NERO PROJECT

Situation Room - AI Advisors Brief Me Aloud

A cabinet meeting about my projects and household, on demand.

Paul Schindler | Built with Claude | November 2026

Executive summary

I run a set of AI agents from a home workstation, and each one reports in its own place and its own format. I designed the Situation Room so they report to me the way a cabinet briefs a head of government. Each advisor reads its own fresh sources, writes a short brief, speaks it aloud in its own voice and then takes my follow-up questions, while a scribe keeps the minutes. The meeting plays on a web page that shows the whole cabinet seated, with the current speaker standing forward. It has run on my home dashboard since September 2026, and the last full meeting took 1 minute 49 seconds to prepare and cost $1.59 in model usage.

8	1 min 49 s	$1.59	Live
permanent seats, one of them the scribe	to prepare a full meeting	model cost of that meeting	on my home dashboard since September 2026

What is in this report

Page	Contents
2	The problem, and the design choices I made (including what I ruled out)
3	How a meeting runs, and why each seat runs on a different AI model
4	The room on screen, and the safeguards that keep the briefs honest
5	Results, status and what is next, plus what this shows about how I work

The idea in one sentence

I wanted the target I had set for my whole agent system - "This is Paul, report" - to be literal. I sit down and ask for a meeting, and every advisor reports in turn, most urgent first.

What I did, and what it runs on

I designed the meeting, ruled on the look, the roster and the rules, and verified every stage by using it. AI agents wrote the code from my written briefs. Claude models write the briefs, Kokoro (an open text-to-speech model) voices them on my own graphics card, and the open Flux and Wan models drew the portraits.

1. The problem

I run a working set of AI agents, each owning one area. One writes my morning brief and tracks the decisions waiting on me. Another watches every scheduled job on the workstation. A third watches the house through Home Assistant. Others scan public social media, match my shopping list against store sales and follow my finances.

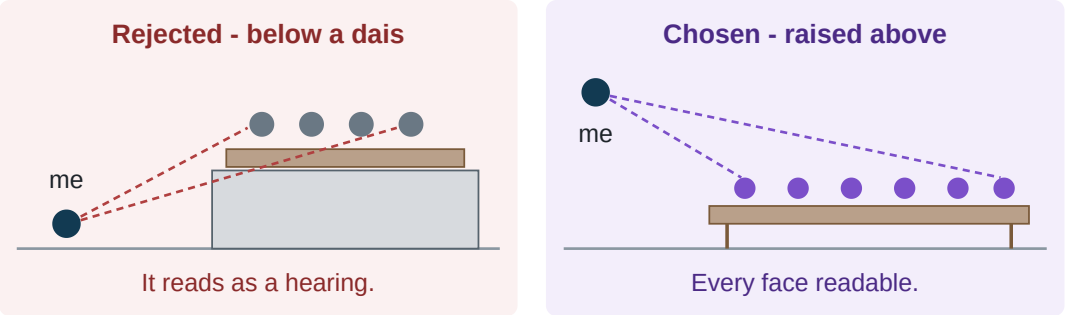
Each wrote its output in its own file and format. I wanted one place where all of them report to me at once, in a form I can take in when I am tired, with the chance to ask the advisor who just spoke a follow-up question. What was missing was a meeting. I began working out its design in my notes in February 2026.

2. Design choices I made

Question	What I chose	Why
When briefs are written	Live, at the moment I call the meeting. The short wait while the cabinet assembles is the writing window.	Pre-written cards go stale. A brief written now reflects the house now.
Who sits at the table	Eight permanent seats. Three guests join only when they bring something waiting on me.	A quiet seat says one line and sits down. An idle guest never appears.
Where I sit	First-person view, raised above a long table, looking down at the cabinet.	The first sketch put me below a raised platform. "It feels like I am the one getting grilled."
How they talk	Plain speech, like a sharp colleague across the table. No system jargon, no stock phrases.	Early briefs fell into stock AI phrasing. I banned those phrases by name.
How long	90 words per brief. Questions to me at most 25 words. Answer buttons at most 8 words, the last one always "Hold". The whole meeting capped at 4 minutes.	Brevity is respect. I raised the cap from 3 to 4 minutes when the roster grew to eight.
Stale information	An empty chair, with how long ago the seat last reported.	An empty chair is honest. A confident recital of an old file is not.
Money	Spoken in percentages only.	Amounts never need to appear on a screen or be spoken in a shared room.

What I ruled out

Option	Why I ruled it out
Briefing cards written in advance	They would recite old news. Live writing at call time costs one short wait.
Voices from the small voice pucks	Too quiet for a meeting. The design sends sound through the TV instead.



The viewpoint ruling, seen from the side. The grandeur of a long table comes from height and framing, not from putting the advisors out of sight.

3. How a meeting runs

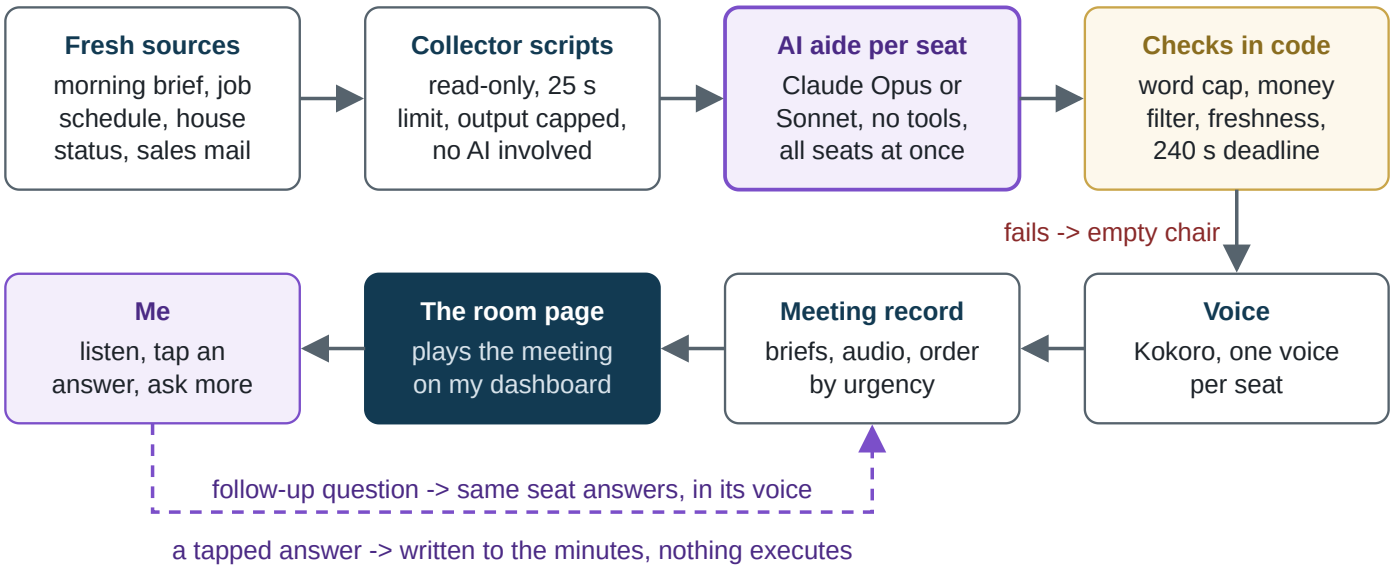
1. Call
I press a button on the room's web page, or run one command.

2. Gather
Plain scripts collect each seat's fresh sources, read-only, 25 seconds per source at most.

3. Write
One AI aide per seat, all at once, writes a brief, up to two decisions for me and an urgency score.

4. Speak
Each brief becomes audio in that seat's own voice. The most urgent seat speaks first.

5. Ask and minute
I question the seat that just spoke. The scribe opens, closes and files the minutes.



The engine. Plain scripts gather the facts; the AI only writes words from them; code checks the words before any voice speaks them.

Model choice per seat - cost and speed

Not every brief needs the most capable model. I put the two seats that weigh judgment on Claude Opus and every other seat on Claude Sonnet, which costs a fraction as much. Seats write in parallel. The last full meeting was ready in 1 minute 49 seconds, although its seats together worked for 8 minutes 25 seconds.

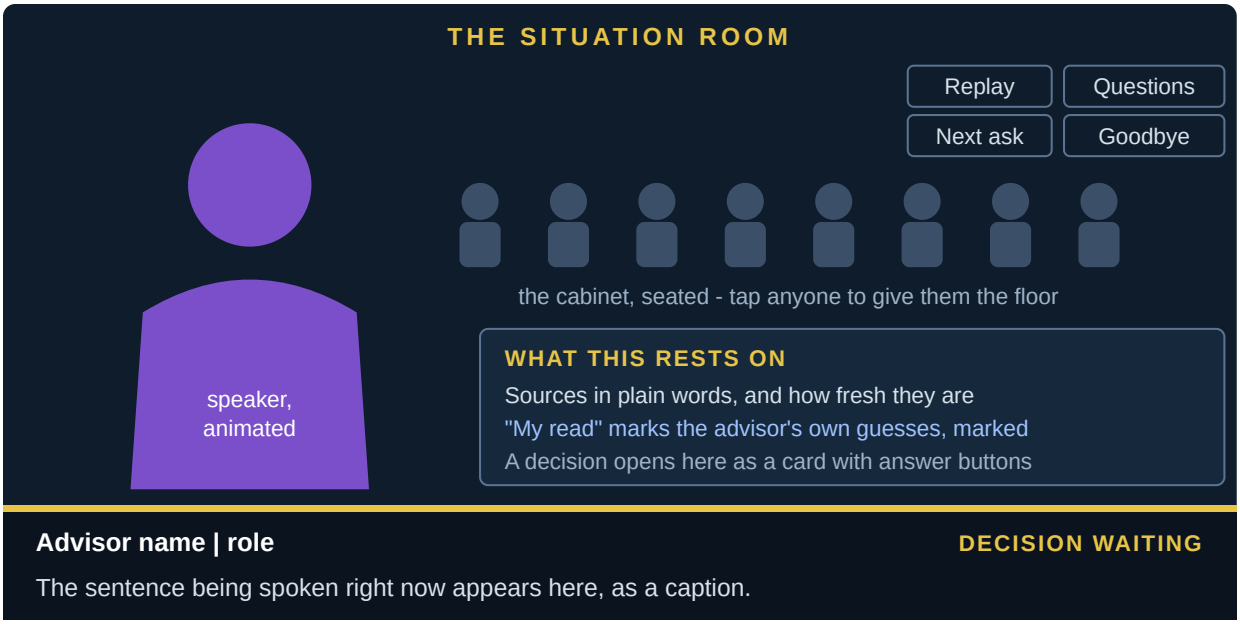
Seat	Model	Why	Cost per seat, last full meeting
Operations chief (morning brief, decisions waiting on me)	Claude Opus	Weighs which of many open items deserves my time	$0.38
Finance advisor	Claude Opus	Reasoning over changing outside facts	$0.32
Scheduled jobs, the house, public signal, deal desk, project seats	Claude Sonnet	Reporting from clear sources is enough	$0.09 to $0.19 each
Scribe	None, code only	Code builds her lines from the meeting's own counts, so she cannot misstate them	$0.00

How the last full meeting opened

"The cabinet is assembled, sir. 7 of 7 reporting, 1 guest seated." The scribe's slower voice opens and closes every meeting.

4. The room on screen

My first version showed the whole cabinet around a table from my seat. Five weeks later I came back to it and was not satisfied. I sent five screenshots of the diplomacy screen in the strategy game Civilization V and ruled that every advisor stays visible in the background and selectable, with one standing in front for whoever has the floor. That is the room today.



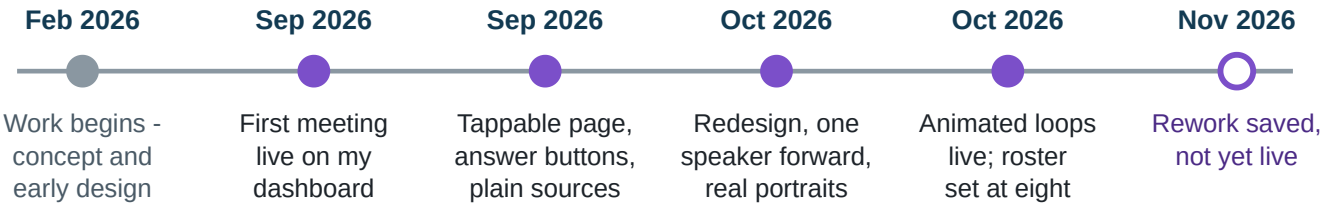
Layout of the live page (diagram, not a screenshot). Status words at the right of the name plate include ALL QUIET, DECISION WAITING, URGENT and EMPTY CHAIR.

Faces and motion. Each seat has one portrait, drawn with the Flux image model on my own graphics card and used both small (seated) and large (standing). Twenty short animated loops, one idle and one speaking per seat, came from the Wan video model. The speaker's mouth moves while its audio plays. Colors follow my ruling of blue, black and white, with yellow as trim only, and nothing I must read is smaller than 20 pixels.

5. Safeguards and quality checks

#	Safeguard	Mechanism
1	**Advisors have no tools**	Each AI aide receives a bundle of text and returns text. It cannot send a message, write a file or run a command.
2	**Facts come from scripts**	Plain read-only scripts gather the sources, with a 25-second limit and capped output. The AI writes only from that bundle.
3	**Empty chair over stale recital**	If a seat's main source is older than that seat's freshness limit, code seats it empty and shows when it last reported.
4	**Every claim shows its source**	Sources sit on screen beside the speaker in plain words. The advisor's own guesses carry a "My read" label.
5	**Money in percentages only**	A code filter rejects any money seat's brief that contains a currency amount. One rewrite, then an empty chair. Follow-up answers get the same check.
6	**Nothing executes**	An answer button records my pick in the minutes, marked sandbox, and the advisor says it back. No action leaves the room.
7	**Hard limits**	Code trims a brief past 90 words. A seat that misses its 240-second deadline gets an empty chair and never delays the meeting.

6. Results, status and what is next



Purple dots are live. Grey is design work in my notes. The open dot is built, awaiting my review.

Measured results

Measure	Result
First full meeting, September 2026	1 min 43 s to prepare, $1.51
Last full meeting, October 2026	1 min 49 s to prepare, $1.59
Five full meetings on record	$1.37 to $1.93 each
Spoken length, last meeting	195 s against the 240 s cap
Brief length, last meeting	55 to 88 words, cap 90
One follow-up question	$0.11 to $0.24

Status, stated plainly

Part	State
Meeting engine, voices, room page	Live since September 2026
Portraits, animated loops	Live since October 2026
Answer buttons	Live, record only
Reworked roster, personas, voices	Saved November 2026, not live
Living-room TV, spoken questions	Designed, not built

My verdict after the first round of testing was "decent progress, not there yet, it is a start." Each round since has come from using it and saying what was wrong.

What is next, in order. First, I retest the reworked version on a separate test copy. It gives each decision its own spoken card that says what happens if I do nothing. It lets an advisor take on small chores unless I stop it, and it brings a revised roster, personas and voices. Then the live page switches over. After that the room moves to the living-room TV, which first needs navigation by remote control. Then the voice pucks listen for my questions after each brief, and a confirmed answer reaches my lists and reminders.

7. What this shows about how I work

I ran this the way I would run any program with outside builders. I wrote each stage as a short brief with a goal and an acceptance test. AI agents built it, and I judged the result by using it, not by reading code. Two habits made the difference. First, I had a second AI model review the builds. On one round that review caught three layout defects the builder's own test missed, because the builder tested on a meeting with fewer seats. The rule now is to test on the busiest meeting. Second, I turned my preferences into enforceable rules - word caps, an empty chair for stale facts, percentages for money - so the system holds them without me watching.

That is the same work as adopting AI inside an organization. The hard part is not the model. It is deciding what people need to hear, in what order and how briefly, and building the checks that keep a confident voice honest.