AI Connector (ChatGPT & Claude)

Connect ChatGPT or Claude to your account so an assistant can read and draft your work directly instead of you copying text back and forth. Covers Service Agreements today and will extend to other modules. Read and draft-write only: finalizing, sending, signing, and cancelling stay human-only in the dashboard. Off by default, Owner-enabled, per-member authorization, every call logged. Requires Premium.

Up to date as of v1.27.0

CONTENTS

Overview

The AI Connector lets ChatGPT or Claude work with your real RCIC App data. Instead of pasting an agreement into a chat window and pasting the answer back, you ask your assistant directly: "list my draft agreements," "start a work-permit agreement for Tomas Novak and his spouse," "add biometrics for all four applicants," "check this draft for problems before I send it." The assistant reads and writes your actual account, under your own sign-in and your own permissions.

It works with both assistants because both speak the same open standard, the Model Context Protocol. There is nothing ChatGPT-specific or Claude-specific to configure, and nothing to install: you paste one URL into the assistant and sign in.

What it covers today. Service Agreements. Other modules will follow. The boundary is deliberate rather than incremental: the connector can do everything up to the point where a document becomes binding, and nothing past it.

What it costs. Nothing beyond Premium, the same subscription that unlocks the Service Agreements module. There is no separate charge and no add-on to buy. You pay OpenAI or Anthropic for the conversation on their side, as you already do.

Confidentiality: read this before you enable it

Read this before you turn the connector on. When ChatGPT or Claude reads an agreement through the connector, your client's name, their fees, and the details of their matter enter that AI conversation. That conversation is governed by YOUR agreement with OpenAI or Anthropic — it is outside RCIC App's systems and outside our contract with you. We cannot see it, control its retention, or delete it on your behalf.

You remain bound by your confidentiality duties under the CICC Code of Professional Conduct. That is a decision only you can make for your practice, which is why the connector is off by default, why only the Owner can turn it on, and why the decision is recorded.

Before you enable it, consider:

- Your AI provider's terms and data-retention settings. Business and enterprise plans generally offer stronger retention and training commitments than consumer plans.
- Whether client consent is appropriate in your practice, and whether your retainer language already covers the use of third-party technology.
- Which team members you want reaching client data this way. Each member authorizes separately, so you can enable the connector for the firm and still have only some people use it.

What we do protect. Client portal links are never exposed through the connector. PDFs are handed over as short-lived signed links rather than raw content. Our own logs record which tool was called, by whom, and on which agreement — never the contents of your AI conversation.

Turning it on and connecting your assistant

Setting up takes two steps, and they are done by different people.

Step 1 — the Owner opens the door (once for the firm)

Go to Settings, then Integrations, then AI Connector, and turn the connector on. This is Owner-only and firm-wide. It does not connect anything by itself; it only makes connection possible. Turning it back off later cuts every member's assistant off at once, regardless of what they have authorized.

Step 2 — each member connects their own assistant

- **In Claude:** Settings, then Connectors, then add a custom connector, and paste the connector URL. Claude handles the sign-in itself. The connector then works across claude.ai, the desktop app, and mobile.
- **In ChatGPT:** add the connector in developer mode, available on Plus, Pro, Team, and Enterprise plans. Note that ChatGPT asks you to re-enable the connector per conversation.

Either way you sign in with your normal RCIC App credentials and approve the connection on a consent screen that names the assistant asking. Nothing is connected until you approve it.

The exact URL, with screenshots, is on the public setup guide at rcicapp.ca/connector, which is also linked from the Help page and from the settings card itself.

What the assistant can and cannot do

The connector is deliberately split at the point where a document becomes binding. Everything up to that point is available; nothing past it is.

What the assistant can read

- List your agreements, with filters for status, client name, and date range.
- Open one agreement in full, including its current validation state.

- Check an agreement's status and lifecycle without pulling the whole document.
- List the templates visible to your firm, global and your own.
- Browse the government-fee catalogue and search it by name.
- Run the validator and get back the same blockers and warnings the Builder shows you.
- Get a short-lived link to a draft or signed PDF.

What the assistant can write, at the draft stage only

- Create a new draft agreement.
- Update client identity and matter details.
- Add people to the agreement and add services to their cards.
- Add and edit government fees.
- Apply AI Populate results, and run an AI review.

What the assistant can never do

Finalize an agreement, send it for signature, sign or counter-sign it, cancel it, amend it, reopen it, or reassign its RCIC. It cannot touch client portal links and it cannot send email to your clients. Those actions carry legal weight and stay in your hands, in the dashboard. The assistant's tool descriptions say so explicitly, so a well-behaved model will tell you to finish in the dashboard rather than hunting for a send button that does not exist.

Government fees are priced by us, not by the assistant. When it picks a fee from the catalogue, the amount, name, authority, per-applicant rule, and effective date are all taken from the current published version on our side; the assistant supplies only which fee it wants. If it proposes an amount that differs from the catalogue, that counts as an override and requires a written reason, which then prints on the agreement. Fees the catalogue does not carry — a Quebec program fee, a one-off disbursement — can be added with a name and amount and are always marked as overrides.

Permissions, logging, and disconnecting

The Owner switch opens the door for the firm, but it does not hand anyone a key. Each team member authorizes their own assistant with their own sign-in, and everything they do through it is bounded by the same permissions they already have in the dashboard.

- A member whose seat does not include Service Agreements cannot reach agreements through an assistant either.
- The same write gate that applies in the dashboard applies here, so an expired trial or a lapsed subscription blocks connector writes exactly as it blocks dashboard writes.
- Every call is logged with the tool used, the member, the agreement, and which assistant made the request.

Disconnecting

To disconnect your own assistant, open Profile, then Connected apps. Every assistant you have authorized is listed with a Revoke button; revoking takes effect immediately for that assistant. To shut the whole thing down for the firm, the Owner turns the switch off in Settings, then Integrations, then AI Connector — that cuts off every member at once.

Turning the firm switch back on later does not silently restore anything. Each assistant has to be authorized again, deliberately, by the person who owns it.

Working with it day to day

A few patterns that work well once the connector is on.

-

Draft from a conversation. Describe the family and the matter in plain language and ask the assistant to create the draft. It can create the agreement, add each person, and put the right service on each card in one pass. You then open the Builder to check it.

- **Validate before you send.** Ask the assistant to run the validator and explain what is blocking. It gets the same blockers and warnings you see in the Send checklist, and can usually fix the mechanical ones itself.
- **Price the disbursements.** Ask it to add the government fees for a family of four. It pulls the current amounts from the catalogue rather than recalling them, which is exactly the failure mode you want removed from an AI conversation.
- **Ask before you commit.** The connector cannot send anything, so there is no path where a misunderstood instruction puts a document in front of your client. The worst case is a draft you disagree with, which you fix or delete.

Always review in the dashboard before sending. The assistant is a drafting aid, not a compliance pass. You are the licensee: the agreement that goes to your client is the one you read and approved, whatever produced the first version of it.