

Connect an AI assistant

Ask Claude, Codex, or any MCP-compatible assistant about legislation, and have it answer from the record instead of from memory.

What this is

Conservative States runs a [Model Context Protocol](#) server. Connecting it gives your AI assistant fifteen read-only tools over the legislative corpus: bills and their full text, amendments, roll calls, legislators, committees, and session-wide aggregates. Coverage is Tennessee today.

Every answer is grounded in a query against the same database the app uses, and results carry citations that link back to the bill. The assistant is not recalling a bill — it is reading one.

Read-only, always. There is no tool that writes, sends, publishes, or changes anything, and the server is built so that adding one would take a deliberate new capability rather than a new list entry.

Before you start

- You need a Conservative States **organization account**. External AI access is an organization feature; individual subscriber accounts cannot connect yet.
- An owner has to switch on **Settings → MCP connections → External AI access**. Doing so does not grant anyone access — it lets each person set up their own.
- Your role needs permission to use it. Owners and editors do; analysts and viewers do not.

The server address is the same for everyone:

```
https://mcp.conservativestates.app/mcp
```

Two ways to connect

This trips people up, so it is worth being blunt: these are **alternatives, not steps**. Do one or the other, and which one you can use depends on your tool, not on your preference.

Sign in from the app

You paste a URL, the app opens a sign-in page, you approve. No credential ever touches your clipboard.

Only this way can reach your organization's tracked bills and positions.

Paste a token

You create a token in settings and configure your tool to send it. For anything that cannot open a browser.

Public legislative data only — never your organization's own positions.

The reason for that split: a stance on a bill is the most strategically sensitive thing in the product. A pasted token can sit in a config file for months; a signed-in connection has consent, ten-minute credentials, and a revoke button.

Claude (claude.ai and the desktop app)

- 1 Open **Settings** → **Connectors**.
- 2 Choose **Add custom connector**.
- 3 Paste `https://mcp.conservativestates.app/mcp` as the server URL.
- 4 Approve the sign-in screen that opens. That is the whole flow — there is no token to enter.

Start a new conversation and ask something like “what moved in the legislature this week?” The assistant will call the tools on its own.

Claude Code

Add the server, then authenticate from inside the session:

```
claude mcp add --transport http conservative-states https://mcp.conservativestates.app/mcp
```

Run `/mcp` and pick the server to sign in. Alternatively, if you would rather use a token, create one in settings and pass it as a header instead — no browser step:

```
claude mcp add --transport http conservative-states https://mcp.conservativestates.app/mcp \
  --header "Authorization: Bearer cs_mcp_..."
```

Codex

Codex takes the token from an environment variable rather than the config file, which is the right way round:

```
export CS_MCP_TOKEN=cs_mcp_...
codex mcp add conservative_states --url https://mcp.conservativestates.app/mcp \
  --bearer-token-env-var CS_MCP_TOKEN
```

OpenAI Responses API

Add it as a hosted MCP tool:

```
{
  "type": "mcp",
  "server_label": "conservative_states",
  "server_url": "https://mcp.conservativestates.app/mcp",
  "authorization": "cs_mcp_...",
  "require_approval": "never"
}
```

MCP Inspector

Useful for looking at the raw tool list and calling one by hand:

```
npx @modelcontextprotocol/inspector --cli https://mcp.conservativestates.app/mcp \
  --transport http --header "Authorization: Bearer cs_mcp_..." \
  --method tools/list
```

Other clients

Any client speaking MCP over Streamable HTTP should work. The server implements the 2026-07-28 revision and also serves the 2025 revision that most shipping clients still use, with OAuth 2.1 discovery via `/well-known/oauth-protected-resource`.

The five above are the ones we have actually tested end to end. ChatGPT's in-product connectors and the Anthropic Messages API connector are **untested** — they may well work, but we would rather say so than imply a guarantee we have not earned.

What you can ask for

TOOL	WHAT IT RETURNS	COST
search_bills	Find bills by structured filters: policy issue, sponsor, committee, status, filing date, significance, or whether the bill is currently moving. Use it when the criteria are concrete.	standard
semantic_search_legislation	Searches the full text of bills, amendments, and fiscal notes by meaning, for when a bill addresses a topic without using its obvious words. Returns passages with the bill they came from, and handles exact strings such as statute citations.	expensiv
get_bill	Full detail for one bill: sponsors, committee, status, action history, companions, policy tags, key provisions, and scores. The right call before any specific factual claim.	standard
get_bill_text	The actual text, by section and page. Use it when the exact wording matters — quoting a provision, checking what a section says, reading an amendment.	standard
get_bill_changes	Amendments filed against a bill, what each does, and whether it was adopted, withdrawn, or is still pending.	standard
compare_bills	Two to five bills aligned side by side — provisions, statutes touched, status, scores — plus which are companions and how textually similar they are.	standard

TOOL	WHAT IT RETURNS	COST
get_similar_bills	Bills whose text and policy resemble a given bill, ranked, with companions labelled.	standard
get_votes	Totals by chamber and party, who voted against, who was absent, and any party defectors. Ask for every member and how they voted to build a whip count.	standard
get_legislator	A member's profile plus sponsorship and voting activity this session: bill counts, top policy areas, and how often they vote against their own party. Omit the name to rank a whole caucus instead of asking member by member.	standard
get_legislator_dissents	The specific roll calls where a member voted against their own party — the bill, the date, how they voted, what their party voted, and what the other party voted. This is the list behind the count.	expensiv
get_committee	Membership with chair and vice chair, bills pending before it, and scheduled hearings. Searching bills tells you what sits in a committee; only this tells you who decides.	standard
get_policy_framework	The conservative position framework this organization measures legislation against — scope, principles, the bright-line support and oppose rules, Tennessee's legal landscape, and the traps — across 13 policy domains. Ask for a domain by name, or pass a bill and get the framework for its own domain.	cheap
get_recent_legislative_activity	Bills filed, actions taken, votes recorded, amendments filed, hearings scheduled. The starting point for "what changed this week", rather than searching text.	standard
aggregate_legislative_data	Counts and averages across the session, computed in the database rather than sampled from a page of results.	expensiv
get_org_portfolio	Your organization's own tracked bills, issue portfolios, and positions — what you support, what you oppose, and how much you care.	cheap

You do not call these by name. Ask a question in ordinary language and the assistant picks; the names matter only when something goes wrong and you want to know what it did.

The [full reference](#) covers each tool's arguments and sharp edges, with worked examples of the chains that answer real questions — and is available as a [PDF](#) to hand to a colleague.

How results come back

- **Citations.** Results carry stable identifiers and links back to each bill, so an assistant can show its work.
- **Freshness and completeness.** Every response says whether it was cut short, and whether any derived analysis is still being generated — so a partial answer cannot be narrated as a complete one.
- **Paging.** Long lists return a cursor. The assistant passes it back to continue; cursors are tied to the exact query and expire after an hour.

Limits

Per connection, per minute:

- 120 cheap calls
- 90 standard calls
- 20 expensive calls — full-text search and session-wide aggregates
- 150 calls total, whatever the mix

Generous by design: a single research question legitimately makes a dozen calls. If you hit a limit, the assistant is told to wait rather than being cut off.

Privacy and control

- **Your questions are not stored.** We log which tool ran, how long it took, and how much came back — never the arguments. What you ask a bill about can itself reveal strategy, so it is kept out of the log by construction.
- **A connection sees exactly what you see.** Permissions are resolved fresh on every request, so a role change takes effect within the minute — access does not linger in an old credential.
- **Revocation is immediate.** Revoke from Settings → MCP connections and the next request fails. There is no cache to wait out.
- **An owner can switch the whole thing off** for the organization at any time, which pauses every connection at once without destroying them.

When something goes wrong

The app says it cannot authenticate, or every request fails with 401.

The connection was revoked, the token expired, or the organization switched external AI access off. Check Settings → MCP connections.

Connecting is refused with 403.

Your role does not include permission to use external AI, or the organization has not switched it on. An owner can do both.

The assistant cannot see your organization's tracked bills.

You are connected with a pasted token. Only a signed-in connection can reach positions — reconnect using the server URL and approve the sign-in screen.

Every request fails with 400, from one network but not another.

Something between you and us is stripping the protocol headers the specification requires — usually a corporate proxy. Try a different network to confirm before digging.

Changes

- **August 2026** — `get_policy_framework`: the position framework this organization analyzes legislation against, across 13 policy domains. The bill tools are descriptive by design and carry no position, so an assistant judging whether something is conservative was previously using its own idea of the word. `search_bills` also gained a government-scope filter and now returns the controversy and fiscal-direction signals.
- **August 2026** — `get_legislator_dissents`: the roll calls behind a member's party-break count, with what their own party and the other party voted on each. Break rates no longer count missed votes — Not Voting and Absent are reported separately, and were previously inflating members who simply miss votes.
- **August 2026** — Public launch. Sign-in and token connections, cursor paging on long lists, usage visible in settings.

Questions, or a client you would like tested? Reach us at support@conservativestates.app.

Tool reference

What each tool returns, what it costs, and the chains that answer questions people actually ask.

How to read this

You never call these by name. You ask a question and the assistant chooses; the names matter when you want to know what it did, or when you are deciding whether the toolset covers a job you care about.

Every tool is scoped to one state and one legislative session, resolved from your membership. Bill numbers may be written naturally — HB809, hb 809, and H.B. 809 all resolve to the stored HB0809, and an unparseable number returns an error rather than a near-miss match.

The 15 tools

search_bills

standard

legislation:read

Find bills by structured filters: policy issue, sponsor, committee, status, filing date, significance, or whether the bill is currently moving. Use it when the criteria are concrete.

Long result sets are paged — the response carries a cursor rather than silently stopping.

query · issueSlugs · statuses · sponsorName · committeeName · filedAfter · minSignificance · onlyMoving · governmentScopeDirection · limit · cursor

semantic_search_legislation

expensive

legislation:read

Searches the full text of bills, amendments, and fiscal notes by meaning, for when a bill addresses a topic without using its obvious words. Returns passages with the bill they came from, and handles exact strings such as statute citations.

This is what keyword tracking systematically misses, and the reason a bill blindsides an organization in April.

query · documentTypes · topK

get_bill

standard

legislation:read

Full detail for one bill: sponsors, committee, status, action history, companions, policy tags, key provisions, and scores. The right call before any specific factual claim.

billNumber · billId

get_bill_text

standard

legislation:read

The actual text, by section and page. Use it when the exact wording matters — quoting a provision, checking what a section says, reading an amendment.

Long documents are paged and every response reports how much remains, so an agent cannot mistake page one for the whole bill.

billNumber · billId · documentRef · documentType · section · page · includeHistorical

get_bill_changes

standard

legislation:read

Amendments filed against a bill, what each does, and whether it was adopted, withdrawn, or is still pending.

In Tennessee amendments are how bills change — there are no separately published revisions — so this is how you know a bill is still the bill you read.

billNumber · billId · adoptedOnly

compare_bills

standard

legislation:read

Two to five bills aligned side by side — provisions, statutes touched, status, scores — plus which are companions and how textually similar they are.

One call instead of several lined up by hand, which is where transcription errors creep in.

billNumbers · billIds

get_similar_bills

standard

legislation:read

Bills whose text and policy resemble a given bill, ranked, with companions labelled.

Similarity shows resemblance only. It is never evidence of coordination or of model legislation.

billNumber · billId · limit

get_votes

standard

votes:read

Totals by chamber and party, who voted against, who was absent, and any party defectors. Ask for every member and how they voted to build a whip count.

Where a source published only a tally, the member list comes back empty and flagged — so an empty dissenter list on a 12–4 vote reads as "the names were not published", never as "nobody dissented".

billNumber · billId · rollCallId · legislatorNames · includeAllVoters

get_legislator

standard

legislators:read

A member's profile plus sponsorship and voting activity this session: bill counts, top policy areas, and how often they vote against their own party. Omit the name to rank a whole caucus instead of asking member by member.

Breaking with the party counts only votes actually cast against it. Not Voting and Absent are reported separately, never as breaks — counting them was inflating members who simply miss votes into apparent mavericks.

name · peopleId · party · chamber · sort · limit · cursor

get_legislator_dissents

expensive

votes:read

The specific roll calls where a member voted against their own party — the bill, the date, how they voted, what their party voted, and what the other party voted. This is the list behind the count.

Each break says whether the member crossed to the other party's majority, and how many of their own party broke with them. A lone crossover and a break twenty colleagues joined are different findings, and a rate alone cannot tell them apart.

name · peopleId · crossedOnly · minSignificance · limit · cursor

get_committee

standard

committees:read

Membership with chair and vice chair, bills pending before it, and scheduled hearings. Searching bills tells you what sits in a committee; only this tells you who decides.

The roster is reconstructed from the roll-call record for the session — who actually sat on the panel, not who sits today. 43% of the member-committee pairs the Tennessee record evidences are absent from present-day rosters.

name · committeeId · include · limit · asOf

get_policy_framework

cheap

legislation:read

The conservative position framework this organization measures legislation against — scope, principles, the bright-line support and oppose rules, Tennessee's legal landscape, and the traps — across 13 policy domains. Ask for a domain by name, or pass a bill and get the framework for its own domain.

The bill tools are deliberately descriptive and carry no position. Without this an assistant judging whether something is "conservative" is using its own idea of the word rather than yours.

domain · billNumber · billId · includePrinciples

get_recent_legislative_activity

standard

legislation:read

Bills filed, actions taken, votes recorded, amendments filed, hearings scheduled. The starting point for "what changed this week", rather than searching text.

sinceDays · issueSlugs · limit · cursor

aggregate_legislative_data

expensive

analytics:read

Counts and averages across the session, computed in the database rather than sampled from a page of results.

Models miscount. Anything that can be arithmetic is done in SQL and handed over as a total.

metric · groupBy · issueSlugs · limit

get_org_portfolio

cheap

org:read

Your organization's own tracked bills, issue portfolios, and positions — what you support, what you oppose, and how much you care.

Reachable only through a connection you signed in to from the AI app. A pasted token can never read it.

portfolioName · onlyTracked · limit · cursor

Playbook

Each of these is one question. The chain underneath is what the assistant runs on its own — you do not orchestrate it.

The Monday morning briefing

"What moved last week on our issues, and what needs a decision from us before Thursday?"

`get_org_portfolio → get_recent_legislative_activity → get_bill ×N → get_committee`

Your tracked bills crossed against the week's actual movement, with hearings before your deadline surfaced first. Replaces the manual pass through a tracking list every Monday.

The whip count

"On HB 809 — who voted for it, who broke with their party, and who is worth a call?"

`get_votes (includeAllVoters) → get_legislator ×N`

Every member named with their vote, cross-referenced against how often each breaks with their party all session. A no vote from a member who never defects is a very different call from one who defects routinely.

Is this still the bill we endorsed?

"HB 1729 has been amended twice since we took a position. What changed, and does our support still hold?"

`get_bill_changes → get_bill_text → get_org_portfolio`

Each amendment, whether it was adopted, and the current text of the sections it touched, set against the position you actually recorded. Catches the quiet gutting amendment that turns a supported bill into one you would oppose.

Companion divergence

"Do the House and Senate versions of the school choice bill still match?"

`search_bills → compare_bills → get_bill_changes ×2`

The two bills aligned provision by provision with similarity scored, plus the amendments that pulled them apart. Answers which chamber's version you should be lobbying.

Finding what you did not know to look for

"Is anything this session quietly expanding civil liability for small employers?"

semantic_search_legislation → get_bill ×N → get_bill_text

Bills matched on meaning rather than keyword, so one that never says "small employer" still surfaces.

Reading the committee before you walk in

"Brief me on House Education Administration before Wednesday's hearing."

get_committee → get_legislator ×N → get_votes

Who sits on it, who chairs it, what it is hearing, and how each member has voted on comparable bills this session. Turns a roster into a vote map.

Session-scale pattern questions

"Which committees are the real chokepoints for our issues, and who sponsors the most bills in them?"

aggregate_legislative_data ×2 → get_committee

Counts computed across the whole session rather than sampled from a page of results. Informs where relationship-building actually pays off next interim.

Drafting grounded in the record

"Draft talking points opposing HB 809, quoting the liability language exactly."

get_bill → get_bill_text → get_votes → the agent writes

Talking points where every quotation is verbatim from the bill and every vote count is the recorded one. The agent writes the argument; it does not invent the facts.

Guardrails worth knowing

The agent is told when the record is silent

Roll calls report whether member-level detail exists. When a source publishes only a tally, the member list comes back empty and flagged — so an empty dissenter list on a 12–4 vote reads as "the names were not published", never as "nobody dissented". Without that distinction a model will confidently report a contested vote as unanimous.

Committee rosters are as of the vote, not as of today

Membership is reconstructed from the roll-call record, so a March vote is analysed against the panel that actually sat in March. Where no roll-call evidence exists, the response says so rather than implying history it does not have.

Citations are grounded by construction

Tools return the entities they legitimately exposed. A citation to a bill, member, or vote the agent was never shown is rejected rather than rendered, so a plausible-looking reference cannot be fabricated.

Retrieved text is data, not instruction

Bill text, amendments, and fiscal notes are material to analyse. Text inside a document that tries to issue instructions is analysed, not obeyed.

Your queries are not retained

Requests are audited for abuse and metering — tool name, cost class, duration — but tool arguments are deliberately never stored. Which bills you are watching and which members you are counting is the sensitive part of an advocacy organization's traffic, and the safest place to keep it is nowhere.

Access follows your staff roster

Entitlements are re-resolved from live membership on every request. A staffer who is removed or demoted loses access at the same moment they lose the interface, without anyone revoking a credential by hand.

Limits, plainly

- One state and one session per connection, resolved from your membership.
- Read-only. No tracking, alerting, position-setting, or scheduling.
- Member-level votes exist only where the source published them.
- Similarity is resemblance, not evidence of coordination.
- Significance is not likelihood of passage — a bill can matter and be going nowhere.
- Organization data requires a signed-in connection plus the organization opt-in.
- Positions are not decided here; full support and oppose analysis lives in the app.

- Rate limits are per connection, per minute, by cost class: 120 cheap, 90 standard, 20 expensive, 150 in total.

conservative-states 2026.8.2 · protocol 2026-07-28 (also serving 2025-06-18). Tool names, arguments, scopes, and cost classes on this page are generated from the running server.

[Connection guide](#).