
The First Corridor Pilot Plan

Will Blair

constellate.science/whitepaper

Companion to The Constellate Architecture · v0.1 · May 26, 2026

ABSTRACT

This pilot plan is the operational companion to *The Constellate Architecture* — the technical whitepaper for an open-protocol stack for cumulative scientific state. The whitepaper specifies the protocol; this document specifies the first deployment.

The plan names the candidate participants, the budget envelope, the regulator counterpart, the credentialing-committee composition, the conditional-deposit grant clause, and the explicit gating conditions under which the project will accept funding. As of this document's release, every name below is a candidate in active conversation; none is yet under signed letter of intent. The plan is written so funders, host candidates, and prospective participating labs can read the proposal against real institutional categories rather than abstractions, and so the project can be held accountable for converting candidates into committed participants within a defined window.

The pilot plan evolves on its own clock as host conversations progress; the technical whitepaper evolves on its clock as the protocol matures. Cross-references between the two documents are by section identifier; the whitepaper is the normative source for protocol mechanics.

0.0.1 About operational mode

The pilot plan describes a future institutional corridor with a foundation host, a reviewer credentialing committee, and a regulator counterpart. As of v0.1, the corpus and technical work are running ahead of that institutional structure in solo-maintainer mode: the maintainer compiles, reviews, and signs all updates personally, and all claims are bounded accordingly. No treatment advice, target validation, clinical validity, external adoption, or scientific-discovery claims are made without returned external evidence. The institutional gates in §8 begin the transition from solo-maintainer to multi-actor; until those gates close, the corpus is read as source material plus solo-maintainer review, not as field-validated state.

1 Frontier

Early Alzheimer's disease treatment-response, framed as the translational-failure map across the anti-amyloid program, with therapy heterogeneity as the entry wedge. The corridor's central object is the failure record: bapineuzumab, solanezumab, semagacestat, verubecestat, lanabecestat, atabecestat, aducanumab's contested approval, and the partial redemptions through lecanemab and donanemab, read as a structured catalog of what the field learned

at trial-level cost and did not absorb into a record the next trial designer could inherit. The frontier covers ten source-expansion lanes: anti-amyloid trials, failed amyloid/BACE programs, ARIA and safety, APOE/subgroup response, biomarker-clinical coupling, tau/amyloid staging, real-world registry evidence, BBB/delivery constraints, model-system transfer, and endpoint/protocol validity. BBB-dysfunction sits inside this scope as one mechanistic lane rather than the whole frontier; APOE4 stratification remains a primary axis. The frontier is chosen because:

1. It has multiple actively contested mechanisms across the ten lanes (vascular-inflammatory, pericyte loss, tight-junction disruption, tau propagation, microglial state, metabolic dysfunction).
2. The data is heterogeneous in exactly the way the protocol is designed to coordinate — post-mortem human cohorts, longitudinal MRI, APOE4-stratified mouse models, organoid perturbation data, trial-outcome records.
3. Several recent high-cost clinical signals (lecanemab ARIA-E findings, donanemab efficacy by APOE genotype, the broader treatment-response heterogeneity question) make decision-relevant updates near-term plausible.
4. The community is small enough that three to five labs can constitute a meaningful coalition once institutional gates close.

As empirical context for this frontier choice: the 2025 AD drug development pipeline (Cummings et al., *Alzheimer's & Dementia: Translational Research & Clinical Interventions*, 2025) holds 138 drugs in 182 active clinical trials across 4,529 sites worldwide, with 50,109 participants needed to populate active trials. The pipeline targets 15 distinct CADRO pathophysiology categories; 33% of agents are repurposed drugs from non-AD indications; 80% of trial sites participate in only one trial. The fragmentation those numbers describe is the absorption problem in concrete form. Cross-trial learning that would let trial N+1 absorb the subgroup outcomes, biomarker-arm divergences, and dose-response data of trial N is structurally not happening at this scale. The corridor's frontier sits inside this active pipeline and is structured to make a measurable subset of that learning happen.

1.0.1 1.1 External validation

The OpenAI Foundation has announced a partnership with Arc Institute to build what the announcement calls a "causal map" of Alzheimer's, training AI models on large-scale organoid experiments that vary genetic and environmental risk factors. The announcement states the program's intent that "researchers can share their findings along the way for others to build on, and nominate mechanistically informed drug targets for further testing." See openaifoundation.org/news/ai-for-alzheimers.

That sentence names the problem this protocol addresses without naming the layer that solves it. Findings shared "along the way" require a substrate that turns activity into signed, replayable, dependency-tracked state. Without one, the sharing degrades into documents, dashboards, and periodic reports, the same forms that produce the absorption gap in the current paper system.

The OpenAI Foundation initiative is not a participant in this pilot and is not on the host candidate shortlist below. This corridor's scope, budget, and host pool remain as specified throughout this document. The announcement is cited because the program it describes (sharing along the way, mechanistically grounded target nomination, and multi-driver causal mapping at patient-subtype resolution) is the operational set a working Constellate corridor enables.

The relationship is upstream and downstream. Large organoid programs of the kind OpenAI Foundation and Arc Institute describe produce the validated reality-contact a substrate absorbs. The first corridor's Early AD treatment-response frontier sits inside that broader Alzheimer's causal map. The pilot would be one early test of whether the substrate makes the "share along the way" promise structurally true rather than aspirationally so.

1.0.2 1.2 Current state

The frontier is live as `vfr_06cfcbe7c449d86a`, cloud-canonical and hub-served at `vela-hub.fly.dev/entries/vfr_06cfcbe7c449d86a`. Pull locally with `vela registry pull vfr_06cfcbe7c449d86a`. As of this revision, the corpus carries roughly:

- 820 finding bundles
- 620 verified sources, climbing toward a 1,000-source v3 target (395 source slots remaining)
- 2,960 evidence atoms
- 826 typed DAG edges
- 420 accepted reviewed state transitions
- 50 contradictions and gaps
- 20 prediction-resolution records
- 7 replications and 80 artifacts
- 1 review return packet, solo-maintainer

All 420 review events are signed by the maintainer. No external review returns are recorded yet. The corpus substantially exceeds the §10 sandbox-demo minimum (30 finding bundles) and is climbing inside the §5 reviewer-event envelope (600–800 events over 24 months that the budget supports once a reviewer pool exists). Daily progress is tracked in the Vela repo under `docs/superpowers/active-goal-status.md` and the project factory at `projects/anti-amyloid-translation/factory/`.

1.0.3 1.3 Scale horizon

The pilot's near-term success criterion is the §9 decision-change gate, not a corpus size. The long-run scale a serious Alzheimer's frontier needs is closer to:

- 5,000 to 50,000+ finding bundles
- 50,000 to 2,000,000 evidence atoms
- 25,000 to 1,000,000 typed DAG edges
- 1,000 to 10,000 verified sources
- 1,000+ accepted reviewed state transitions
- 100+ predictions/resolutions with closed calibration

The current corpus sits two to three orders of magnitude below the field-grade target. The Vela frontier-completion v3 target inside the repo (1,000 accepted updates, 5,000 finding bundles, 50,000 evidence atoms, 1,000 sources, 25,000+ DAG edges, 300 contradictions, 500 gaps, 100+ predictions, 3–5 real external review returns) is the intermediate milestone between the current solo-maintainer corpus and the field-grade horizon above. The pilot plan does not commit to those numbers as a v0.1 deliverable because hitting them requires the institutional structure the §8 gates establish; they are noted here so the current corpus is read as a starting state, not a finishing line.

1.0.4 1.4 Realistic ambition

The historical model for what cumulative scientific progress in a disease domain actually looks like is cardiovascular disease. US age-standardized cardiovascular mortality fell roughly 75% between 1950 and 2023 (Dattani, *Our World in Data*, 2025). No single intervention produced this. Drugs (statins, beta blockers, ACE inhibitors, clot busters), devices (pacemakers, stents, AEDs), surgery (bypass, valve replacement, transplant), diagnostics (echocardiography, CT, MRI), emergency response (911 systems, CPR, public AEDs), and lifestyle changes (smoking decline, routine blood pressure screening, dietary policy) layered over seventy years to produce a civilizational mortality decline. The cardiovascular field accomplished this absorption without a substrate. The next generation of equivalent declines, for AD and other diseases, can be compressed by building the absorption layer as infrastructure from the start rather than as decades of cultural practice.

The pilot does not promise to cure Alzheimer's. The honest realistic ambition is to compress the same trajectory: make the layered absorption of validated findings that the CVD field accomplished over seventy years happen for AD on a meaningfully shorter timeline (twenty to thirty years is the plausible target with substrate plus AI plus the current AD pipeline scale). Lecanemab and donanemab are the first modest inflection points in that trajectory. The next thirty to fifty inflection points, each independently modest, are what the absorption layer is built to enable. The corridor's value is to be the first proof that the absorption layer for AD can be built and operated. Everything beyond that runs at the pace biology and institutional adoption allow.

The corridor's scope is bounded against the kind of "end of disease" claim Lowe (2025, *Science*) and other working scientists have rightly critiqued. Machine learning rearranges existing evidence; it does not generate new biology. The substrate's contribution is to make the new biology that does get generated, when it gets generated, cumulative across institutions and time. That contribution is bounded but real, and it is the operation the documented seventy-year CVD trajectory shows actually produces civilizational mortality decline.

2 Host Candidates

Cure Alzheimer's Fund, the Alzheimer's Drug Discovery Foundation, BrightFocus, the Chan Zuckerberg Neurodegeneration Initiative, or an FRO chartered for this purpose (Convergent Research has indicated openness to such a chartering conversation). The host's commitments are: convene the lab coalition, underwrite reviewer labor at the corridor's defined rate, adopt the conditional-deposit grant clause (§7), and own the corridor-level governance policy through the pilot period.

3 Target Lab Candidates

The corridor will recruit from the active NIH-funded neurovascular Alzheimer's research community, with candidate labs working on the relevant phenotypes (post-mortem neurovascular phenotyping, longitudinal BBB-permeability imaging, APOE4-stratified pericyte and tight-junction biology, multi-modal connectomics with neurovascular endpoints). Specific PIs are not named in this document without their prior written consent; the project commits to publishing the named candidate roster only after each named PI has been notified of inclusion and has affirmed willingness to be publicly listed. Public disclosure of *committed* participants

under signed LOI follows within 90 days of host commitment. This sequencing — notification before publication, LOI before commitment — protects target PIs from being named as candidates in a public document before they have agreed.

4 Regulator Counterpart

The FDA's Center for Drug Evaluation and Research neuroscience division, engaged through a Type C scientific advice meeting under FDA's formal-meeting program. The engagement does not seek protocol blessing or agency-wide acceptance; it seeks confirmation that, for the corridor's first IND-enabling or real-world-evidence submission, signed proof packets are an acceptable form of supporting evidence alongside existing submission materials. A Type C meeting is a corridor-specific instrument, not a general policy lever, and this document will not claim agency-wide acceptance from one meeting outcome. EMA engagement follows for cross-jurisdictional applicability.

5 Budget Envelope, 24 Months

Primary number: \$750K total. The envelope can stretch to \$900K (plus \$150K) if a fourth lab joins or if TypeScript-reducer remediation requires contracted engineering, and can compress to \$600K (minus \$150K) if the depositor-side staff line is absorbed by participating institutions through existing grants — but the project will not pursue compression by cutting the depositor-side line in isolation.

Allocation:

- **\$200K reviewer-pool underwriting.** Per-event payments at \$250 per substantive review event (calibrated against NIH study-section half-day honoraria of roughly \$200/half-day), with higher rates (\$400–\$600) for high-stakes merges involving active trials, animal studies, or manufacturing transitions. Envelope funds approximately 600–800 reviewer-events over 24 months.
- **\$150K infrastructure and host stewardship.** Vela instance operation, manifest signing infrastructure, transparency-log integration with at least two independent witnesses, sandbox demo at `demo.constellate.science`.
- **\$100K translation work.** Funds a 0.5–1.0 FTE corridor translator under the host's payroll, with a target hire date within 90 days of host commitment. The translator's job: convert canonical state into decision packets the regulator, foundation program officers, and clinical teams can act on (evidence packet generation, regulator-readable export tooling, decision-context briefings).
- **\$50–100K cross-implementation conformance testing.** Python and TypeScript reducer parity audits per the whitepaper's conformance regime.
- **\$100–150K depositor-side staff support** at participating labs. Typically a 0.2–0.4 FTE staff scientist or technician per lab over 24 months for structured deposit work (negative-result formatting, condition-record curation, trajectory step capture, *plus* research-administrator time to draft and route the IRB/IACUC protocol amendments required for deposit, which typically runs 20–40 hours per site). The single line covers both depositor-facing and research-administrator-facing work because both are required to make deposits actually happen. Without this line item, the substrate's burden falls invisibly on postdocs

whose time is already grant-encumbered, and the deposits become noise. This line is the difference between a pilot that runs and a pilot that drifts.

Lab work is funded separately through participating institutions' existing grant pipelines; the pilot does not pay for science, only for the substrate layer that records it and the staff support to do the recording properly.

6 Reviewer Credentialing Committee Composition

The credentialing committee that decides who enters the reviewer pool is the political surface that determines whether the corridor is captured at birth. The default composition specified for the BBB-dysfunction corridor: **five seats, all serving renewable 18-month terms**, composed of:

1. **Two PIs with substantial publication records in *competing* mechanistic camps** within the corridor (e.g., one vascular-inflammatory, one tight-junction), so no single camp dominates credentialing.
2. **One statistical reviewer** drawn from outside the corridor's mechanistic schools. Candidate pools include the Alzheimer's Disease Cooperative Study (ADCS) statistical core, biostatistics groups at ADRCs whose primary work is on AD trial design rather than mechanistic biology, and the ATRI (Alzheimer's Therapeutic Research Institute) statistical group. This seat may take longer to credential than the others because the pool of mechanistic-camp-unaligned AD statisticians with substrate fluency is small.
3. **One provenance/methods reviewer** (typically from a research-integrity office or an open-science methods program).
4. **One early-career representative** from the credentialed reviewer pool below the PI tier. *Non-voting* on credentialing decisions involving PIs from their own training lineage or current institution; *voting* on early-career credentialing decisions. The non-voting restriction protects the early-career representative from the career risk of being asked to gate access for senior PIs in their own field.

Conflict-of-interest disclosure follows NIH study-section practice: financial disclosures, recent collaboration disclosures, and recusal from credentialing decisions involving direct co-authors within the prior three years. The committee roster is public; rotations are recorded as governance events.

7 Conditional-Deposit Grant Clause

The conditional-deposit mechanism — milestone capital tied to signed state transitions — is the single most operationally novel commitment in the adoption path. The clause is structured as four obligations on the grantee, four obligations on the grantor, and three named operational mechanisms.

Grantee obligations.

1. The grantee shall, at each milestone defined in the grant schedule, produce signed state transitions under their registered actor record for findings, negative results, replications, or trajectories agreed at grant award.
2. The grantee shall use the corridor-designated Vela instance and shall sign under keys registered to the grantee's institutional actor record.

3. Deposits shall conform to the Carina kernel version pinned at award time; kernel migrations during the grant period are governed by the corridor's RegistryGovernancePolicy.
4. Trade-secret-sensitive content may be deposited under Restricted or Classified tiers per the protocol's access-tier model; the existence and signing of the transition is public, the underlying content may be tiered.

Grantor obligations.

1. The grantor shall release milestone payments within 14 days of milestone-event signature verification by the corridor's reviewer quorum.
2. Disputes about whether a milestone deposit meets specification are resolved by the corridor's named arbiter panel (see below).
3. A 30-day cure period follows any non-deposit finding; the grantee may remedy the deposit or contest the finding before payment is withheld.
4. Withheld payments do not retroactively void prior milestone payments; the grant continues unless the grantor exercises termination under separately negotiated terms.

Operational mechanisms.

1. The corridor maintains a public list of grant-clause-bound actor records, so any reader can verify which deposits were grant-conditional.
2. The corridor's arbiter publishes anonymized dispute summaries quarterly so the clause's operation is inspectable.
3. The clause includes a sunset and renegotiation provision: after 24 months of operation, host and grantees jointly review the clause's mechanics and may amend it through a governance event.

Default arbiter panel. The arbiter is a three-seat rotating panel, not a single person. The default composition: three senior reviewers credentialed in the corridor, drawn by rotation from a standing pool of at least nine eligible reviewers, with the rotation excluding any reviewer affiliated with the grantee institution or the grantor's program for that specific dispute. Panel composition for a given dispute is determined by deterministic rotation order, recorded as a governance event, and posted to the corridor's transparency log; the grantee and grantor each have one peremptory challenge per dispute. The panel's decision is by simple majority and published as a signed `dispute.resolved` event with reasoning attached. If the standing pool falls below nine eligible reviewers, dispute resolution pauses pending pool replenishment rather than collapsing to a single arbiter.

This clause template is provisional and explicitly invites counsel review by hosts adopting it. The protocol does not specify the legal form of the clause — that is properly the host's work — but commits to the operational mechanisms above as the minimum surface a host would need to make the clause auditable.

8 Gating Commitments

The project will not accept funding above the pilot envelope until:

1. **Steward gate (90 days).** The external steward is named per the whitepaper's §10.4 commitment.
2. **Coalition gate (180 days).** At least one host and three labs are under signed letters of intent.

3. **Grant-clause gate (180 days)**. The conditional-deposit grant clause (§7 above) is reviewed by counsel and adopted by the host.
4. **Conformance gate (270 days)**. The cross-implementation conformance pass rate reaches a defined threshold across Rust, Python, and TypeScript reducers.
5. **Decision-change gate (24 months)**. The first signed state transition produced under a real grant condition closes a real downstream decision (defined below).

The pilot is engineered to fail visibly if any of these gates are missed.

Why this ordering — steward first. The sequencing above leads with the external steward, not with named host or lab commitments. This is deliberate. A steward gives funders and host candidates a counterparty that survives the project's individuals; without one, every committed host is implicitly committing to a single-author project. Several host candidates have specifically asked for steward identification before discussions move from informational to commitment-shaped. The 90-day steward gate is therefore the unblocking move for the entire downstream chain.

9 Terminal-Gate Definition: "Real Downstream Decision"

Gate (5) — the 24-month commitment — requires that "the first signed state transition produced under a real grant condition closes a real downstream decision." The qualifying decision categories are:

- A foundation program pausing, redirecting, or terminating a funded program citing the transition.
- A regulator citing the transition in a submission decision (acceptance of evidence, request for additional data, conditional approval modification).
- A hospital or trial network modifying an active protocol citing the transition.
- An NIH study section recommending allocation against a posture the transition supports or undermines.
- A participating lab pausing or pivoting an active program citing the transition.

Anonymous "we considered it" verbal acknowledgments do not satisfy the gate; the qualifying decision must be documented in writing and referenced by the citing institution. The corridor's arbiter panel certifies satisfaction of the gate.

10 Sandbox Demo

The Early AD treatment-response frontier is already live as `vfr_06cfcbe7c449d86a`, hub-served at `vela-hub.fly.dev`, reachable through the Vela CLI (`vela registry pull vfr_06cfcbe7c449d86a`). The current corpus (820 finding bundles, 620 sources, 2,960 evidence atoms, 826 DAG edges, 420 solo-maintainer-signed state transitions; see §1.2) substantially exceeds what the v0.1 plan had committed to as a minimum sandbox surface.

The remaining commitment for `demo.constellate.science` is a public web read-side over the live frontier, standing up within 60 days of this document's release: a curated subset rendering with at least one signed correction propagating to a dependent finding visible end-to-end, one downloadable proof packet, and one signed `frontier.conflict_resolved` event illustrating maintainer arbitration per the whitepaper's §4.6. The sandbox does not carry canonical state separate from the live frontier; it is the public read-side of the live frontier under the

demonstration-infrastructure label. Its purpose is to give funders, regulators, journalists, and prospective participating labs a click-through artifact rather than a textual claim, and to make the corpus inspectable without requiring the reader to install the Vela CLI.

Because the corpus is live, the current `demo.constellate.science` commitment is a UI and curation task, not a corpus task. The 30-bundle minimum from earlier drafts is retained only as a floor for the curated view; the actual corpus exposed will be substantially larger.

11 IRB and IACUC Interaction

Structured deposit of human-cohort or animal-model trajectory information may require modification to participating institutions' existing IRB or IACUC protocols. The project commits to working with each participating institution's research-integrity office to determine the minimum protocol modification required, and to publishing a template IRB protocol amendment for human-cohort deposit and a template IACUC modification for animal-model deposit within 90 days of first-lab commitment. The expected modifications are narrow (adding the corridor's signed-deposit mechanism to existing data-sharing provisions) but require institutional review at every participating site. Research-administrator time to draft and route the amendments is covered by the depositor-side staff support line in the budget envelope. This is real friction the pilot will hit, named explicitly so the timeline accommodates it.

12 Pivot Criterion

If the BBB-dysfunction corridor's decision-relevance hook collapses before 12 months — for example, if the lecanemab and donanemab efficacy questions resolve cleanly via traditional means and the corridor loses its "what decision moves" target — the pilot will pivot rather than run a corridor without a closable loop. The alternative-frontier candidates under preliminary consideration:

- Pediatric high-grade glioma perturbation-response state (host candidates: Pediatric Brain Tumor Foundation, Alex's Lemonade Stand, Children's Brain Tumor Network).
- Tauopathy-spectrum disorder differential diagnosis (host candidates: CurePSP, Rainwater Charitable Foundation).
- Pathogen-surveillance state for an emerging infectious disease (host candidates: CEPI, the Wellcome Trust, or the Open Philanthropy Biosecurity portfolio).

A pivot decision is itself recorded as a corridor governance event and is contestable by participating labs and the host; pivot is intended as a discipline against running theater, not as an exit from the protocol's commitments.

13 Early-Corridor Non-Colluding-Depositor Relaxation

The protocol's privacy floor requires condition records to be aggregated across non-colluding depositors with at least one outside the host's funding portfolio. This binds tightly when the corridor has three to five labs all working in a small funded community — in the BBB-AD case, several of the candidate labs are likely to be CAF or NIH-funded. To prevent this from forcing all early-corridor deposits into Restricted tier by default, the corridor host may file a `non_colluding_relaxation` governance event certifying that the depositor pool is too small to satisfy the strict criterion. The relaxation is itself a transparency-log-witnessed event subject

to maintainer-quorum review, time-bounded to a 12-month window, and reviewable at corridor maturity.