seal is the harder problem.

O. Reg. 941 s.53(2) requires a practitioner to seal an engineering document where they prepared its engineering content or “otherwise assumes responsibility” for any part of it. Responsibility, not authorship, is the trigger. Section 53(3)(b)(ii) then requires the electronic seal image to carry the practitioner’s licence number.

CONTRARIAN THESIS. A stolen seal image is a working credential, not a graphic. Because s.53(3)(b)(ii) puts the licence number inside the image, a leaked seal file gives an attacker everything needed to pass a casual check. In our practice the seal image belongs under the same handling rules as a signing key, and almost no firm treats it that way.

The oversight gap for professional records and project audits

Where the session runs decides the record. 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 design 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 monitoring, 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 CAD and BIM software

Blast radius is the difference. Design-platform AI is embedded in a CAD, BIM, or project system and scoped to that system’s data. Claude Cowork reaches across your own files instead, so a clash-detection model sees 1 dataset while a desktop agent sees whatever folder and connectors an owner 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. In my experience a BIM vendor 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.

Choosing the provider who will run those guardrails is a separate decision. Our [comparison of IT providers for Canadian architecture and engineering firms](#) scores 6 provider types on data residency and restore proof.

How much does Claude Cowork cost a design firm?

Less than the scoping work around it. 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 5 times the usage is US\$100. Enterprise is US\$20 per seat plus usage at API rates.

Claude plan pricing, 2026 (US dollars).

	Team, annual billing: US\$20 per seat per month.
	Team, monthly billing: US\$25 per seat per month.
	Premium seat: US\$100 per seat per month.
	Enterprise: US\$20 per seat plus usage.

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

In our practice the licence is rarely the expensive part. We measured the real cost sitting in scoping, policy drafting, and the monitoring wiring, at 3 to 6 hours for a Toronto design firm rather than a per-seat charge.

Plan tier and a setup checklist for a design firm

Plan tier is the first decision. 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.

Need the policy first? Use our [AI acceptable use policy template](#).

That 5-year gap is the whole argument for buying 2 seats, 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 where AI tools touch your project files →](#)

The 8-step rollout.

1. **Choose Team or Enterprise.** Client project files on a personal account is the first risk to fix.
2. **Scope to 1 project folder.** Never the whole archive or drawing library. 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 a licensed practitioner signing off.** Nothing sealed or client-facing ships without review.
8. **Lock down the seal image.** It carries a licence number, so handle it like a signing key.

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 **law firms** and **manufacturers** under their own rules.

Where a design firm should start this week

Buy 2 Team seats and scope them to 1 live project folder. Set approvals to Manually approve, and put the cloud-session decision in writing before anyone opens a second folder. Fusion Computing runs that scoping session in an afternoon. 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 project files and drawings?

Claude Cowork can be safe for project files on a Team or Enterprise plan, with access scoped to 1 project folder and a practitioner 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. Can Claude Cowork draft specs and reports?

Yes. Cowork can draft specifications, reports, and proposals from your templates and project notes, which saves a design team time on the first pass. Treat every draft as a starting point for a licensed practitioner to review. Keep the source documents in a scoped project folder so the agent sees only the 1 active job rather than the whole archive.

Q. Can a practitioner seal a document an AI agent drafted? —

The seal follows responsibility rather than authorship. O. Reg. 941 s.53(2) requires a practitioner to sign, date and seal an engineering document where they prepared its engineering content or otherwise assume responsibility for any part of it. Drafting help does not change that duty, so a practitioner who seals AI-assisted content owns it. Section 53(3)(b)(ii) also requires the electronic seal image to carry the licence number.

Q. What plan does a design 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 studio only, the answer is Enterprise. Client project files on a consumer account is the first risk our CISSP-led review looks for.

Want an AI use policy that fits your client contracts and professional duties? →

Q. How is Claude Cowork different from design-specific AI? —

Design-specific AI is usually built into CAD, BIM, or project-management platforms and scoped to those systems. Claude Cowork is a general agent that works across your own files and apps, which suits specs, proposals, and document work more than modelling. The practical difference is blast radius: a clash-detection model sees 1 dataset, while an agent sees whatever folder you granted it.

Q. Does Claude Cowork help with building-code research? —

Yes, with a caveat. Cowork can summarize building codes and standards for a project and pull the relevant clauses together, which speeds up research. A practitioner should verify every reference against the official source before it informs a design decision, because a summary is a starting point rather than the authority. Keep the verification step in the file.

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 of project managers and senior staff who understand client confidentiality, never the whole firm. 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 the written policy Mike Pearlstein reviews before a pilot.

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