

Policy Explorer — Beta Onboarding

Policy Explorer is a research tool you use from inside Claude (claude.ai). You ask a question about municipal codes, and it reads the actual ordinance text and answers with citations you can click straight back to the source. This page gets you connected and covers a few house rules for the beta.

1. Connect it in Claude (one time, about two minutes)

Before you start: you'll need a Claude account. If you don't have one, go to **claude.ai** and sign up first — the free plan works fine, no paid plan required. If you've used Claude before, just sign in.

This is a **two-part process**: first you add Policy Explorer to your connector library, then you actually connect to it. You need both — adding it alone won't sign you in.

Part 1 — add it to your library:

1. In a chat window, click the + button.
2. Go to **Connectors > + Connector > Add custom connector**.
3. Enter the name **Policy Explorer MCP**, paste this address, and click **Connect**:

```
https://policy-explorer-production.up.railway.app/mcp
```

This just adds it to your list — you won't be asked to sign in yet. That's expected.

Part 2 — connect to it:

1. Click the + again, then **Manage connectors**.
2. Find **Policy Explorer MCP** and click **Connect**.
3. Now Claude sends you to a **sign-in page**. Use the login we sent you, then click **Allow** on the permission screen.

That's it — you only do this once; it stays attached.

Optional, but it saves a lot of clicking: by default Claude asks permission every time it uses a tool — and a single research run uses many, so that's a lot of "Allow" taps. To switch it off in one place, go to **Settings > Connectors > Policy Explorer MCP** and set the tool permission dropdown to **always allow**. (Doing it from the in-chat prompt only clears that one tool, so you'd keep getting asked — the connector setting covers all of them at once.) Entirely optional; declining just means extra clicks, not a worse answer.

If any step doesn't look like this or you hit an error, skip to "If something breaks."

2. Ask your first question

Start a normal chat and ask a research question. For example:

"How do towns in DuPage County regulate coach houses (accessory dwelling units)?"

One thing to expect, and it's the important part: **before it searches, it checks two things with you.**

1. **What you're actually asking** — the kind of answer you want and which places it should cover.
2. **The specific points to compare** — it lists these and asks you to confirm.

Read each one and either confirm it or correct it. It won't read anything until you've okayed both. Those two quick confirmations are what keep the answer on-target instead of off researching the wrong thing. After that it does the reading and comes back with findings, each tied to the exact section of code it came from.

As it works, it'll also pause at each stage and offer you next-step choices — keep going, dig deeper, or wrap up. Pick one; it's keeping you in the driver's seat rather than running off on its own.

3. What to expect on longer questions

Bigger questions take real work — it may read dozens of ordinance sections, and that can run several minutes.

- **It reports as it goes.** Rather than going silent, it'll hand you progress updates along the way. That's normal, not a stall.
- **Claude may pause and show a “Continue” button.** Claude chat is limited in the number of tools it can use in a single turn. Complex questions may result in the limit being hit. If this occurs, you will be asked to click **Continue** — it then picks up exactly where it left off. This is a normal Claude behavior, not an error and not the tool breaking. Just click it.

4. Reading the answers

Every finding traces back to the source. A few things worth knowing about how to read what comes back:

- **Citations are clickable.** Each finding links to the exact section of municipal code it came from — click through to read the original language yourself.
- **“Not found” is not the same as “doesn't exist.”** If a town isn't listed in an answer, that means the tool didn't confirm a provision there — not that the town has none. It's careful about that distinction on purpose.
- **It tells you what it did and didn't check.** For bigger comparisons it'll say how many places it sampled and why, and flag what's still unverified. Take that framing seriously — it's the difference between “here's the pattern” and “here's a complete count.”

5. Coordinate before a big question

We all share **one research workspace** — findings that one person confirms are saved and visible to everyone. That's intentional; we're building up a shared, checked record. It does mean two rules:

- **One research question per topic area at a time.** If a colleague is already researching, say, parking rules, hold your parking question until they've finished.
- **Give a heads-up before a big run** in [team channel], so two people aren't researching the same topic at once.

Quick one-off look-ups are fine anytime. It's the larger “compare these towns on this topic” runs that we take turns on.

6. If something breaks

It's a beta, so now and then something will look wrong or stop partway. When that happens, send Jason three things:

1. A **screenshot** of what you saw.
2. The **exact question** you typed.
3. **Roughly what time** it happened.

Those three let us find your session in our activity log and see what went wrong. The time and the exact wording help more than they might seem to — please include them both.

Login not working, or a question about any of this? Contact Jason.