SPERT® Forecaster: Connect AI

User Guide

Pair a compatible AI assistant with your project so it can read your forecast and explain it · forecaster.spertsuite.com

WHAT IS CONNECT AI?

Connect AI is a feature in SPERT® Forecaster that lets a compatible AI assistant — Claude, ChatGPT, Copilot, Gemini, Grok and others — read the project you have open and help you make sense of it. It is the difference between asking *"what does P80 mean?"* in the abstract and asking *"why does **my** P80 land in February?"*

This connection is read-only, and that is the whole shape of the feature. Your assistant can read your forecast; it cannot change anything in SPERT Forecaster. There is no tool it could call to edit a project — the server does not define one. If your assistant concludes something in your model should change, it will tell you what and where, and you make the change yourself. If you have used Connect AI in SPERT Story Map or SPERT Scheduler, this is the difference worth carrying over: those apps let an assistant build and edit. This one lets it *explain*.

Once connected, your assistant can:

Read your forecast	Your project setup, sprint history and velocity statistics, the inputs the forecast actually ran on, the percentile results for every distribution and every milestone, your productivity adjustments, and your deadline probability.
Explain what the numbers mean	Why P90 sits where it does, what your velocity spread is doing to the spread of outcomes, whether a milestone date is credible, what changed between two runs.
Explain the method itself	Two further tools answer *how does this work?* and *what does this word mean here?* — eight method topics and a 16-term glossary. Both work with no pairing at all, so you can ask about the method before you ever connect a project.
Tell you what it cannot see	Every read carries a freshness verdict, the list of distributions actually on your screen, and an explicit list of what the data does not include — so a good assistant can say *"I can't see that"* instead of guessing.

HOW IT WORKS — EACH SESSION

#	Action	Detail
1	Click Connect AI	The **Connect AI** button sits in the app header, just left of the theme toggle, and is visible from every tab. It appears once you have at least one project that is not archived. A consent dialog opens the first time, and again after you disconnect or sign out.
2	Choose Read Mode	**Read Mode** is the only toggle, and it starts ticked — there is no write consent to give, because there is nothing to write. With it on, a copy of the project you currently have open is uploaded to SPERT's cloud so the assistant can read it. With it off, the assistant can still explain how forecasting works but cannot see your data. Click **Connect**.
3	Copy the pairing code and prompt	A pairing code appears — a word and ten digits, such as CRANE-0384216791 — with two buttons: **Copy code** (the code alone) and **Copy prompt for your AI** (a ready-to-paste message containing the code and the instructions your assistant needs to read the data correctly). Use the prompt unless you have a reason not to. The code is single-use: unused, it expires in 15 minutes, and once your assistant has used it, it stops working within about a minute.
4	Paste the prompt into your assistant	Your assistant must already have the SPERT MCP server configured — see One-Time Setup below. Paste the prompt and send it. The assistant redeems the code, confirms which project is open, and reads your forecast.

5	Ask your questions	The button now reads **AI** with a status dot: grey when an assistant is paired but idle, pulsing blue while one is actually connected and reading — it goes back to grey about two minutes after the assistant's last read. Your answers update as you work — change an input and the assistant sees the change on its next read.
6	Disconnect	Click **Disconnect** in the panel when you are done. That ends the pairing and deletes the uploaded copy. Otherwise the pairing expires 7 days after your last activity.

Your pairing follows whichever project you are viewing — the one the Sprint History and Forecast tabs show. It does not belong to one project, which is why the button never names one — the honest claim is *"an assistant is paired to this browser"*. Switch projects and the uploaded copy is **replaced**, not added to: only ever one project is shared at a time. If you switch mid-conversation, tell your assistant to re-read before it answers. An archived project is never the one you are viewing: archive the project you have open and the pairing follows the first unarchived project in your list, as if you had switched. With every project archived, your assistant sees no project: the uploaded copy is deleted, and the Connect AI button is hidden until you unarchive one. The pairing itself is not ended; sign out, or unarchive a project and click **Disconnect**, to end it.

ONE-TIME SETUP: CONNECT YOUR AI CLIENT

Connect AI requires an AI client that supports MCP servers. Complete this setup once; every subsequent session — in Forecaster or any other SPERT app — starts at step 1 above. Claude (on the web, in the desktop app and in Claude Code), Grok, ChatGPT and Copilot Studio work today.

Claude — web and desktop app (any plan, including Free)	Lowest-friction path for non-developers — no terminal or config file required. The Free plan allows one custom connector. Go to **Customize** → **Connectors**, click **+ Add**, then **Add custom connector**. Name it **SPERT Suite** and paste https://mcp.spertsuite.com as the URL. Click **Continue** through the next screens, choosing **No sign in** for Authentication, and finish with **Add**. On a Team or Enterprise plan, an Owner adds the connector first: **Organization settings** → **Connectors** → **Add** → **Custom** → **Web**, then the same name, URL and No sign in. Each member then opens Customize → Connectors, finds SPERT Suite (marked Custom) and clicks Connect. A connector added on claude.ai also appears in the Claude desktop app automatically. There is no second setup, no Node.js and no config file. You can also add it from inside the desktop app, the same way. To use it: open a new conversation, click **+** (lower left) → **Connectors**, and make sure SPERT Suite is switched on, open SPERT Forecaster in your browser, click Connect AI, and paste the copied prompt.

Claude Code (Pro, Max, Ultra, Team, Enterprise or Console account — not the Free plan)	Best for developers — one command. Run this once, on one line:
	`claude mcp add --transport http --scope user spert-suite https://mcp.spertsuite.com`
	`--scope user` makes it available in every folder. Without it, it is added only for the folder you run the command in. Confirm with:
	`claude mcp list`
	spert-suite should be listed as Connected. If you already added the connector on claude.ai and sign in to Claude Code with that account, it is there already — check with:
	`/mcp`
	To use it: start a Claude Code session, open SPERT Forecaster, click Connect AI, and paste the copied prompt.

Grok	Go to grok.com/connectors → **New Connector** → **Custom**, paste https://mcp.spertsuite.com and finish — no sign-in is needed. On Grok Business or Enterprise, a team admin adds the connector in xAI's cloud console first. To use it: open a Grok conversation, open SPERT Forecaster in your browser, click Connect AI, and paste the copied prompt.

ChatGPT (Plus / Pro / Business / Enterprise / Edu — web only)	On Plus or Pro, you add it yourself. On Business, a workspace admin or owner adds it once — on Enterprise and Edu, so can a member they authorise. Go to **Plugins** → **+** → **Add custom MCP server**. Name it **SPERT Suite**, paste https://mcp.spertsuite.com under Connection, choose **No authentication**, tick **I understand and want to continue**, and click **Create as a plugin**. To use it: install the SPERT Suite plugin, start a new chat, type @ and pick SPERT Suite, then paste the copied prompt. Pick SPERT Suite with @ again in any later message that should re-read your project — ChatGPT uses a plugin only in the message where you pick it. **Note:** if ChatGPT shows fewer tools than expected or behaves unexpectedly, check that you picked SPERT Suite with @ in that message.
Copilot Studio (Microsoft 365, advanced)	In a Copilot Studio agent, open **Tools** → **Add a tool** → **New tool** → **Model Context Protocol**. Fill in the three required fields: **Server name** (SPERT Suite), **Server description** (for example: Reads SPERT Forecaster projects) and **Server URL** (https://mcp.spertsuite.com). Set Authentication to **None** and click Create, then **Create a new connection** → **Add to agent**. Test it in the agent's test panel. It needs a Copilot Studio licence or the free Copilot Studio trial — sign up with a work or school email; a trial agent can be tested but not published. The Copilot Studio features bundled with ordinary Microsoft 365 plans don't include the connectors that MCP tools run on. **Microsoft 365 Copilot Chat** (the standard chat experience) also connects and works well here, since reading a forecast takes only a few calls; deploying an agent built here to Copilot Chat takes additional admin configuration.

Any of them is fine for Forecaster. Reading a forecast takes a handful of tool calls and stops — there is no long build sequence to stall on, which is the one place assistants differ meaningfully. Pick whichever you already use.

Not supported today: ChatGPT Free and Go, and any other AI assistant without MCP support configured. **Not yet tested with SPERT: Gemini web app** (US, personal Google accounts): Google now lets you add a custom MCP app under Settings → Connected Apps → Custom apps → Add a custom app. **Antigravity CLI** (Google's replacement for Gemini CLI): add SPERT to `~/.gemini/config/mcp_config.json`:

```
{ "mcpServers": { "spert-suite": { "serverUrl": "https://mcp.spertsuite.com" } } }
```

READ MODE & PRIVACY

Read Mode is the only consent	There is no write consent to grant. The pairing is created with writes refused permanently, and the server re-checks that before any write could happen — so read-only is enforced in three independent places rather than promised once.
Only the open project is shared	Never your whole workspace. Switching projects replaces the shared copy; it does not accumulate. Projects you are not looking at are not uploaded and are not readable.
Turning Read Mode off	Deletes the uploaded copy immediately and keeps the pairing alive. Your assistant can still explain how forecasting works; it just cannot see your data. Toggle it back on at any time from the Connect AI panel.
Disconnecting and signing out	**Disconnect** deletes both the pairing and the uploaded copy. Signing out does the same, and so does signing in as a different account. If Disconnect or signing out cannot reach the server — for example, you are offline — a notice says *"Your AI assistant may still be connected."* and Forecaster keeps trying (see Troubleshooting). Switching from cloud to local storage deletes the uploaded copy and keeps the pairing.

Expiry	A pairing stops being readable 7 days after your last activity and is deleted shortly after. Pairing codes are single-use: an unused code expires in 15 minutes, and a used one stops working within about a minute.

Read Mode uploads regardless of your storage setting. If you keep your projects in **Local** storage, they normally never leave your browser — but enabling Read Mode uploads a copy of the open project to SPERT's cloud anyway, because that copy is how your assistant reads it. This is the one thing about Connect AI most likely to surprise a local-storage user. The consent dialog says so too. Turning Read Mode off, or disconnecting, removes it.

WHAT YOUR ASSISTANT CAN SEE

With Read Mode on, your assistant reads a snapshot of the open project. Knowing what is in it — and what is not — is what separates a useful answer from a confident wrong one.

Included	Not included
Project name, unit of measure, cadence and key dates · sprint history (the most recent 60 sprints, with velocity statistics computed over *all* included sprints) · milestones with their remaining work and completion state · productivity adjustments, including disabled ones · the forecast inputs the run actually consumed · percentile results for each distribution and each scope · your deadline probability · which options your three summary dropdowns are displaying.	The raw per-trial simulation data · the charts as drawn on your screen · any project other than the one open · the cumulative-probability curve behind the CDF chart · **the large percentage in the "Your forecast" headline** · anything not yet forecast.

The headline percentage is deliberately not shared, and this is the disclosure people ask about most. The big number in the blue "Your forecast" callout is the *true cumulative probability at the displayed date* — a different calculation from the percentile you selected, and it reads higher, often materially so. Your assistant sees the percentile behind the sentence *below* the headline, not the headline figure. If you quote the headline at your assistant and its numbers disagree, neither of you is wrong; you are looking at two different quantities.

READING YOUR FORECAST CORRECTLY

For SPERT Story Map and SPERT Scheduler, the risk is an assistant writing something wrong. Here the risk is an assistant **quoting** something wrong — numbers that are out of date, a distribution you have hidden, or a milestone that is not the one on your screen. The shared data is built to make all three of those hard to state confidently, and the copied prompt already tells your assistant to check them. These are the four things worth knowing yourself.

Freshness — check it first	Every read is labelled. **fresh**: the results match your current inputs. **stale**: you changed something after the run, *and it says what* — a milestone was added, removed or resized, the set of included sprints changed, a productivity adjustment was added, removed or edited, or a specific input changed from one value to another. **recomputing**: a simulation is running right now. **absent**: no forecast has run since the page was loaded, or none can run yet because an input is missing or a saved sprint date is not a valid date — with the reason. A good assistant reads this before quoting a date; if yours does not, ask it to.
Visible versus computed	The app computes up to six distributions, but you may only be looking at one. The data carries both sets, and splits the difference into distributions that are *invalid for your current mode* (Bootstrap needs History mode and 5+ included sprints; Uniform is Subjective-mode only) and ones you simply chose to hide in Settings. Ask your assistant to name which distribution a number comes from.

Scopes and milestones	Results are grouped by scope: the whole project, and each milestone. Milestones are cumulative — a milestone is reached when the running total of every milestone up to and including it has been delivered. Each scope is flagged with whether your Forecast tab is showing it, so a milestone you have hidden from the charts is still distinguishable from the ones you are reading.
Independently derived values	The deadline probability is recomputed from the stored run, for the Entire Project and every computed distribution, rather than read off the on-screen panel, which has its own scope and distribution pickers. It is marked as such. If it differs slightly from what your panel shows, that is why — check the panel's own scope selector before assuming either is wrong.

SPERT Forecaster is a Monte Carlo simulation over sprint velocity. It is NOT three-point PERT — there is no (O + 4M + P) / 6 anywhere in it. This is the single most common thing an AI gets wrong about this app, and it gets it wrong confidently: assistants pattern-match "PERT" to the classical weighted-average formula and explain your forecast in terms of a calculation the app does not perform. Each trial draws a velocity for every future sprint from a probability distribution, subtracts it from the remaining backlog, and counts sprints until the backlog reaches zero; ten thousand trials (the default) produce a distribution of completion sprints, read as percentiles. The copied prompt says so, and a dedicated tool exists purely to correct it — if your assistant starts describing weighted averages of optimistic, most likely and pessimistic estimates, stop it and ask it to look up the method.

GOOD QUESTIONS TO ASK

Connect AI is most useful for the questions the screen answers slowly. A few that work well, and what a good answer draws on:

Interpreting a number	*"Why is my P90 four sprints later than my P50?"* · *"Is my forecast spread wide or narrow for a team like mine?"* · *"What would P80 look like if my velocity were steadier?"*
Sanity-checking the model	*"Do my inputs look internally consistent?"* · *"My backlog is 400 but my milestones sum to 350 — which is right?"* · *"Are any of my productivity adjustments doing nothing?"* (disabled and already-past adjustments are shared precisely so it can answer this).
Choosing settings	*"Should I be using Bootstrap instead of Lognormal?"* · *"I have three sprints of history — is History mode trustworthy yet?"* · *"Is my trial count high enough for a stakeholder presentation?"*
Explaining it to someone else	*"Draft a paragraph explaining this forecast to a sponsor who has never seen a Monte Carlo schedule."* · *"What should I say when someone asks why I won't commit to the P50 date?"*
Learning the method	*"How does SPERT Forecaster actually work?"* · *"What does 'coefficient of variation' mean here?"* — both answerable with no pairing at all, so you can ask before you share anything.

WHAT STILL DOESN'T WORK

Editing anything	By design, permanently — not a gap waiting to be filled. Your assistant cannot create a project, change an input, add a milestone, or adjust a setting. It tells you what to change; you change it.
Running a forecast for you	Also deliberate. The simulation uses an unseeded random number generator, so a run triggered by an assistant would return numbers that differ from the ones on your screen for no visible reason. Run it yourself and ask again.
More than one project at a time	Only the open project is shared. To compare two projects, ask about one, switch, tell your assistant to re-read, and ask about the other.
Seeing your charts	The Burn-Up, CDF and Histogram charts are not shared as images, and the cumulative-probability curve behind the CDF is not shared as data. Your assistant reasons from the percentile results instead.

Reading a forecast you haven't run	There is nothing to read until a forecast has run since the page was loaded: results are kept in memory, not saved, so reloading the page clears them. Open the Forecast tab and it runs one by itself while Auto-recalculate is on (the default); otherwise click **Run Forecast**. The status comes back **absent** with the reason — including the cases where a run is impossible because no usable remaining backlog has been entered, or because a saved sprint date is not a valid date.

TROUBLESHOOTING

The Connect AI button is missing	It shows only while you have at least one project that is not archived. If every project is archived, the Projects tab says so: tick **Show archived (N)**, then click a tile to unarchive it and open it, or use its **Unarchive** button. Archiving does not end a pairing, so the button comes back as it was.
"No forecast data yet"	The pairing is fine but nothing has been shared. Check that **Read Mode** is on in the Connect AI panel and that you have a project open (one that is not archived), then ask your assistant to retry.
The assistant says it cannot read your project	Read Mode is off. Turn it on from the Connect AI panel — no need to re-pair.
The pairing looks healthy but there is no data	The shared copy has its own expiry and is not refreshed by your assistant reading it, so a pairing left open for days with no edits can outlive its data. Make any edit, or toggle Read Mode off and back on, and it is re-uploaded.
"Session not found"	The pairing has expired, been disconnected, or been ended by signing out or by signing in as a different account. Click Connect AI again for a fresh code — this one cannot be recovered.
"Your AI assistant may still be connected."	Disconnect or signing out could not reach the server to end your session, for example because you were offline. Until it does, a paired assistant can still read the last project you had open. Forecaster tries again each time the page loads and whenever you are back online; **Try again** retries now. **Dismiss** hides the notice and stops the retries, leaving the session to expire on its own 7 days after its last activity.
The assistant says the session belongs to another app	You pasted a pairing code from SPERT Story Map or SPERT Scheduler. Each app's pairing only works with that app's tools. Generate a code from Forecaster.
The pairing code doesn't work	Codes are single-use: an unused code expires after 15 minutes, and a used one stops working within about a minute. The panel gets a fresh code each time it opens, and a new one every 10 minutes while it stays open, so a code copied a while ago may have expired: close and reopen the panel, and copy again.
The numbers are out of date	Ask your assistant to re-read, and to check the freshness status before it answers. If it reports **stale**, the reason names what you changed — open the Forecast tab, which recomputes by itself while Auto-recalculate is on; otherwise click **Run Forecast**.
It quotes a distribution you don't have on screen	Ask which distribution each number comes from, and ask it to restrict itself to the ones actually visible to you. It is given both lists and should be using the visible one.
It describes SPERT as three-point PERT	It is pattern-matching the name rather than reading the method. Ask it to look up how the app actually works before explaining further — a tool exists specifically for this, and it needs no pairing.

The status dot never turns blue	Grey means paired but idle — normal until your assistant actually reads. If it stays grey after a successful read, check the assistant really is connected to the SPERT server rather than answering from memory: ask it to state your project name.
The connection drops after a tab switch	Leave the Forecaster tab open in the foreground while you work with your assistant. Background browser tabs get throttled, which can make an idle pairing read as disconnected; returning to the tab re-announces it.
Your client shows no SPERT tools	The client is caching an old tool list from before the server was updated. Remove and re-add the connector, or restart the client. In ChatGPT, check that you picked SPERT Suite with @ in that message; on Business, Enterprise or Edu, an admin or owner refreshes the tool list from the workspace's app settings.

SPERT® Forecaster · Connect AI Guide · © 2026 William W. Davis, MSPM, PMP · Licensed under GNU GPL v3 forecaster.spertsuite.com