FUSION COMPUTING FIELD GUIDE

Claude Cowork for transport and logistics firms: secure operations and document work

Mike Pearlstein, CISSP, MSc Computer Science (AI)

Published May 27, 2026; Updated August 12, 2026

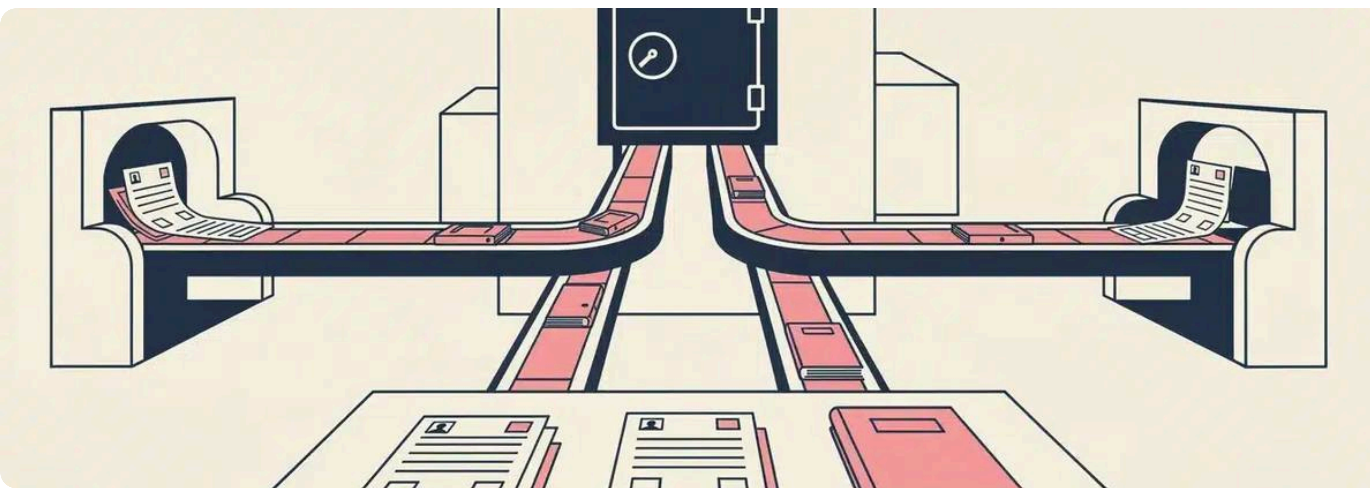

WHAT'S INSIDE

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

Transport and logistics operators want to know whether Claude Cowork can clear the document backlog without exposing a customer contract or a rate card. According to Statistics Canada, 12.2% of Canadian businesses now use AI, double the rate a year earlier, so back-office staff are already trying these tools. The duty to protect customer data and contract terms stays with the firm.

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

Key takeaways

Book a Consultation

- A transport firm can use Claude Cowork on a Team or Enterprise plan with access scoped to one account or lane folder and a written policy.
- Scope it to one account, never the whole TMS or customer drive.
- Local sessions stay on the laptop and admins cannot export them. Sessions run via web and mobile are captured in the Compliance API. Decide which surface your customer work runs on.
- Cloud sessions are on by default for Team and off by default for Enterprise. Check the toggle before you write the policy.
- A person reviews anything sent to a customer or a carrier.

Can transport and logistics firms use Claude Cowork with customer and operational data?

Yes, a transport firm can use Claude Cowork on a Team or Enterprise plan, with access scoped to one account or lane folder and a written policy. The firm owns customer-contract confidentiality and its PIPEDA duties, and no vendor setting removes them. On the business plans, your content is not used to train Anthropic’s models by default, which is why customer work belongs there and never on a personal account.

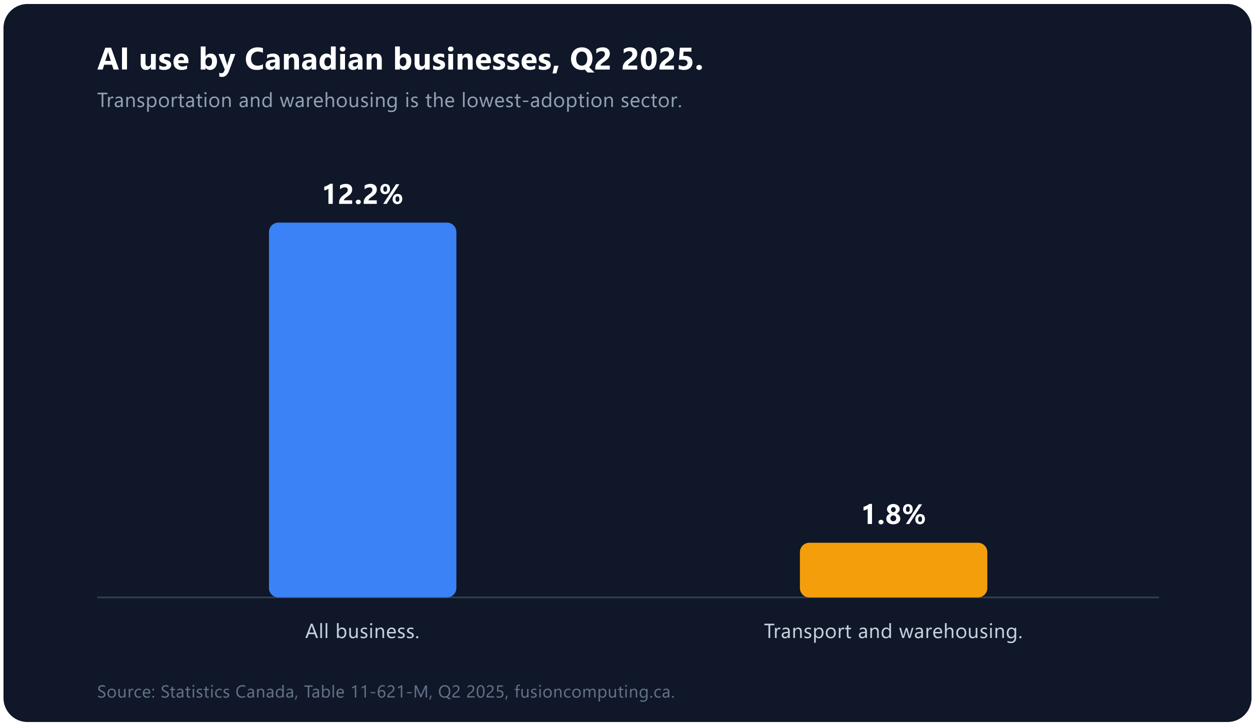

Transport and warehousing sits at 1.8% AI use against a 12.2% all-business rate. Source: Statistics Canada (2025).

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 customer. Customer trust depends on keeping rates and terms controlled, so the setup around the tool is what keeps it safe.

It’s the same secure-adoption logic from the pillar guide on using Claude Cowork securely in your business, applied to a carrier or broker. It sits alongside our broader IT for transport and logistics work.

What is Claude Cowork in a transport back office? The difference from logistics AI, explained

> Cowork completes multi-step document work. Logistics AI is different: it usually lives inside a TMS or visibility platform and is scoped to those systems. For a transport firm the practical Cowork jobs are customer and carrier documents, rate and quote drafts, dispatch summaries, safety filing, and claims cleanup. Each output is a draft a person verifies before it leaves the office.

Here's how I map those jobs to the work, with the guardrail that protects the customer. Fusion Computing walks operators 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 operation safely.

Task.	What to test.	The guardrail.
Customer and carrier documents	Reads BOLs, rate confirmations, and contracts, then extracts terms.	Scope to one account. Ops verifies.
Rate and quote drafting	Drafts rates and quotes from your templates.	Rate cards stay in a scoped folder.
Dispatch and shipment summaries	Drafts shipment and status summaries from records.	A draft for review, never the system of record.
Safety and compliance docs	Organizes National Safety Code and safety documents.	Internal documents, reviewed before filing.
Claims and invoice cleanup	Deduplicates and categorizes claims and invoices.	A draft for review.

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 customer-data and contract guardrails: why least privilege requires one folder

> Least privilege is the whole guardrail. Scope Cowork to one account or lane folder, never the whole TMS or customer drive. Classify what is allowed in, meaning working documents for the active account, and what stays out, meaning rate cards and customer identifiers beyond that folder. Keep a person reviewing anything customer-facing. Local sessions run 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.

Where operators get the scope wrong

The first thing I check is scope. When a firm connects the agent to the whole TMS, a single task can read every customer's rates. Across our Canadian SMB client base that is the single change that moves a pilot from reckless to routine. Scope it to the active account and you have cut most of the risk.

> **Field note.** The first thing I change is access. I've watched a dispatcher point an agent at a drive holding every customer's negotiated rate. We scoped it to one account 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 names the approved tool, the data that may go in, and who may run it. Our guide on what belongs in an AI acceptable use policy is the template I hand operators. Fusion Computing pairs it with a cybersecurity review, in line with PIPEDA.

The oversight gap for compliance records and audits

> It depends where the session ran. Local sessions keep their history on the user's own computer, and admins cannot centrally manage or export them. Sessions run via web and mobile are captured in the Compliance API. Safety programs and customer audits assume a firm can trace its records, so decide which surface your regulated work uses before the pilot, not after.

Anthropic's guidance on using Cowork on Team and Enterprise plans splits it cleanly. Local history "cannot be centrally managed or exported by admins." The other half is just as plain. "Cowork via mobile and web is captured in the Compliance API." Enterprise audit logs capture metadata, not the work itself.

If the pilot needs centralized monitoring, define the destination, filtering, access and retention before enabling OpenTelemetry. If a document supports a claim or a safety record, the firm keeps that trace on purpose.

> **Field note from Mike**
>
> In my experience the traceability question decides more transport rollouts than the security question does. A carrier can live with an agent that reads one folder. It cannot live with a claims file it cannot reconstruct 14 months later. So I ask operators which documents would have to be produced in a dispute, and we route that work to a surface the Compliance API can retrieve.

Choosing who secures the systems underneath

Choosing the managed IT provider that secures those systems is a separate decision. Our buyer's comparison of IT providers for Canadian transportation and logistics companies scores 5 provider types against CCCS ITSM.50.030 and the CARM access controls a carrier already carries.

> *A rate card is the most valuable unencrypted document in a brokerage, and it is usually sitting in the folder everyone can open. Scope the agent before you scope the ambition.*
>
> Mike Pearlstein, CISSP, Fusion Computing

> **Field note from Mike**
>
> I ask one question before any transport pilot: who else has this folder? At a Hamilton-area carrier of roughly 60 staff, the answer was every dispatcher, two former contractors, and a shared login on the yard terminal. Our engineers found that before we connected anything. Fixing the share took an afternoon. Connecting the agent first would have made every one of those accounts a route to the rate book.

Plan tier and a setup checklist for a transport firm

> Plan tier is the first decision. Only Team and Enterprise carry the "not trained on by default" commitment plus the admin controls a firm needs. From there the rollout is short. Scope to one account folder and keep "Manually approve" on. Write a usage policy, turn on OpenTelemetry monitoring, and keep a person signing off on customer-facing output. Then review the vendor terms against your customer contracts.

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. On the business tiers your content is not used to train models by default, as Anthropic's privacy centre sets out.

One sizing note that catches smaller carriers. Per-team control lives on Enterprise, through groups and custom roles, and Anthropic sells Enterprise at a 20-seat minimum self-serve and 50 sales-assisted. A 15-person brokerage is therefore choosing between buying 20 seats and accepting that Team is all-or-nothing. Here's the checklist Fusion Computing runs with an operator.

Get a CISSP-led review of where AI tools touch your customer and rate data.

> **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.** Customer rate data on a personal account is the first risk to fix.
2. **Scope to one account folder.** Never the whole TMS or customer drive. Widen only with a reason.
3. **Default to "Manually approve."** Cowork always asks before deleting files; keep approvals on.
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 only after deciding what may be logged.** It gives the firm a structured visibility stream into what the agent did.
6. **Keep a person signing off.** Nothing customer-facing ships without review.
7. **Map the terms to your contracts.** Check Anthropic's data handling against your customer agreements before go-live.

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, and the same pattern carries to other document-heavy operations. If you want a second set of eyes before your firm pilots Cowork, talk to us or read more about how we work.

Claude Cowork is worth adopting for the document and customer work that fills a transport back office.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. Is Claude Cowork safe for customer and rate data? –

Claude Cowork can be safe for customer and rate data on a Team or Enterprise plan, with access scoped to one account folder and a person 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. Cowork sessions run remotely by default on Anthropic's servers, while existing desktop deployments may still run locally, with code in an isolated virtual machine on the member's device, so expose only the documents a task needs. On the business plans your content is not used to train models by default, which is why customer data belongs there.

Q. Can Claude Cowork work with our TMS exports? –

Cowork can read files and use approved connectors, so it can work with TMS exports or a scoped folder rather than a live connection to your whole system. Point it at one account, never the entire customer drive. A person should verify anything Cowork extracts before it informs a rate, a quote, or a customer message.

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

A transport firm 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 firm needs. Customer rate data on a personal account is the first risk to remediate.

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

On Team and Enterprise plans, your content is not used to train Anthropic's models by default, so documents processed under a business plan stay out of training. Personal Pro and Max plans follow individual privacy settings, which differ from the business default. For a firm guarding customer relationships, that difference is the reason to use a business plan.

Want an AI use policy that protects your customer contracts?

Q. How is Claude Cowork different from logistics AI? –

Logistics AI is usually built into a TMS, routing, or visibility platform and scoped to those systems. Claude Cowork is a general desktop agent that works across your own files and apps, which suits document, rate, and claims work more than routing or tracking. The practical differences are where the data lives and how broadly the agent can reach.

Q. Does Claude Cowork help with safety paperwork? –

Yes. Cowork can draft and organize National Safety Code records, driver files, and safety policies from your own materials. Treat the output as a draft for a safety lead to review. Keep your own record of how it was produced if the session ran locally, because local history cannot be exported by an admin.

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

Both, and more besides. 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 mobile apps, in beta for Team and Enterprise, where sessions execute on Anthropic's infrastructure instead of the laptop. Confirm the current feature list inside the app, because the surfaces are still moving.

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

Start with a small group in operations or the back office who understand customer confidentiality, never the whole firm. The master Cowork switch is organization-wide, so on a Team plan it is all-or-nothing. Enterprise adds groups and custom roles, which is how you enable it for dispatch and withhold it from everyone else. Pair either route with training and a written policy.

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