

Powerbuilder

A multi-agent AI tool that helps nonprofits and campaigns build field plans grounded in real research, real data, and real organizing practice.

powerbuilder.app

BUILT BY

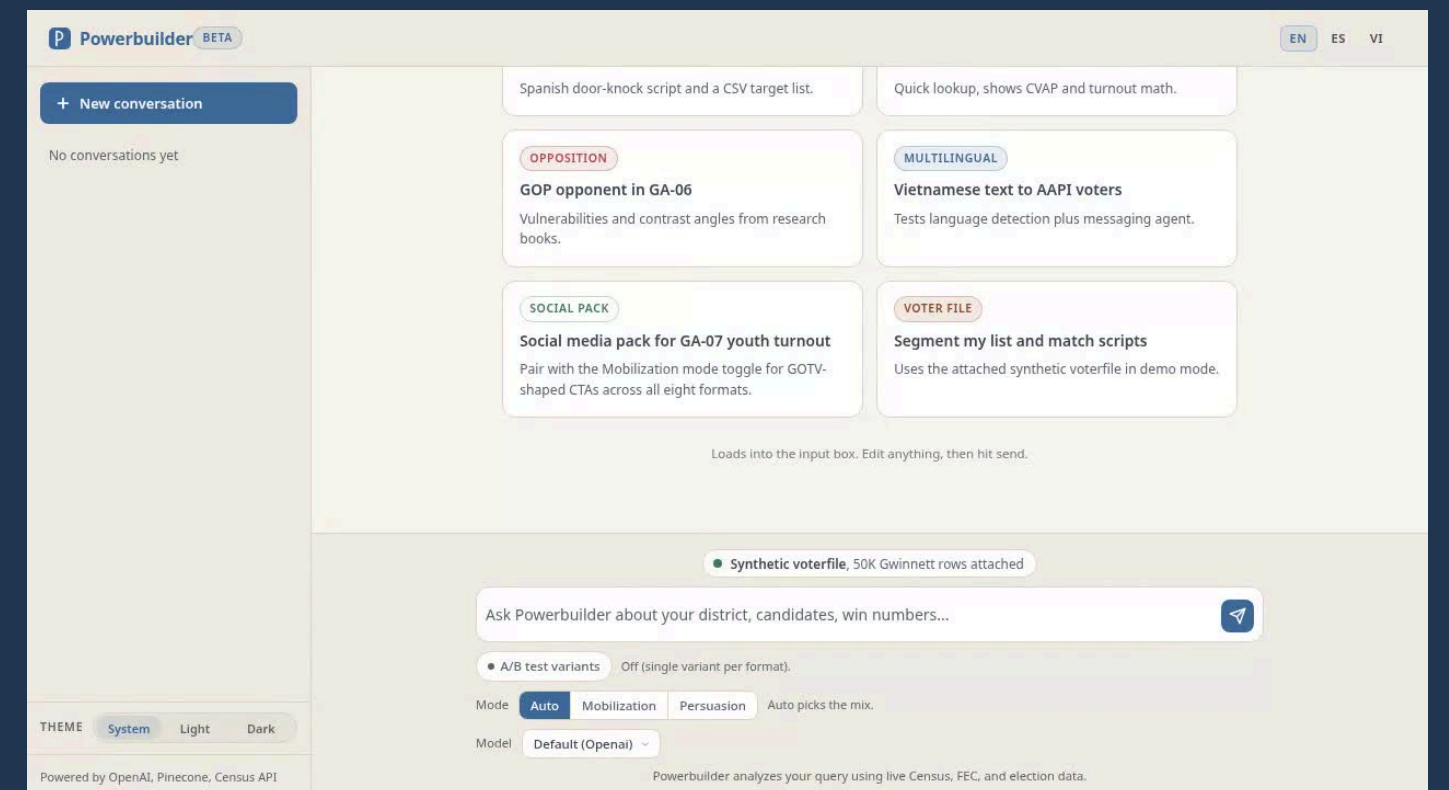
Benjamin Oh

architect, data + agents

Rosario Palacios

product, UX, testing

DxP (Re:)Power Partners



Bios



Benjamin Oh

- Senior Data Director – Alliance for Youth Organizing/Action
- 8 years of campaign/political organizing and data experience



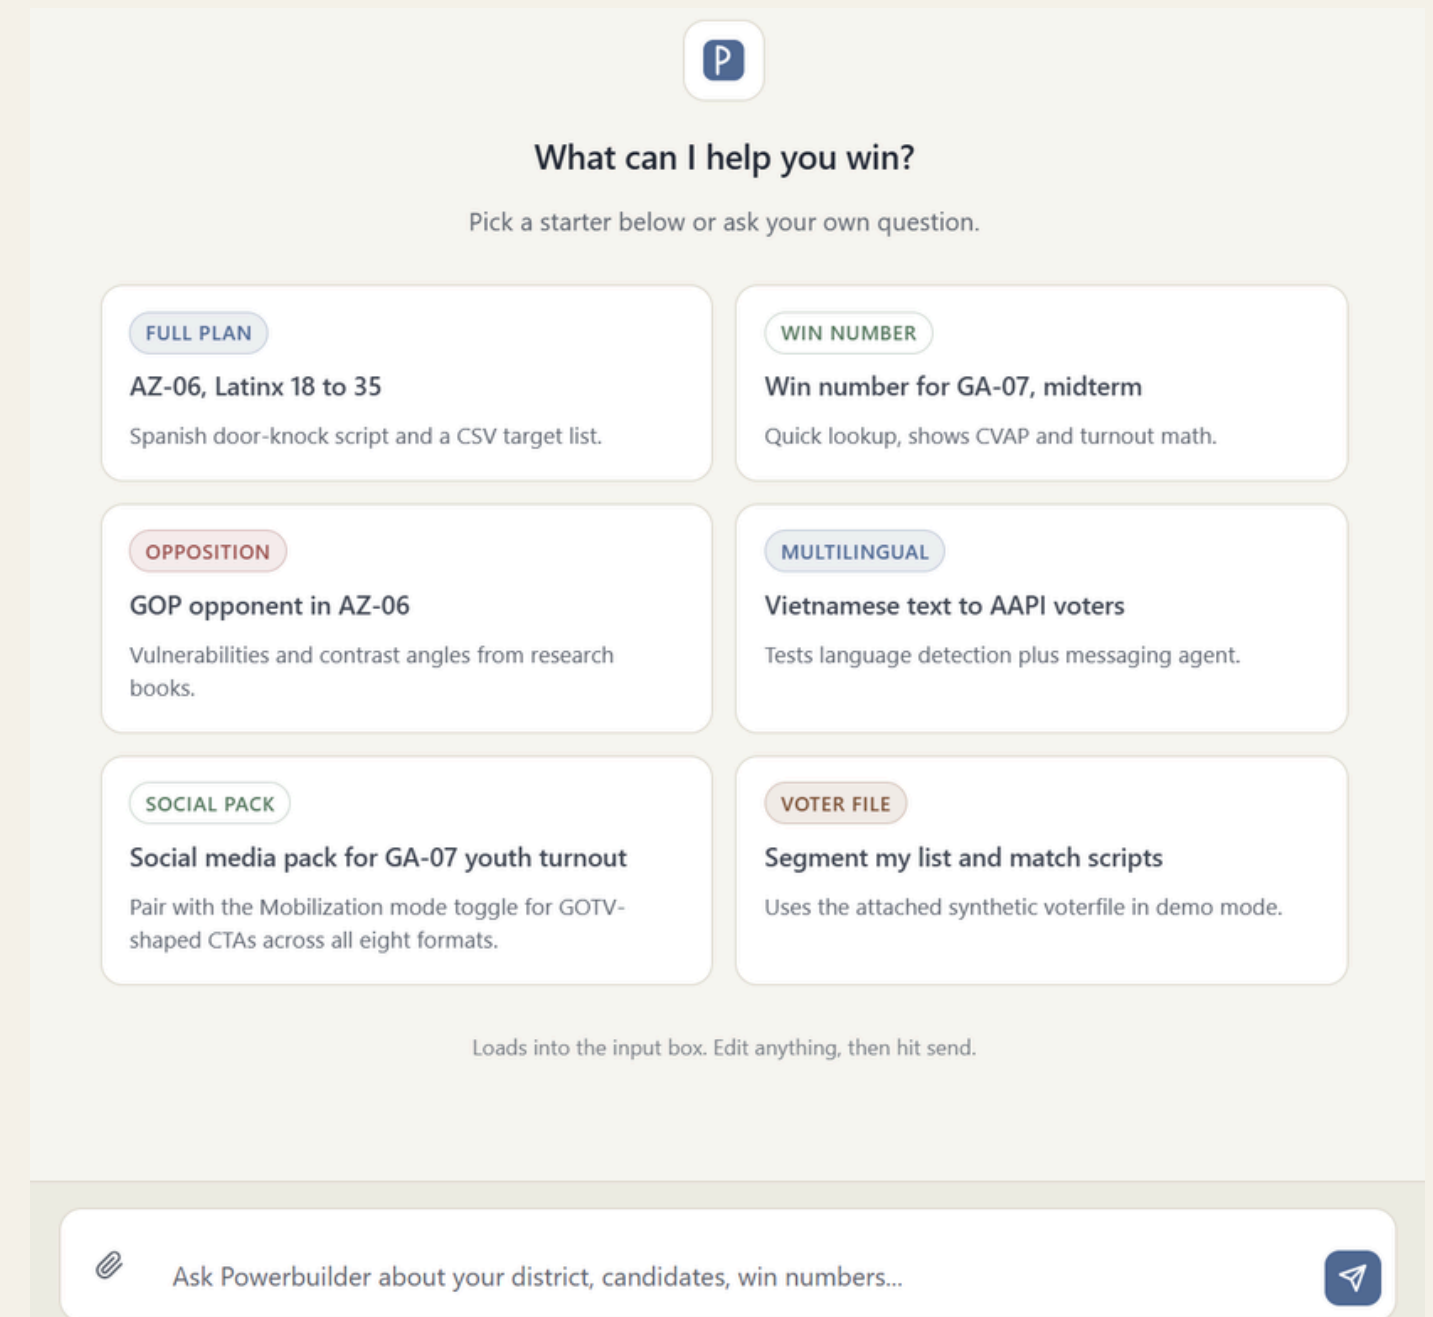
Rosario Palacios

- Principal, Palacios Contigo LLC
- 18 years organizing/political experience with 8 years data engineering/full stack dev.

Introducing: Powerbuilder

Powerbuilder is a human-centered agentic web app that turns weeks of data collection into an editable program plan draft in minutes.

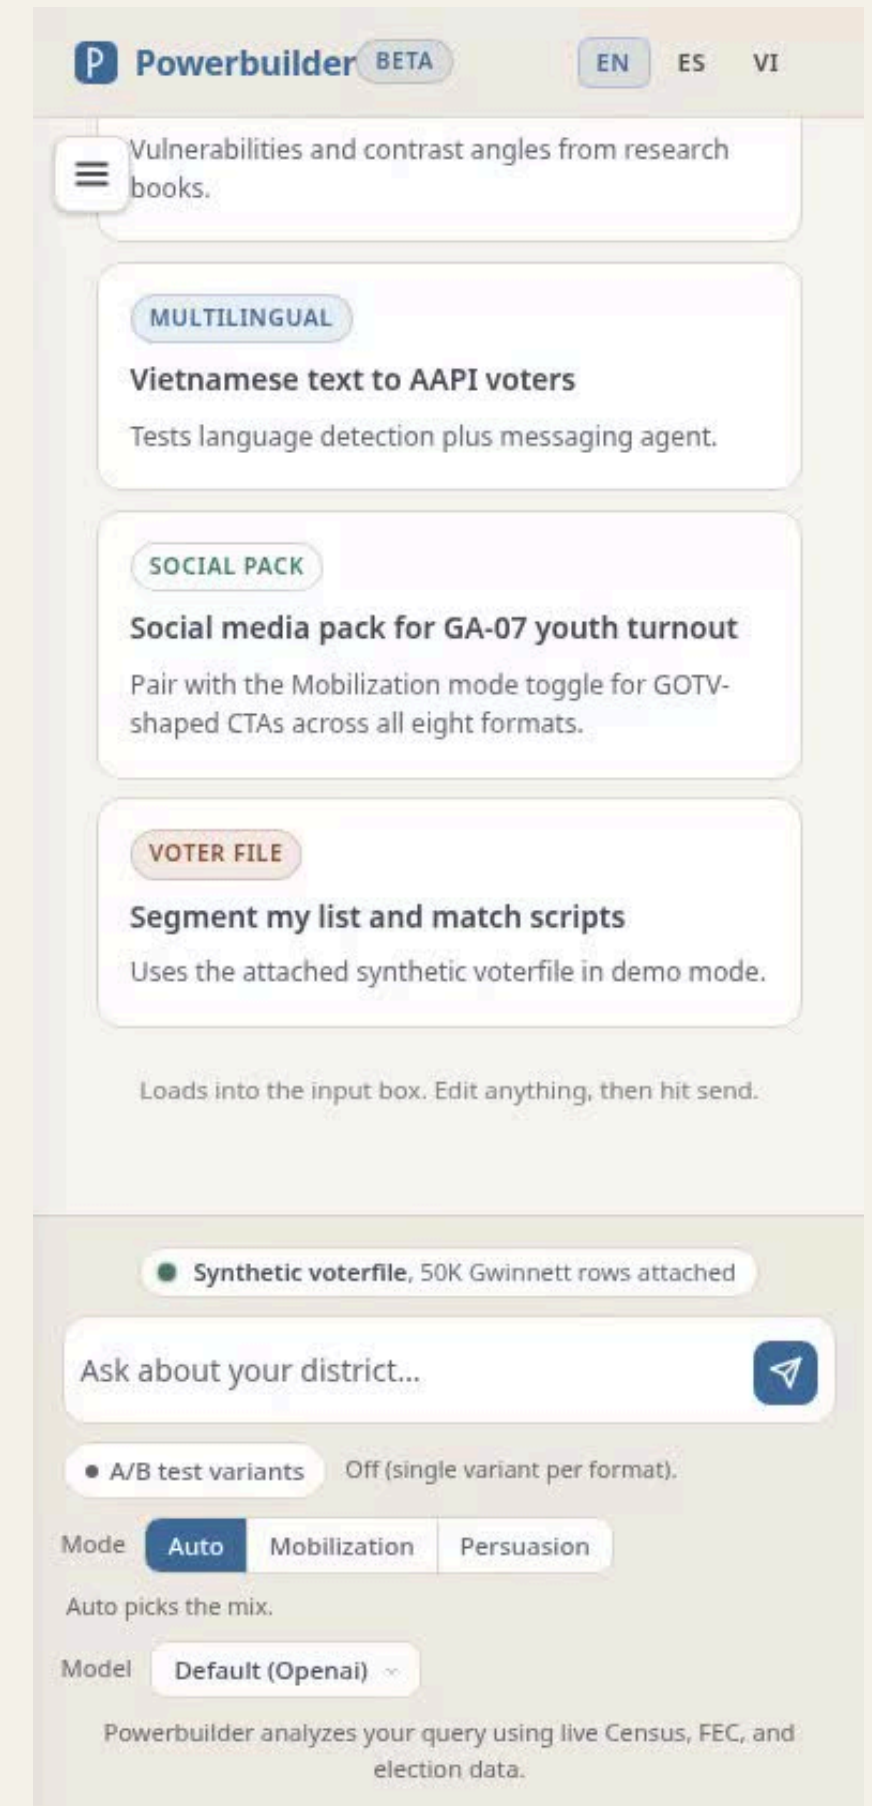
Built by movement for movement, Powerbuilder handles the compilation so your team can focus on what only you can do — telling the story of your organization and making the case for your program.



Imagine you were just hired as a Field Director in Arizona's 6th Congressional District.

It's August. You need a finalized GOTV plan for young Latino voters by Friday. What do you do?

You line up tabs of spreadsheets and tools. You call at least 3-4 of your staff or consultants to pitch in. You face a blank page and spend long hours stitching together numbers and ideas. Constrained by time, you might leave good data and research on the shelf or loosely rely on outdated information that won't meet the moment.



Watch what happens with Powerbuilder.

PROMPT

"Build a GOTV plan for Arizona's 6th Congressional District targeting Latinx voters age 18 to 35. Generate a Spanish door-knock script and give me a CSV of the target list."

POWERBUILDER'S TWO-CLICK SOLUTION

Enter the prompt. Hit send. Let it run while you check emails or grab a coffee.

The screenshot shows the Powerbuilder interface. At the top, there is a logo with the letter 'P' and the text "What can I help you win?". Below this is the instruction "Pick a starter below or ask your own question." There are six starter cards arranged in a 3x2 grid:

- FULL PLAN**: AZ-06, Latinx 18 to 35. Spanish door-knock script and a CSV target list.
- WIN NUMBER**: Win number for GA-07, midterm. Quick lookup, shows CVAP and turnout math.
- OPPOSITION**: GOP opponent in AZ-06. Vulnerabilities and contrast angles from research books.
- MULTILINGUAL**: Vietnamese text to AAPI voters. Tests language detection plus messaging agent.
- SOCIAL PACK**: Social media pack for GA-07 youth turnout. Pair with the Mobilization mode toggle for GOTV-shaped CTAs across all eight formats.
- VOTER FILE**: Segment my list and match scripts. Uses the attached synthetic voterfile in demo mode.

Below the grid is the text "Loads into the input box. Edit anything, then hit send." At the bottom, there is an input box containing the prompt: "Build a GOTV plan for Arizona's 6th Congressional District targeting Latinx voters age 18 to 35. Generate a Spanish door-knock script and give me a CSV of the target list." with a send button on the right.

Twelve specialized agents, real data, real plans

01

What is an agent?

A specialized program focused on one job, like pulling research, estimating a win number, or drafting a script. Each one writes its answer to a shared whiteboard.

02

A manager + 11 specialists

A manager agent reads your question, picks the right agents, runs them in order, and stitches the answer together. No black box. Anything from a sample script to a full plan.

03

Validated sources, not the entirety of the internet

Census, FEC, historical election results, Research Books, Cook Political Ratings, Pinecone database of memos & research curated by you. Human-discerned, not AI-crawled.

Build a GOTV plan for Arizona's 6th Congressional District targeting Latinx voters age 18 to 35. Generate a Spanish door-knock script and give me a CSV of the target list.

○ Searching the corpus
✓ Sources gathered

Researcher Electionresults Oppositionresearch
Winnumber Precincts Messaging Costcalculator

• Powered by OpenAI · gpt-4o

- 879-AI-GOTV-Mintt-Analyses-External-Final.pdf 2026-02-01
- Equis_2025 Elections and New Equis Polling_October.pdf 2025-01-01
- 2025-Civic-Power-Report-1.pdf 2025-01-01
- Equis Latino Research March 2026.pdf 2023-03-10

Make a real plan, and download it.

PRECINCTS + CENSUS

Real voting precincts, prioritized by your target populations

DOWNLOADABLE ASSETS

Every artifact is a named card with size and date. Three taps on a phone.

RESEARCH REFERENCES

Every response cites its sources, allowing audits without all the tabs.

OTHER FEATURE HIGHLIGHTS

Analysis of Your Universe

Upload your universe to center your plan around, or get quick summaries and suggested messaging

Prompts in Draft Plans

Use prompted questions throughout the plan to fill it with your org's mission and personality

Editable Prompts, Easy Regeneration

Edit or update your prompts, and regenerate responses from the app with one click

People-centered at every step.

RESEARCH, BACKED BY YOU

Training and research data curated by organizers and experts like you, not extracted from the open internet. Localized, so you are not paying for tokens you don't need nor left with hanging queries.

CITE YOUR SOURCES

Every response cites the memos behind its content. A reviewer can see what Powerbuilder prioritizes, trace content to sources and download documents to read.

ELIMINATE WRITER'S BLOCK, NOT THE WRITER

It's easier to edit than to create from a blank page. Powerbuilder sets you up to infuse a plan with your organization's personality and purpose, because you know your strengths and talents better than AI.

GET YOUR TIME BACK

Liberation of time, not replacement of people. We give you back time to focus on succeeding instead of compiling because we value your time and expertise. Fewer late nights and more capacity for what matters.

Mobile view, multilingual, and model picker.

ACCESSIBILITY · MOBILE · MULTILINGUAL

Field directors work on phones in Spanish and Vietnamese.

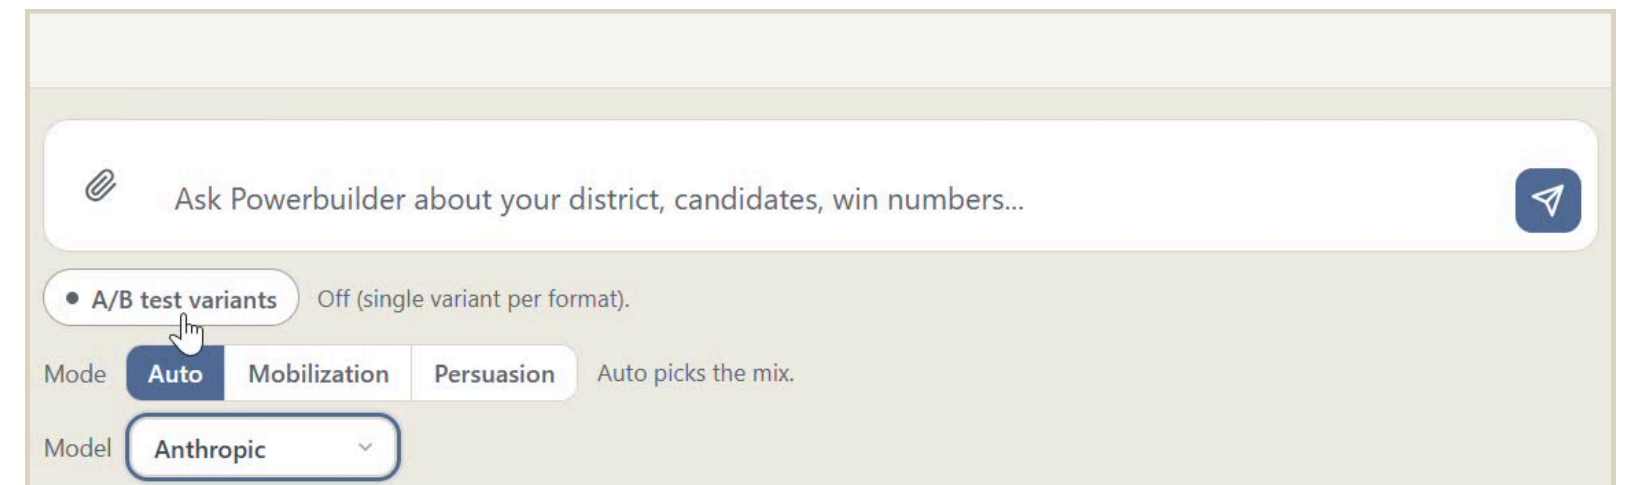
Mobile drawer, full ES + VI translations, focus rings, skip links, light + dark themes. In-language outreach is the easy path, not a buried setting.



MODEL PICKER

No single LLM is neutral. So we let orgs choose.

Seven providers, one default that works, a small chip on every answer that says which model wrote it. Auditable, not magic.



Next steps: There is no limit.

Partners

Pilot with progressive partners who want a values-aligned tool. Co-design, not co-marketing.

More APIs, More Data

Downballot races, expanded research, DDHQ, AP results, texting and phonebank tools, Voter file CRMs, better models. Live data, not stale exports.

More languages

Beyond ES and VI: Korean, Mandarin, Tagalog. Same content, no translation toggle hunt.

Org-level security upgrades

Org-level logins, Pinecone namespace isolation, user authentication, audit logs.

Appendix follows: architecture, agents in detail, model picker, accessibility, language audit.

Be in touch.

Benjamin Oh

benjaminjoh19@gmail.com

Rosario Palacios

rosario@palacios.community

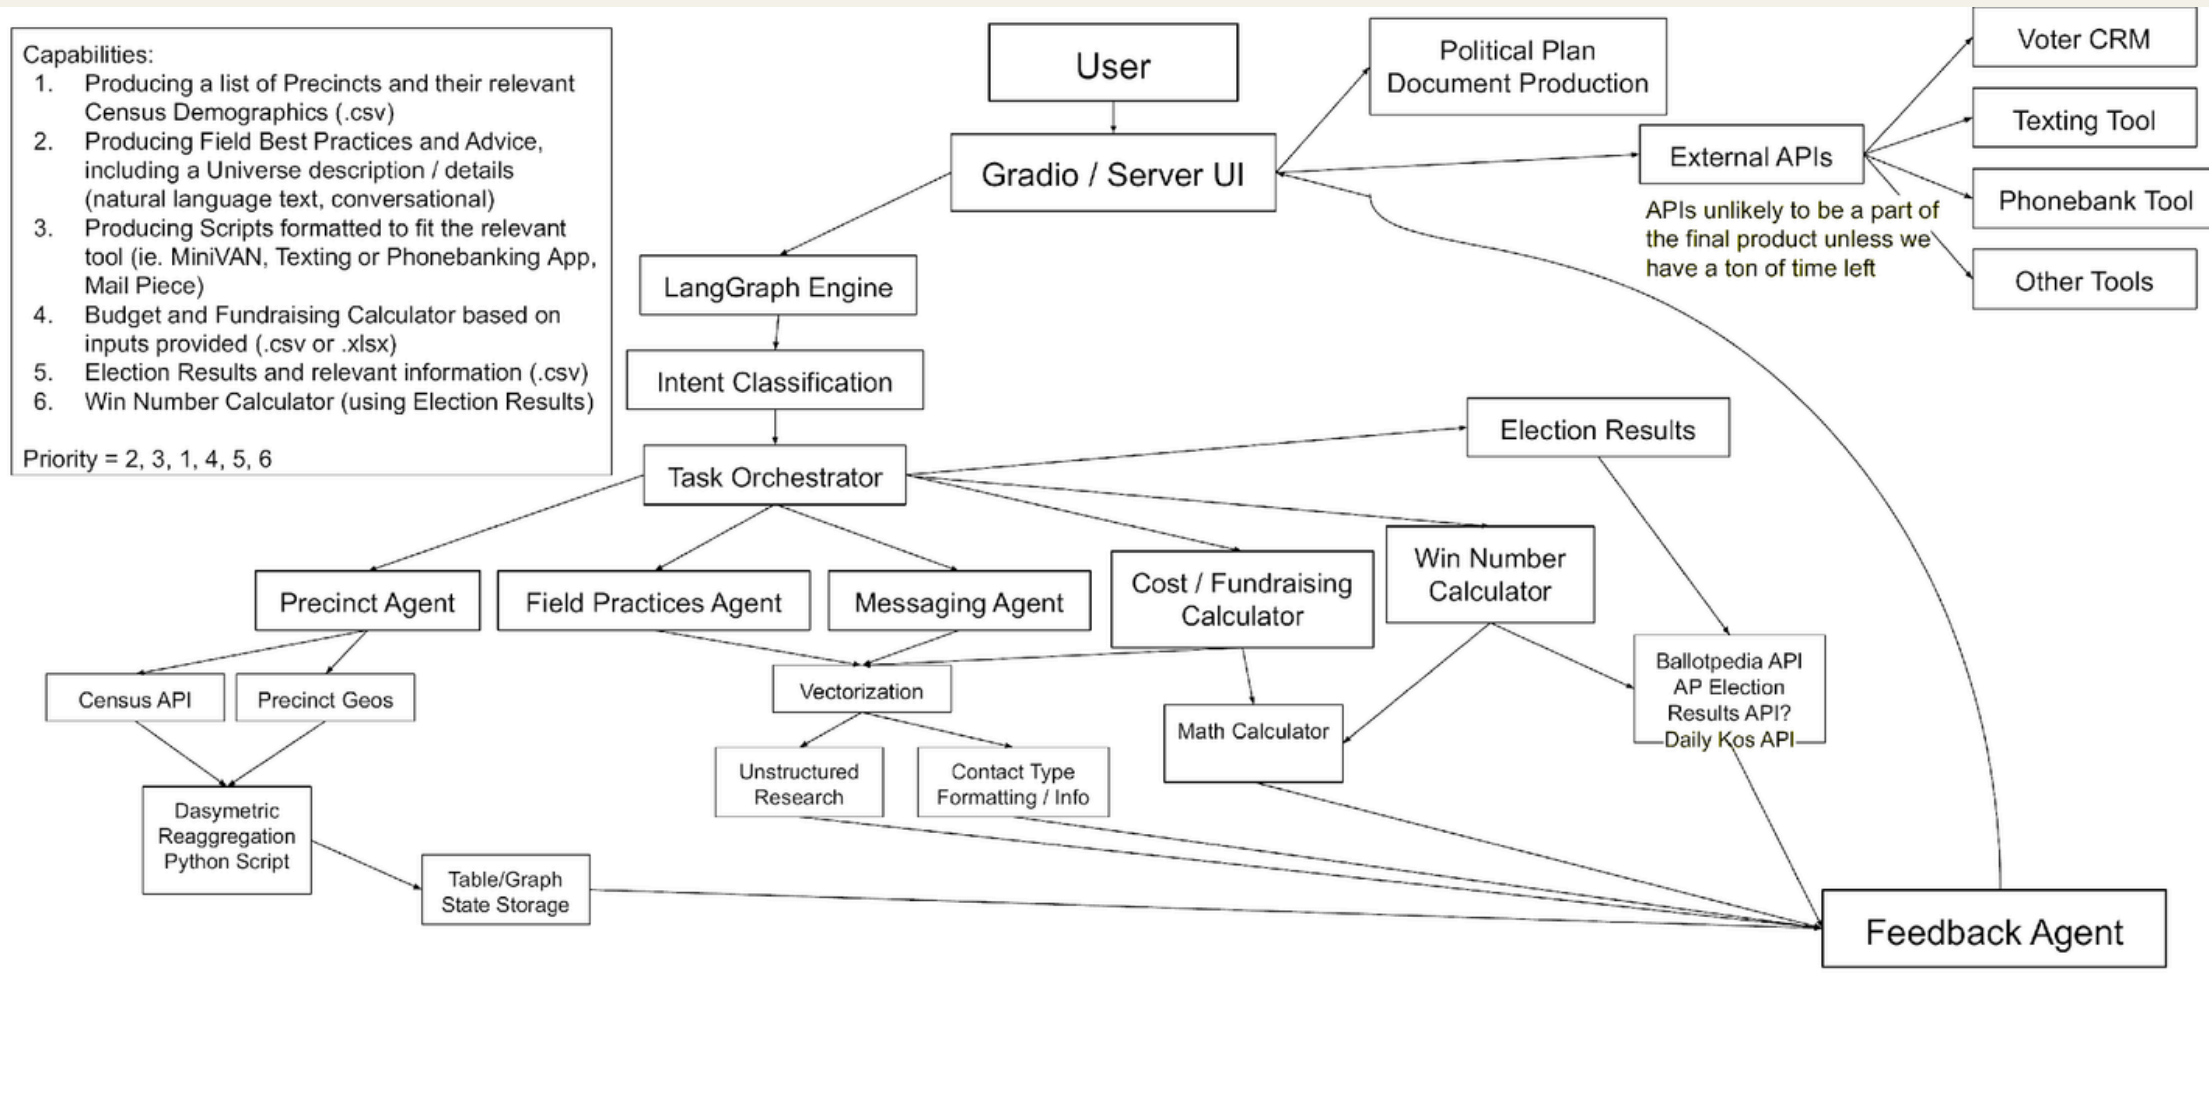
Thank you.

Powerbuilder · powerbuilder.app

Built by Benjamin Oh and Rosario Palacios for the DXP Fellowship.

User → orchestrator → specialists → real data → exported plan.

User asks a question. The manager picks specialists. Specialists pull from real sources. The export agent compiles the plan. Most of what was on the original sketch shipped, and some of the internal structure deviates from our original plan below.



- SHIPPED**
- Multi-agent orchestrator
 - Census + election data
 - Precinct + win-number math
 - Messaging in five formats
 - Plan document export

- SET ASIDE FOR LATER**
- Live VAN / EveryAction sync
 - Texting / phonebank tool APIs
 - Ballotpedia, AP, Daily Kos, DecisionDesk HQ
 - Feedback agent
 - User-uploads to org-level Pinecone database

Twelve specialists, real data, real deployment.

Powerbuilder runs on twelve specialist agents, each fulfilling a critical part of the whole app. Specialist agents are orchestrated by a Manager that reads the user's question, picks the right specialists, and runs them in order. State stitches the answer together and Export prepares the compiled response for delivery to the user.

Manager: orchestrator that picks specialists and coordinates their runs

Ingestor: uploads new research and documents to the Pinecone vector database

State: The whiteboard, where every agent writes its response before being passed to the Export

Export: Compiles and deduplicates specialist outputs, formats results into .docx or .csv files for downloads

Researcher: Extracts and compiles research based on relevancy from the Pinecone vector database

Contrast: Extracts information on opposing candidates from American Bridge's Research Books

Messaging: Uses Researcher data to create sample door, phone, mail, text and digital scripts

Voter File: Ingests and standardizes voterfile universes via CSV upload, advises on messaging

Precincts: Maps Census demographics on Voting Precincts using dasymetric reaggregation

Election Results: Sources historical election results to identify trends and classify district competitiveness

Win Number: Compiles CVAP and Election Results data to calculate a district target win number

Finance: Live FEC Data calls to inform program costs and calculate budgets based on program goals

Underneath: Census API with 24-hour cache, TIGER/Line shapefiles, district crosswalks across redistricting years, Cook Political Report ratings, American Bridge Research Books, MEDSL Election Results, FEC fetcher, S3, Pinecone, Django + nginx + systemd on a DigitalOcean droplet.

Safe demos, obvious downloads, mobile-first.

A · DEMO MODE

Pre-approved answers.

Curated, safe-to-show generations for rooms we can't afford to surprise.

B · DOWNLOAD CARDS

Named cards, size and date.

Three taps to a walk script or a full plan.
No scrolling chat history.

C · MOBILE DRAWER

Settings slide in.

Language as a top-level control. Mobile-first by default.

D + E · ITERATION

Rerun chips, copy, error chips.

Regenerate just the messaging. Errors say what failed, plainly.

F · SIDEBAR

Rename, reorder, dedup sources.

Plans evolve. Sidebar treats each plan as reopen-and-refine.

G + H · EMPTY + SOCIAL

Useful empty pages.

In-language tiles and sample prompts.
Platform-shaped social scripts (Meta, YT, TikTok).

No single model is neutral.

Different LLMs refuse different prompts, write in different voices, and price differently. A model trained to avoid "political content" can soften a turnout pitch. A model tuned for compliance can flag organizing language as activism. Your organization greenlights some LLMs but not others. So we built in the power to pick the model that works best for you and your program.

Contextvar override

Picker pins the provider for the whole run. Visible to every downstream call. No agent signature changed.

Powered-by chip

Every answer says which model wrote it. A small chip on the response. Auditable, not magic.

Choice log

Every pick logged to `exports/model_choices.jsonl`. We learn which model serves which task best (no database needed).

One prompt, multiple outcomes

We prompted three major LLM APIs with the same request to compare how each model performed given the same task, resources and instructions. While all three successfully completed the task with the same win number, CVAP calculation, top 5 target precincts and election data, there were significant differences and trade-offs in their performances, detailed below:

Prompt: Write me a complete political program plan for Arizona's 6th Congressional District targeting young and Hispanic voters in 2026.

OpenAI

- **Model: gpt-4o**
- 9,552-character draft plan (1 minute, 17 second run time)
- Issue choices: clean energy, youth engagement (no details)
- Often vague and generic, with 1-2 sentences per section
- Did not name the actual date of the 2026 November Election
- Aggregated youth and Latino voters

Anthropic

- **Model: claude-sonnet-4-5**
- 43,380-character draft plan (6 minute, 23 second run time)
- Issue choices: healthcare, clean energy, economy, Social Security
- Most mathematically-robust and contextually nuanced
- Spoke about youth and Latino engagement separately

Gemini

- **Model: gemini-2.5-flash**
- 24,157-character draft plan (2 minute, 31 second run time)
- Issue choices: healthcare, clean energy, economy, Social Security, Epstein Files, abortion, water crises
- Sufficient detail, lacked some nuance on messaging and targets
- Spoke about youth and Latino engagement separately

Run of show: 6:00 total

Time	Slide	Speaker	What's on screen
0:00 – 0:30	1. Title + Vision	Benjamin	Title slide, ambient app loop
0:30 – 1:00	2. Imagine yourself	Rosario	Pull quote + mobile video loop
1:00 – 1:15	3. Kick off the demo	Benjamin	Live screen share, Tab 1 (Powerbuilder)
1:15 – 2:45	4. While it runs	Rosario	Agents · 12 specialists · sources
2:45 – 3:45	5. Result + tab-flip	Benjamin	Tabs 1 → 2 → 3 → 1 (Powerbuilder, ChatGPT, ChangeAgent)
3:45 – 4:15	6. Why ours is different	Benjamin	Four-tile narrative
4:15 – 4:45	7. Two other things	Rosario	Mobile + Model picker, two video loops
4:45 – 5:30	8. What's next	Benjamin	Partners · APIs · Languages · Security
5:30 – 6:00	9. Close	Both	Be in touch slide, take questions

An org-friendly partner ladder.

Pilot with progressive partners that already do voter engagement. Co-design, not co-marketing. We meet them in their tooling, their language, and their compliance posture.

TIER 1 · STATE c3s

In-state pilots

PARTNERS WE HAVE IN MIND

Common Cause GA, Alliance for Youth Organizing, Migrant Equity Southeast, Nonprofit VOTE state affiliates, State Voices and America Votes tables.

WHY THEY FIT

Already doing voter education and engagement. Smallest contracts, fastest feedback loop.

TIER 2 · IDENTITY c3s

In-community pilots

PARTNERS WE HAVE IN MIND

Voto Latino Foundation, APIAVote, Asian Americans Advancing Justice (AAJC), VoteRiders, Native Vote.

WHY THEY FIT

Multilingual is a feature for them, not an add-on. Their organizers already work in ES, VI, KO, ZH.

TIER 3 · INFRASTRUCTURE

Data + tooling

PARTNERS WE HAVE IN MIND

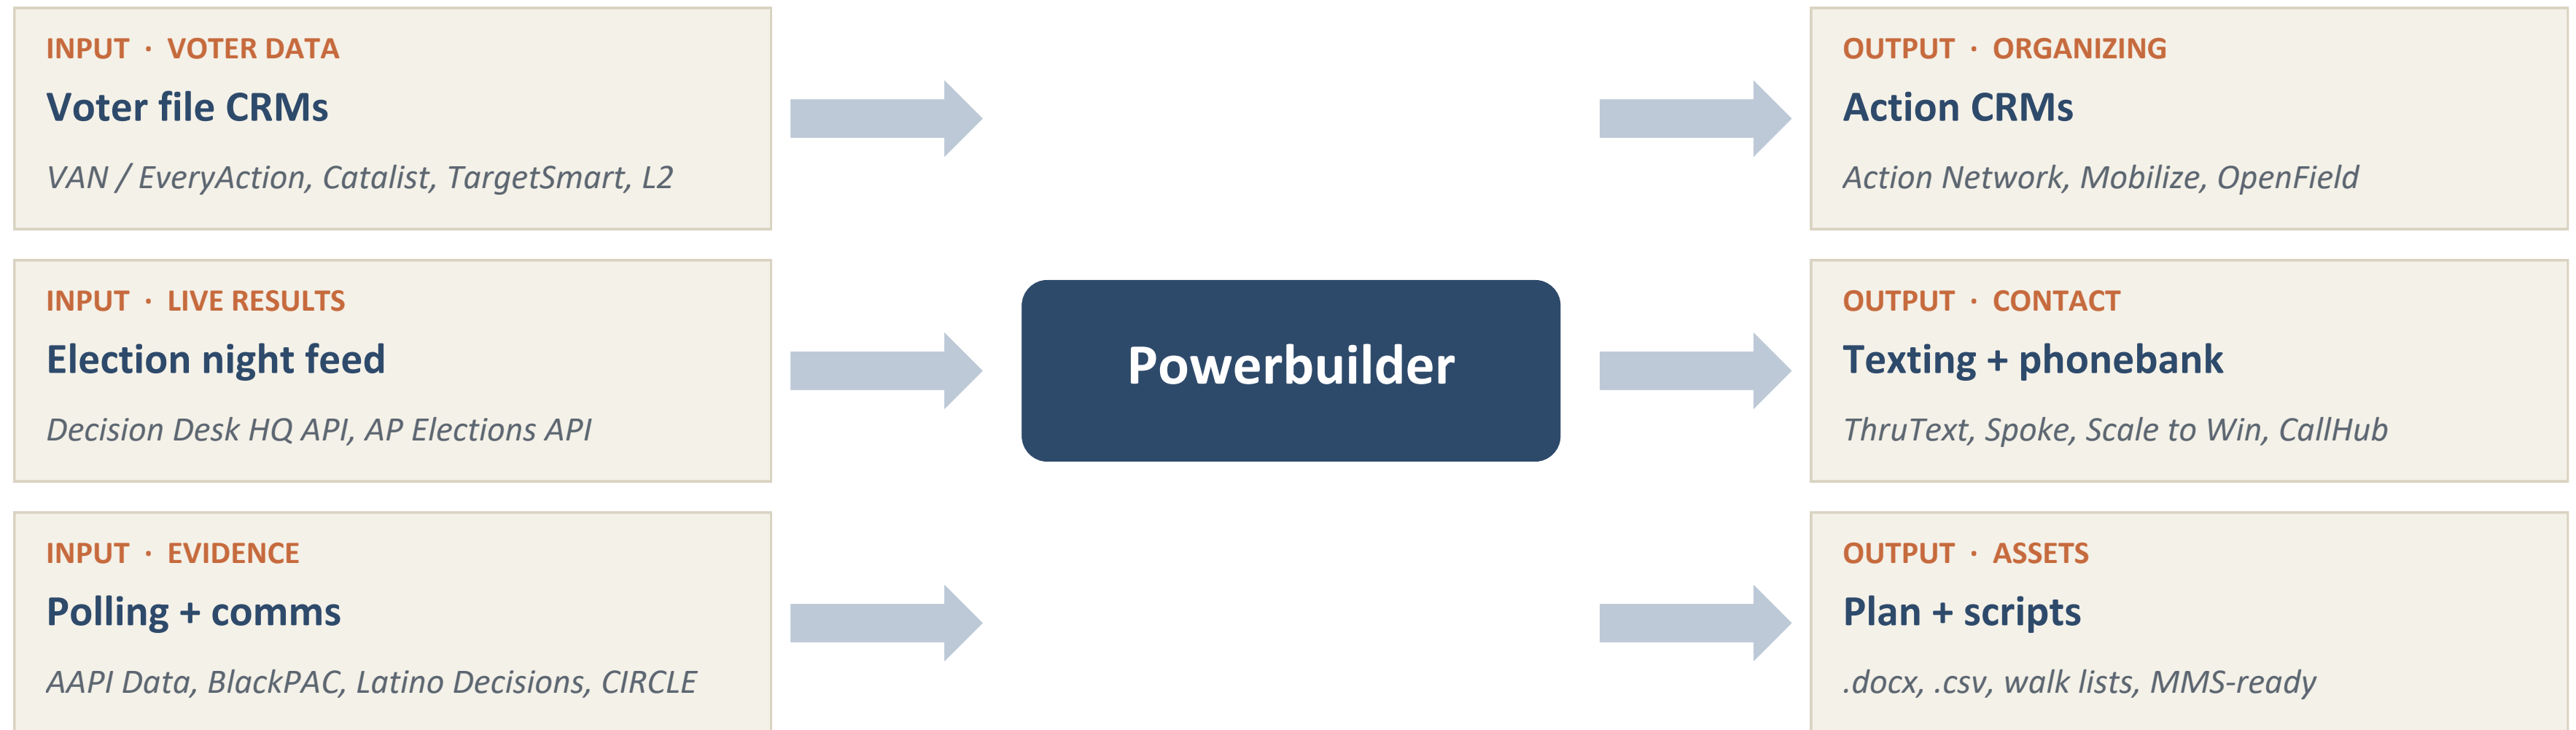
Catalist (data utility), The Movement Cooperative, Rock the Vote (civic tech), Civis Analytics (federated data).

WHY THEY FIT

They host the data layer we already touch. Integrating with them, not around them.

Powerbuilder, in the middle of the stack.

Today, Powerbuilder ingests data from the Census, FEC, Cook Political Report, a custom research corpus, and curated American Bridge files, then exports plans. Tomorrow, it reads from the CRMs organizers already use and writes back the universes, scripts, and assets it produced.



In-language by default, not as a setting.

Section 203 of the Voting Rights Act already requires bilingual voting assistance where language minorities cross thresholds. Our roadmap follows the same map: ship the languages organizers are already federally required to serve.

<p>SHIPPED</p> <p>EN</p> <p>English</p> <hr/> <p>WHERE IT MATTERS</p> <p>Default. Every plan, script, and asset.</p>	<p>SHIPPED</p> <p>ES</p> <p>Spanish</p> <hr/> <p>WHERE IT MATTERS</p> <p>AZ-06, NV, CA, TX, FL, GA. ~13% of voters nationally.</p>	<p>SHIPPED</p> <p>VI</p> <p>Vietnamese</p> <hr/> <p>WHERE IT MATTERS</p> <p>LA, OC, Santa Clara, Houston, Atlanta, NoVA, Seattle. Section 203.</p>	<p>NEXT</p> <p>KO</p> <p>Korean</p> <hr/> <p>WHERE IT MATTERS</p> <p>LA, NJ, NYC, Atlanta, NoVA. Required under Section 203 in Cook + LA Co.</p>	<p>NEXT</p> <p>ZH</p> <p>Chinese</p> <hr/> <p>WHERE IT MATTERS</p> <p>SF, LA, Chicago, NYC, Boston, Honolulu, Philadelphia, TX, WA. Largest LEP Asian electorate.</p>
---	---	---	---	--

Also on the list: Tagalog (FIL), Bengali (BN), Hindi (HI), Khmer (KM). Same content pipeline, no toggle hunt.

Every org gets its own walls.

Today the corpus is shared. Tomorrow each partner org has its own namespace in Pinecone to upload data or documents, its own SSO, its own audit log. One org's research can never leak into another org's plan.

TENANT ISOLATION

<p>Common Cause GA <i>namespace: ccga-prod</i></p>	isolated
<p>Voto Latino Found. <i>namespace: vlf-prod</i></p>	isolated
<p>APIAVote <i>namespace: apiavote-prod</i></p>	isolated

Shared Pinecone index · per-tenant namespace

UPGRADES ON DECK

SSO + per-user roles

Org logs in through Google or Microsoft. Admin, organizer, and reviewer roles. Auth, not honor system.

Audit log on every read

Every prompt, every source pulled, every export logged with user + timestamp. Reviewable by org admin.

Encryption + key rotation

At rest and in transit. Quarterly key rotation. SOC 2 Type II posture as the target, not the floor.

Right to delete

Offboarding drops the namespace, the audit trail, and the exports in one step. Clean exit.

Curated memos, not a crawl.

Our corpus is human-discerned: an organizer or researcher reads it, tags it, and uploads it. The roadmap is more sources, more recency, more accountability, never a scrape.

In-community evidence · partner-uploaded

AAPI Data, BlackPAC, Latino Decisions, Way to Win, Galvanize Action, CIRCLE, in-language qualitative + focus groups.

Movement research · partner-shared

American Bridge Research Books, Center for American Progress, Brennan Center, EPI, NDRC redistricting briefs, HIT Strategies, Equis

Ratings + reference · curated public

Cook PVI, Sabato's Crystal Ball, Daily Kos Elections data, Ballotpedia, MIT Election Lab, DecisionDeskHQ.

Foundational · open + structured

Census ACS + precincts, FEC, FCC, EIA, state SOS election results, Election Administration data, BLS, IRS Stats of Income.

THE INGESTION FLOW

- 1 Organizer reads + tags**
A human picks the source. Tags it by issue, geography, audience, recency.
- 2 Ingestor agent embeds**
Chunks the memo, embeds it, writes to the org's namespace.
- 3 Researcher cites at query**
Plans link to the exact memo + page. Auditable.
- 4 Reviewer flags or retires**
If a memo goes stale, a reviewer retires it. Versioned, not deleted.

From "give me a plan" to a real back-and-forth.

Today, Powerbuilder answers one prompt at a time. Next, it asks back. Clarifying questions, partial drafts, in-line edits. The plan becomes a conversation, not a download.

MOCKUP · CLARIFYING TURN

Write me a plan for AZ-06, young + nonwhite voters.

you · 10:42

Two quick questions before I run:

1. GOTV only, or persuasion + GOTV?
2. Is your universe under 18 to 35, or under 30?

powerbuilder · 10:42

agents thinking...

WHAT WE ARE BUILDING

● Clarifying questions

Agent asks before answering when the prompt is ambiguous. Two questions max, never a quiz.

● Inline edit + rerun

Highlight a paragraph, say "make this Spanish-first" or "shorter, for a text". Just that paragraph reruns.

● Plan memory across turns

The plan is the conversation. Refining one piece never blows away the rest.

● Voice + multimodal in

Drop a flyer photo, paste a poll PDF, or speak the prompt on a phone. Field organizers do not type fast.