

Claude Cowork for manufacturers: secure document and supplier work without exposing IP

Mike Pearlstein, CISSP, MSc Computer Science (AI)

Published May 27, 2026; Updated August 12, 2026



WHAT'S INSIDE

1. Can manufacturers use Claude Cowork without exposing trade secrets?
2. What does Claude Cowork actually do on the back office and shop floor, explained
3. The IP and supplier-data guardrails
4. What the CCCS baseline controls require of a plant adopting AI
5. The oversight gap for quality records and customer audits
6. Cowork versus the AI already inside your ERP and MES
7. How much does Claude Cowork cost a manufacturer?

8. Plan tier and a setup checklist for a manufacturer

9. Where a manufacturer 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.

Manufacturers want to know whether Claude Cowork can speed up supplier and document work without leaking a drawing, a price, or a process to a third party. According to **Statistics Canada** (2026), 13.1% of Canadian manufacturers used AI in the 12 months to the second quarter of 2026, against 19.2% across all industries. The duty to protect the firm's intellectual property stays with the firm.

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

A plant does not need a different tool than everyone else. It needs the agent scoped to 1 project folder, and a decision on record about whether sessions run in Anthropic's cloud. Then a buyer or quality lead reads anything a supplier or a customer will ever see.

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 buyer and 1 quality lead can pilot without a plant-wide purchase.
- Scope Cowork to 1 project folder, never the whole engineering or supplier drive.
- 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 CCCS baseline is 13 controls scoped to organizations under 499 employees, which covers most Canadian plants.
- Keep the agent on the business network and away from the plant floor. Segmentation, not patching, is the control that survives a validated line.

Can manufacturers use Claude Cowork without exposing trade secrets?

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 duty to protect designs, pricing, and customer NDAs survives that promise, so scope the agent to 1 project folder first.

The control that matters is scope: which files the agent opens, which plan governs the data, and who reviews the output before it reaches a supplier or a customer. Trade-secret protection depends on keeping the information controlled, so the setup around the tool is what keeps it safe.

This spoke applies the secure-adoption logic from the pillar guide on [using Claude Cowork securely in your business](#) to a plant, and it sits alongside our broader [IT for manufacturers](#) work.

Plants that also design and stamp their own equipment will want the companion [Claude Cowork guide for architecture and engineering firms](#), which covers sealing duties.

AI use by Canadian businesses, Q2 2026.

Share reporting AI use in producing goods or delivering services.

All industries.

19.2%

Manufacturing.

13.1%

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

What does Claude Cowork actually do on the back office and shop floor, 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 plant that means quotes, RFQs, part lists, audit files, and supplier correspondence.

Here's how those 5 jobs map to the work, with the guardrail that protects the firm's IP. Fusion Computing walks plants 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 plant safely →](#)


Task	What to test	The guardrail that keeps it safe.
Supplier and PO processing	Reads quotes and POs, extracts terms, flags mismatches	Scope to 1 project; a buyer verifies.
RFQ and quote drafting	Drafts an RFQ or quote from your templates and specs	Pricing and designs stay in the scoped folder.
BOM and inventory cleanup	Deduplicates and reconciles part lists	A draft for review, never the system of record.
Certification and compliance docs	Organizes ISO and customer-audit documents	Internal documents, no proprietary designs.
Vendor communication drafts	Drafts emails and follow-ups to suppliers	A person reviews before anything is sent.

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 IP and supplier-data 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 job rather than the engineering share.

The first thing I check is scope. When a plant points the agent at the whole engineering share, 1 task can read every drawing and price. Scope it to the active project and the file-exposure risk drops sharply, though prompt-injection, connector and web-access risk remain.



FIELD NOTE FROM MIKE. The first thing I change is access. I've watched an operations lead point an agent at a drive holding every customer drawing under NDA. We scoped it to 1 job 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 job, supplier quotes for it, and your own draft correspondence.
- **Out:** customer drawings under NDA, process know-how, and pricing beyond the scoped folder.
- **Out:** the engineering share, the ERP export, and anything reachable from the plant network.
- **Always:** a buyer or quality lead on anything a supplier or a customer will read.

Anthropic documents 3 approval modes, and the one a manufacturer 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 plant. 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 protects its IP and customer NDAs.

What the CCCS baseline controls require of a plant adopting AI

The baseline is built for your size. According to the Canadian Centre for Cyber Security's **Baseline Cyber Security Controls** (2026), control OC.1 says organizations using them "should have less than 499 employees". That scope covers most Canadian plants, and there are 13 baseline controls.

REGULATOR QUOTE. Baseline control 9.1 tells organizations to run "dedicated firewalls at the boundaries between its corporate network and the Internet" and to isolate Internet-facing servers from the rest of the corporate network. That boundary logic is the one to borrow when you decide where a desktop agent is allowed to run.

Read alongside the 18 controls, 153 safeguards, and 3 implementation groups of CIS Controls v8.1, the baseline gives a plant a defensible answer for a customer audit. I map the agent to both before a pilot. Neither framework has a control written for an AI agent, so you map it to access and boundaries instead.

Where the frameworks stop and judgement starts.

Baseline control 2.1 asks you to "enable automatic patching for all software and hardware" or run full vulnerability management. On the business network that is straightforward. On a validated line it is not, because a controller needs a scheduled shutdown before anything changes, and the baseline sets no timeframe.

CONTRARIAN THESIS. Patching is the wrong first question for a plant. The CCCS baseline contains no control for isolating operational technology, so the honest criterion is segmentation: can the agent, or

anything it opens, reach the controller network at all? In our practice a plant that answers that cleanly can defer a patch window without deferring the risk decision.

The oversight gap for quality records and customer 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 plant 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. A plant that wants AI-assisted work traceable for an [ISO 9001](#) review builds that record itself. Our engineers found the OpenTelemetry 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 ERP and MES

Blast radius is the difference. Manufacturing AI is embedded in an MES, ERP, or quality platform and scoped to that system's data. Claude Cowork reaches across your own files instead, so a scheduling model sees 1 dataset while a desktop agent sees whatever folder and connectors an owner granted it.

That difference cuts both ways, and the question I ask first is which system already hands you an audit trail. In my experience an ERP vendor carries part of that burden for you. With Cowork the plant 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 a manufacturer?

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 plant, and my order is always the same. The floor is 2 licences, so a buyer and 1 quality lead can pilot without a plant-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 Hamilton plant rather than a per-seat charge.

Plan tier and a setup checklist for a manufacturer

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 plant 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 designs and pricing →

The 8-step rollout.

1. **Choose Team or Enterprise.** Proprietary work on a personal account is the first risk to fix.

2. **Scope to 1 project folder.** Never the whole engineering or supplier drive. Widen only with a reason.
3. **Set the approval mode to Manually approve.** Deletion always prompts, and customer drawings 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 buyer or quality lead signing off.** Nothing supplier-facing or customer-facing ships without review.
8. **Prove the segmentation.** Confirm no machine running the agent can reach the controller network.

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 plant 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 customer audit 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 plants across **Hamilton** and **Toronto and the GTA**. The same pattern carries to **construction firms** and **accounting firms** under their own rules.

Where a manufacturer should start this week

Buy 2 Team seats and scope them to 1 live job 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](#).

Related reading. Carriers and brokers face the same scoping problem with different documents. See [Claude Cowork for transport and logistics firms](#) for the rate-card and BOL version of this rollout.

Frequently Asked Questions

Q. Is Claude Cowork safe for proprietary designs? —

Keep proprietary designs out of Cowork unless they sit in a scoped folder you have chosen to share. 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. The safe default is still to expose only the documents a task needs, because scope limits exposure in a way a contract clause cannot.

Q. Can Claude Cowork read our ERP or supplier data? —

Cowork can read files and use approved connectors, so it can work with exported supplier documents or a scoped folder rather than your whole ERP. Point it at 1 project, never the entire engineering or supplier drive. A buyer should verify anything Cowork extracts before it informs a purchase order or a quote.

Q. Can Claude Cowork touch the plant floor?

—

It should not, and segmentation is how you prove it. The CCCS baseline has no control written for operational technology, so borrow baseline control 9.1 and treat the controller network as a boundary the agent never crosses. Confirm that no machine running Cowork can reach a PLC or an HMI. Patch scheduling on a validated line is a separate problem with its own shutdown window.

Q. What plan does a manufacturer 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 purchasing only, the answer is Enterprise. Proprietary work on a consumer account is the first risk our CISSP-led review looks for.

Want an AI use policy that protects your IP and customer NDAs? →

Q. How is Claude Cowork different from manufacturing AI tools?

—

Manufacturing AI is usually built into an MES, ERP, or quality platform and scoped to those systems. Claude Cowork is a general agent that works across your own files and apps, which suits document and supplier work more than line-level control. The practical difference is blast radius: a scheduling model sees 1 dataset, while an agent sees whatever folder you granted it.

Q. Does Claude Cowork help with ISO documentation?

—

Yes. Cowork can draft, update, and organize ISO 9001 procedures, work instructions, and audit documents from your own materials. Treat the output as a draft for a quality lead to review, and keep the record of how it was produced. A local Cowork session history stays on the laptop and is not captured in central audit logs or the Compliance API.

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 plant should run Claude Cowork?

—

Start with a small group in purchasing or quality who understand the IP and customer NDAs, never the whole plant. 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