DISTRICT CIVIC INTELLIGENCE DASHBOARD

A Civic Intelligence, Policy Measurement, and Responsible AI Framework

Philadelphia 4th District Demonstration Model

PUBLIC REVIEW DRAFT | VERSION 0.9 | JULY 2026

Produced and Conceptualized by C. Michael Robinson

Product Owner | Founder, Troubadour Solutions

Prepared for public peer review through cmichaelrobinson.com

Publication and Review Notice

This document is a condensed public-review edition of the District Civic Intellegence Dashboard Master Artifact Compendium. It preserves the project thesis, architecture, data model, governance standards, demonstration scope, and core acceptance requirements while removing repetitive artifact history and operational detail that is unnecessary for public technical review.

What this edition intentionally omits

Private credentials, exact exploit-detection rules, break-glass procedures, unpublished client strategy, sensitive administrative configurations, detailed contractor assignments, and other information that could reduce platform security or disclose confidential business operations.

The project is owned and developed by Troubadour Solutions. The Philadelphia 4th District is the initial demonstration model. It is not an official City of Philadelphia or election-administration system. No open-source license is granted by publication of this document.

Reviewers are invited to submit comments on technical feasibility, methodological integrity, geographic design, responsible AI, accessibility, privacy, security, and public usefulness. Feedback may be submitted through the contact mechanisms at cmichaelrobinson.com.

Reading convention

Statements describing the proposed system are design requirements unless explicitly identified as implemented, demonstrated, or validated. The platform remains in pre-production development.

Contents

1. Abstract and Executive Summary

Abstract

The District Civic Intellegence Dashboard (DCID) is a proposed AI-enabled, geospatial decision-support platform designed to convert fragmented public information into measurable civic strategy. It organizes data by place, connects conditions to policy interventions, distinguishes implementation activity from resident outcomes, and requires every public finding to retain its source, geography, reporting period, methodology, and limitations. The Philadelphia 4th District serves as the initial demonstration model, but the system is designed as a reusable Troubadour Solutions product for campaigns, elected offices, civic organizations, advocacy institutions, and public-interest partnerships.

The platform combines a governed PostgreSQL/PostGIS warehouse, ArcGIS mapping and feature services, a custom web application, controlled APIs, and permission-aware AI tools. Its distinctive governance features include a nonpartisan analytical standard, a pre-generation Secondary Bias and Integrity Gate, source and methodology provenance, strict separation of public and internal data, and fail-secure controls for suspected exploitation. The central research question is whether these components can produce a civic intelligence system that is simultaneously useful for political decision-making, transparent enough for public scrutiny, and disciplined enough to resist partisan or external manipulation.

Executive summary

DCID treats geography as the organizing framework for evidence, policy, participation, and accountability.

The governing chain is: Condition -> Baseline -> Geographic concentration -> Contributing factors -> Intervention -> Target -> Measurement -> Public narrative.

PostgreSQL/PostGIS is the authoritative analytical repository; ArcGIS provides managed maps and feature services; public and internal applications consume approved publication views.

The first demonstration combines district, neighborhood, ward, division, corridor, polling-place, committee, and Census geographies.

The Policy Outcome Engine requires a baseline, target, deadline, geography, responsible entity, measurement source, and separate implementation and outcome statuses.

Public polling and civic research may be used only with sponsor, methodology, geographic fitness, and reuse-rights disclosure.

AI is restricted to approved tools and evidence packets. It cannot query arbitrary databases or treat inference as fact.

A nonpartisan standard requires equivalent evidence and claims to receive equivalent treatment across parties, candidates, and clients.

Every generated result must pass a structurally separate Secondary Bias and Integrity Gate. Failed checks block result generation.

The public edition invites peer review before the Philadelphia 4th District demonstration is treated as a production-ready civic platform.

The governing analytical chain

Each public claim should be traceable backward from narrative to measurement and source.

Figure 1. The governing analytical chain used throughout DCID.

2. The Public Problem and Product Thesis

Public institutions, campaigns, and neighborhood organizations often work from disconnected data systems. Demographic estimates, property records, service requests, election results, capital projects, business activity, community assets, and resident surveys may each be available, but they are rarely organized around a common geography, common definitions, or a shared accountability model.

The result is a familiar gap: decision-makers can produce charts, reports, proposals, and announcements without establishing what condition is being addressed, where it is concentrated, what baseline applies, who is responsible, or how success will be measured. In political communication, the same fragmentation makes it easy to overstate evidence, confuse activity with impact, or present citywide findings as though they describe every neighborhood.

Product thesis

Core proposition

A civic intelligence platform should not merely display information. It should preserve a traceable chain from condition to evidence, from evidence to intervention, and from intervention to measurable public outcome.

DCID is therefore designed as a decision-support and public-accountability system. A map selection should update indicators, sources, policies, participation structures, and AI context together. A policy statement should identify its baseline and target. An AI answer should disclose what the evidence supports and what it does not establish. A public narrative should remain traceable to the approved evidence used when it was written.

What DCID is not

It is not a voter-profiling or individual persuasion engine.

It is not an official government record system.

It is not a replacement for public agencies, election authorities, or professional advice.

It is not a general-purpose chatbot with unrestricted database access.

It is not a promise that every public dataset is complete, current, or causally explanatory.

It is not a guarantee that a proposed policy will achieve its target.

3. Design Principles and Intended Users

Principle	Operational meaning
Geography is the organizing framework	Every metric, finding, policy objective, participation record, and narrative should identify the geography to which it applies.
Measures must lead to decisions	Indicators are selected because they support diagnosis, intervention design, comparison, or accountability - not merely because data exists.
Outputs are not outcomes	A hearing, grant, program, contract, or press release is an implementation event. Resident conditions are outcomes and must be measured separately.
Public claims require provenance	Material claims should identify source, period, geography, methodology, limitations, and whether the value is a count, estimate, rate, model, or interpretation.
AI assists analysis; it does not replace evidence	Models may summarize, compare, and draft, but factual district claims must come through approved tools and governed data.

Principle	Operational meaning
Political usefulness cannot compromise analytical integrity	The system may serve partisan clients, but evidence and methods must remain symmetric and source driven.
Public and internal information remain separated	Campaign strategy, stakeholder notes, private contact information, and restricted data must never leak into public outputs.
Accessibility is a product requirement	Mapped information must have keyboard, screen-reader, and table alternatives, with non-color status communication and responsive layouts.

Primary users

User	Primary need
District resident	Understand local conditions, policy commitments, sources, polling places, and public progress.
Community leader or organization	Compare neighborhoods, locate service gaps, evaluate corridors, and prepare evidence-based advocacy.
Campaign or policy staff	Develop measurable policy, test claims, prepare district briefings, and connect strategy to geography.
Data steward or administrator	Approve sources, methods, corrections, publication status, and data-quality classifications.
Journalist, researcher, or analyst	Inspect provenance, compare areas, understand limitations, and reproduce public findings.
Candidate or senior decision-maker	Receive concise, source-grounded intelligence while preserving the distinction between evidence and judgment.

4. Platform Modules and User Environments

Four environments

Environment	Purpose	Typical contents
Public dashboard	Public explanation and accountability	District overview, neighborhood profiles, corridors, community assets, policy objectives, sources, public AI.
Internal strategy workspace	Authorized analysis and drafting	Unpublished findings, policy workbench, narrative lab, field and stakeholder intelligence, research queue.
Administrative environment	Governance and publication control	Source approval, methodology registry, user permissions, data quality, corrections, publication review, audit logs.
Conversational interface	Guided access through ChatGPT or another approved client	Approved analytical tools, structured evidence packets, accessible summaries, interactive result widgets.

Public module set

District Overview: headline conditions, priority findings, active objectives, and map context.

Neighborhood Intelligence: demographic, housing, economic, service, and infrastructure profiles.

Policy Outcomes: measurable objectives, baselines, targets, milestones, and progress.

Infrastructure and Capital Projects: project locations, status, responsible entities, and related conditions.

Commercial Corridors: business activity, property conditions, permits, transit, environmental risk, and investment context.

Community Assets and Service Gaps: facilities, programs, access areas, underserved populations, and geographic coverage.

Hyper-Local Participation: wards, divisions, committee structures, polling places, registration, turnout, and representation gaps.

Accountability: implementation status, outcome status, corrections, updates, and public evidence.

Ask the District AI: constrained question answering through approved tools.

Sources and Methodology: definitions, calculations, rights, refresh dates, quality labels, and limitations.

Shared application state

A user selection should persist across compatible modules. At minimum, the application maintains the active geography, topic, indicator, reporting period, comparison geography, filters, access scope, and AI context. This prevents a common dashboard failure in which each page silently applies a different geographic or temporal frame.

5. Reference Technical Architecture

The technical architecture separates authoritative analysis, geospatial publication, user experience, AI tools, and cloud operations. This separation is intended to support reproducibility, least-privilege access, and replacement of individual vendors without redesigning the civic intelligence model.

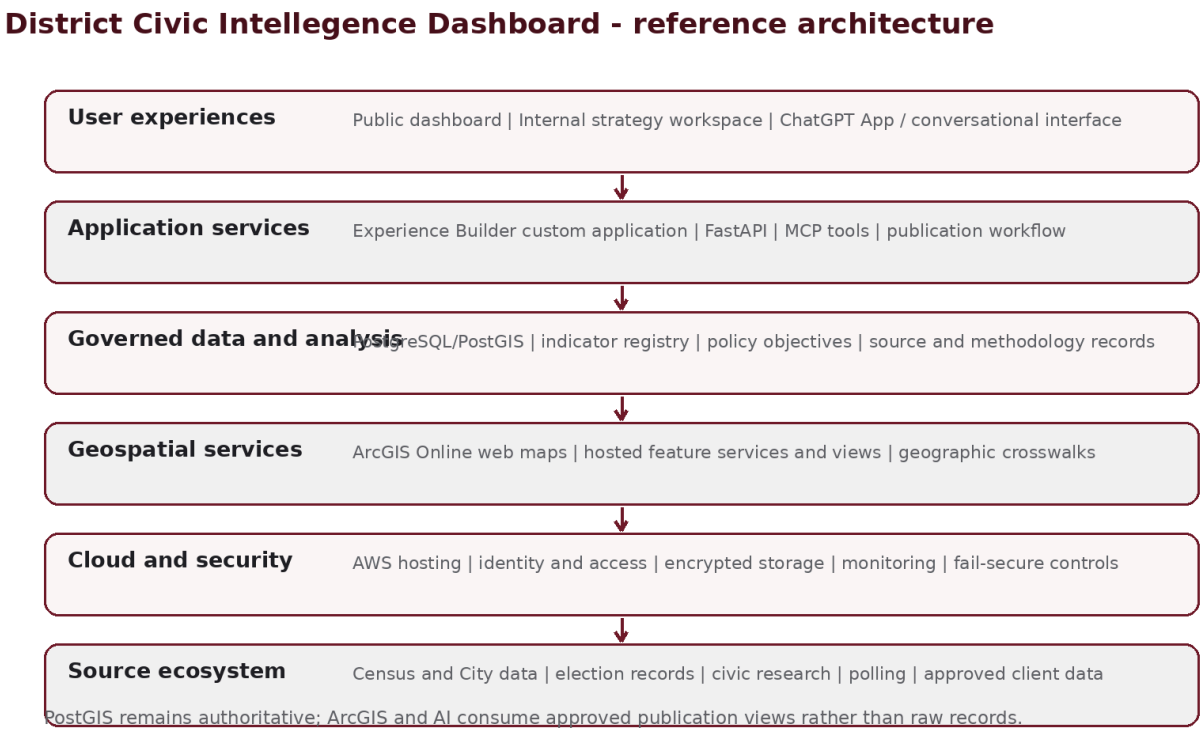

Figure 2. Reference architecture. The public peer-review edition omits sensitive implementation details.

Locked reference components

Layer	Reference technology	Role
Interactive application	ArcGIS Experience Builder Developer Edition	Custom responsive experiences, shared map state, accessible widgets, and independent web hosting.
Geospatial services	ArcGIS Online	Managed web maps, hosted feature services and views, symbology, pop-ups, groups, and permissions.

Layer	Reference technology	Role
Authoritative warehouse	PostgreSQL with PostGIS	Normalized geographic records, indicators, source metadata, historical snapshots, and spatial analysis.
Application services	Python FastAPI	Public, internal, administrative, publication, and tool-specific APIs.
AI integration	MCP tools and approved model clients	Permission-aware analytical functions, structured outputs, evidence packets, and controlled widgets.
Cloud infrastructure	AWS reference deployment	Object storage, container execution, database hosting, identity, encryption, logging, and web protection.

Esri documents that custom widgets and themes are developed through Experience Builder Developer Edition and that custom widgets or themes are not deployed directly to the ArcGIS Online edition. Downloaded experiences can be hosted independently while continuing to use ArcGIS Online content and services [1]-[3].

Amazon RDS for PostgreSQL supports PostGIS for spatial data, while ECS with Fargate provides managed container execution and CloudFront can restrict access to S3 origins through origin access control [4]-[6]. The selected architecture is therefore technically conventional at the infrastructure layer; the proposed novelty lies primarily in the governed civic data, policy, AI, and integrity workflows rather than in the mere combination of cloud products.

Governed publication flow

Retrieve and archive an original source response or file.

Validate schema, dates, categories, identifiers, geographic coverage, and rights.

Normalize records into governed core tables.

Calculate indicators using versioned methodology.

Run primary data-quality tests.

Run the independent Secondary Bias and Integrity Gate.

Publish only approved records to public APIs and ArcGIS feature-layer views.

Retain audit records and the result-generation fingerprint.

6. Geographic and Data Framework

Geographic hierarchy

The platform must relate several geographic systems without pretending they are interchangeable. The Philadelphia demonstration includes the official City Council district, Census blocks, block groups and tracts, neighborhoods, wards, divisions, polling places, commercial corridors, parcels, projects, service areas, flood zones, transit assets, police districts, and school catchments.

Each geography receives a persistent identifier, boundary vintage, source, valid dates, and relationships to parent or overlapping geographies. Crosswalk records may contain land-area, population, household, or housing-unit weights. Distance and area calculations must use an appropriate projected coordinate reference system rather than uncritically calculating in latitude and longitude.

Custom geography calculations

Counts are aggregated from the smallest reliable units using documented population, household, housing-unit, or land-area allocation.

Rates are reconstructed from weighted numerators and denominators, not averaged across source geographies.

Medians are not averaged. When a custom median is required, it should be reconstructed or interpolated from an underlying distribution where feasible.

Point and parcel records are assigned spatially and deduplicated according to indicator-specific rules.

Every estimated or allocated value carries a quality label, methodology version, and limitation statement.

Source families

Domain	Initial source families
Population and housing	U.S. Census geography and American Community Survey five-year data.
Property and development	Office of Property Assessment, vacancy indicators, L&I violations, unsafe buildings, permits, zoning, and business licenses.
Services and infrastructure	Philly311, capital projects, public facilities, transportation and environmental layers.
Public safety	Published incident data with location, classification, delay, masking, and comparability limitations.
Elections and participation	Official district maps, wards, divisions, registration, turnout, results, committee contests, polling places, and accessibility codes.
Civic research	Committee of Seventy, Pew, SSRS, universities, research firms, and other publicly accessible research with rights and methodology records.
Client-authorized data	Restricted field, stakeholder, volunteer, or organizational records isolated from public publication.

The current ACS reference baseline is the 2020-2024 five-year release, whose detailed tables and supported geographies are available through the Census API [7]. Philadelphia City Commissioners publish political maps, election data, ward and division information, and downloadable election resources that can support the demonstration [8]-[10].

Data-quality labels

Label	Meaning
Verified	Source, geography, method, and reconciliation checks passed.
Verified with limitation	Usable, but a material uncertainty or allocation limitation must be shown.
Provisional	Available for internal evaluation but not approved as a public baseline.
Stale	Past the approved refresh threshold.
Incomplete	Known geographic, temporal, categorical, or field coverage gap.
Blocked	Failed an integrity, rights, security, or publication requirement.

7. Policy Outcome and Accountability Model

DCID treats policy as a measurable object rather than a page of promises. Each objective is linked to a defined challenge, intervention, target population, target geography, baseline indicator, baseline period, target value, deadline, measurement frequency, responsible entity, sources, and limitations.

Required objective structure

Element	Question answered
Challenge	What condition is occurring, why does it matter, and who is affected?
Principle	What public responsibility or value governs the response?
Intervention	What will government, a campaign, or a partner actually do?

Element	Question answered
Baseline	What is the current measurable condition?
Target	What specific change is sought, by when, and within which geography?
Responsible entity	Who has authority or operational responsibility?
Implementation status	Has the authorized action, funding, procurement, or delivery occurred?
Outcome status	Has the underlying resident or neighborhood condition changed?
Evidence	Which source and method determine progress?

The controlling accountability distinction

Implementation answers whether the institution did what it said it would do. Outcome answers whether the relevant condition improved. The platform will not collapse these into one progress score.

Outcome ladder

Authorization - legislation, budget, executive action, agreement, or formal commitment.

Funding - identified and legally available resources.

Procurement and staffing - operational capacity to deliver.

Delivery - service, project, program, enforcement, or public asset implemented.

Output - immediate count of activities or beneficiaries.

Intermediate outcome - measurable behavioral, access, or condition change.

Long-term outcome - sustained improvement in the public condition.

8. Hyper-Local Participation and Committee Intelligence

The Philadelphia demonstration adds political geography because policy and participation operate through structures below the council-district level. The platform will support configurable terms so that Philadelphia wards, divisions, committee seats, and polling places are implementation values rather than permanent assumptions about every jurisdiction.

Public functions

Display ward and division boundaries and identify the units associated with an address.

Display polling places with the applicable election date and available accessibility information.

Describe the formal committee-seat structure and distinguish elected, appointed, acting, vacant, contested, pending, and unverified status.

Identify publicly documented committee representatives and ward leadership with a verification date.

Display aggregated registration, ballots cast, turnout, and participation change by comparable election type.

Compare participation with one community condition at a time while warning that overlap does not establish causation.

Connect divisions to neighborhoods, Census geography, corridors, policy objectives, and service conditions.

Public and restricted separation

Public or publishable	Restricted internal
Boundaries, polling places, official structures, verified public representatives, vacancies, aggregated turnout, registration and public meeting information.	Campaign contact status, volunteer capacity, canvass completion, relationship notes, supporter classifications,

Public or publishable	Restricted internal
	private contacts, internal organizational ratings, recruitment priorities.
Address lookup may return district, ward, division, neighborhood, tract, polling place, public committee information and related policy objectives.	Address lookup must never return individual voter history, private support data, health information, or restricted field records.

Philadelphia City Commissioners identify divisions as the fundamental election geography and publish current maps and downloadable election information. Committee of Seventy also publishes civic education materials on wards, divisions, ward leaders, and committeepeople [8], [11]. Official records take precedence for election status, while civic sources provide explanation and supplementary context.

9. Research, Polling, and Source Governance

The source framework extends beyond government APIs. Public reports, polls, surveys, toplines, crosstabs, questionnaires, academic studies, civic research, and machine-readable supplements may be indexed when their methodology, geography, sponsor, and use rights are recorded.

Public access is not the same as an open license

Status	Permitted treatment
Openly licensed	Reuse, adaptation, and redistribution according to the stated license.
Publicly accessible / citation permitted	Link, cite, and summarize within applicable terms; do not assume bulk redistribution rights.
Machine accessible	Automated retrieval may be technically possible, but terms, rate limits, and rights still govern use.
Restricted reuse	Use only as permitted; seek permission for republication, commercial redistribution, or derivative products.
Unclear rights	Index as citation or hold for review; do not publish the underlying material as open data.

Research-quality record

- Sponsor and fielding organization
- Field dates and publication date
- Target population and geographic universe
- Sampling frame, sample size, interview mode, weighting, response rate and design effect when available
- Margin of error and subgroup sample sizes
- Question wording, order, response options, toplines and crosstabs
- Funding source, orientation classification, corrections, and known limitations
- Rights, licensing, attribution, redistribution, derivative-use and automated-access status

A citywide poll cannot be converted into ward, division, neighborhood, or district estimates unless the source provides adequate geographic observations and an approved estimation method. Polling estimates must remain separate from administrative conditions and from a report author’s interpretation.

The demonstration proposes Committee of Seventy and Pew as initial Philadelphia-specific research sources. Committee of Seventy publishes nonpartisan ward and division resources, while Pew’s 2025 Philadelphia Resident Survey disclosed its sponsor, fielding firm, field dates, address-based sample, modes, sample size, weighting, and margin of error [11]-[13].

10. AI and Conversational Interface

AI is a governed access layer over approved data and methods. It is not the system of record. The model receives structured tool results containing geography, reporting period, measurements, comparisons, findings, limitations, source identifiers, methodology versions, access classification, and generation time.

Tool architecture

A remote Model Context Protocol service exposes a constrained catalog of functions. It authenticates the user where required, verifies scope, validates inputs, calls the DCID API, removes unauthorized fields, records tool use, and returns structured results. It does not expose arbitrary SQL, raw database credentials, unrestricted source retrieval, or administrative controls.

OpenAI currently describes the Apps SDK as a method for extending ChatGPT with applications built on MCP, and the API platform supports remote MCP tools, function calling, and structured output patterns [14]-[15]. The ChatGPT App should therefore be treated as one conversational client of DCID, not as a replacement for the authoritative warehouse, public dashboard, or administrative environment.

AI response contract

Use approved tools for factual district claims.

Name the geography and reporting period.

Distinguish counts, rates, estimates, models, survey findings, and interpretations.

Disclose interpolation, allocation, suppression, missing data, and conflicting evidence.

Do not infer causation from correlation.

Do not expose restricted data or cross client boundaries.

Do not describe an output as an outcome.

Return insufficient data when the evidence cannot answer the question.

Separate fact, inference, policy judgment, and political interpretation.

Require human review for sensitive public narratives and named-person claims.

Standard answer structure

Section	Purpose
Direct answer	Answer the user's actual question without overstating evidence.
Key evidence	Present the most relevant approved measurements.
Geographic context	Identify the active geography and valid comparisons.
What the evidence supports	Describe defensible findings.
What it does not establish	Prevent causal, individual-level, or geographic overreach.
Limitations	State uncertainty, freshness, rights, or methodological concerns.
Sources	Identify datasets, research, period, and methodology.
Suggested next analysis	Offer a governed follow-on query rather than speculation.

11. Nonpartisan Analytical Integrity

The platform may be commissioned by a candidate, party committee, officeholder, advocacy organization, labor organization, foundation, or civic institution. That client may have an explicit political position. The platform's evidence and analytical treatment must nevertheless remain partisan-neutral.

Controlling standard

Equivalent evidence, conduct, claims, sources, and methodological conditions must be evaluated under equivalent rules regardless of party, candidate, faction, sponsor, ideology, or political benefit.

Core controls

Control	Requirement
Symmetrical evidence	A source is neither rejected because it is inconvenient nor accepted because it is useful to the client.
Equal inclusion	Comparable records for applicable parties are included when available; party-specific views are explicit filters.
Comparable treatment	Geography, period, election type, denominator, suppression, quality and classification rules must match.
Time-window neutrality	Trend windows require a documented reason and cannot be selected solely to create a favorable narrative.
Geographic neutrality	Custom areas and boundary changes are disclosed and cannot be manipulated to manufacture partisan findings.
Source balance	Relevant official, civic, academic, commercial, advocacy, campaign and contrary sources are sought; quality governs weight.
No false balance	Neutrality does not require unsupported claims to receive equal weight with verified evidence.
Counterfactual symmetry	Party and candidate labels are swapped in evaluations to test whether the evidentiary classification changes.
Public-service neutrality	Party affiliation and turnout cannot determine whether residents deserve services or investment.

Advocacy versus analysis

A campaign may choose its policy, argument, audience, and strategy. Those choices are advocacy. The system must still disclose evidence that conflicts with the preferred narrative, label political interpretation as interpretation, and prevent client configuration from changing the standard applied to an opponent.

12. Secondary Bias and Integrity Gate

Primary validation determines whether a source may enter the governed system. The Secondary Bias and Integrity Gate determines whether the prepared data may generate a result. It operates immediately before a metric, comparison, map classification, AI answer, Claim Audit, briefing, export, or published narrative is released.

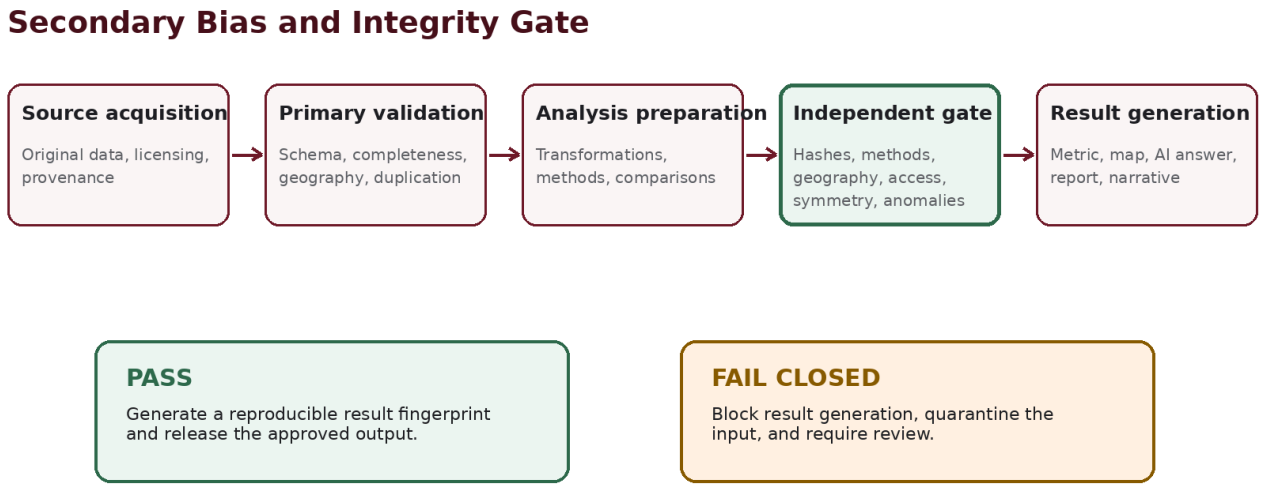

Figure 3. The gate is independent of the process that prepared the result and operates on a fail-closed basis.

Required checks

Check	Question
Source provenance	Is the publisher, retrieval, original hash, date, rights status and correction history intact?
Source change	Did fields, categories, row counts, geographic codes, or historical records change unexpectedly?
Completeness	Are required fields, parties, periods and geographic coverage present?
Geographic integrity	Are the boundary vintage, crosswalk and spatial assignment approved?
Methodology integrity	Are the numerator, denominator, weighting, suppression, benchmark and method version approved?
Partisan symmetry	Would materially equivalent facts produce the same treatment after party or candidate substitution?
Sponsor and interest	Are sponsor, funding and organizational orientation accurately disclosed?
Access control	Is the user, tenant, client, tool and publication scope authorized?
Reproducibility	Can the result be regenerated from approved versions and configuration?
Anomaly detection	Are there implausible changes, category disappearances, unusual asymmetry, or model-tool inconsistencies?

Fail-closed outcome

If the gate cannot confirm provenance, method, geography, access, symmetry, or reproducibility, the system does not silently substitute another value. It returns a blocked or pending status, retains the prior approved publication when appropriate, quarantines the affected input, records the failed test, and requires review. High-risk outputs - including named-person allegations, election-integrity claims, racial or health disparities, and material accusations of government failure - also require human review.

Result fingerprint

Each released result retains identifiers for its sources, source hashes, extraction dates, geography and boundary versions, methodology version, configuration version, validation status, gate execution, symmetry

status, requesting user or service, and generation time. This record is intended to make public corrections and internal audits reproducible.

13. Security, Privacy, and Public Accountability

Security is organized around least privilege, separation of environments, protected secrets, encrypted storage and transport, audit logging, source hashing, backups, rate controls, web protection, incident response, and the ability to suspend affected functions. Public transparency about the existence of these controls does not require disclosure of exact detection thresholds, administrative credentials, or emergency procedures.

Data classes

Class	Examples	Public treatment
Public	Approved aggregate indicators, published objectives, public maps, source records, official participation data.	Available through public views and read-only APIs.
Internal	Unpublished analysis, draft policy, research queue, publication review.	Authenticated access only.
Restricted	Stakeholder, field, strategy, private contact, proprietary polling and sensitive administrative information.	Role and client scoped; excluded from public tools.
Prohibited	Credentials, individual vote choices, unnecessary health records, unlawfully obtained data, unrestricted voter records.	Do not ingest or retain unless a separate lawful and approved basis exists; never expose publicly.

Mandatory public disclosures

- Ownership and non-government status
- Data use, source attribution, reporting period, methodology and limitations
- AI use and the fact that AI output is not independent evidence
- Polling sponsor, fielding organization, sample and geographic limitations
- Privacy practices, tracking technologies, retention and third-party services
- Corrections, responsible security reporting, accessibility, and contact information
- Demonstration or simulated status for features not yet connected to production systems

Exploit tracking and fail-secure access

Attempts to bypass authentication, escalate privileges, access administrative configuration, cross tenant boundaries, manipulate sources, disable bias controls, inject prompts, execute arbitrary SQL or code, or retrieve restricted data must generate security records. When exploitation is imminent or evident, the platform may revoke sessions, suspend credentials, isolate a client environment or module, enter read-only mode, or temporarily disengage a portal. Security monitoring applies to conduct rather than political identity or criticism.

Public accountability rule

A user should not have to guess who created a result, what evidence supports it, how AI influenced it, whether the platform is governmental, or whose interests and rights are involved.

14. Philadelphia 4th District Demonstration Model

The first demonstration is intentionally bounded. Its purpose is to prove the product thesis, not to build every future feature. It uses a controlled combination of live, cached, static, and simulated components, each clearly labeled.

Demonstration geography

The official Philadelphia 4th Councilmanic District boundary.

Three to five selected subareas, such as Wynnefield, Overbrook, Parkside, Roxborough or Manayunk, and at least one tract or block-group example.

One selected commercial corridor, with 52nd Street and Lancaster Avenue identified as a preferred initial case and West Girard as an alternative.

Ward and division overlays, polling places, committee-seat records, registration or turnout, and at least one participation-condition comparison.

Required demonstration modules

Module	Minimum proof
District overview	Five headline indicators, district map, findings, sources and active objectives.
Neighborhood intelligence	Selection and comparison across approved geographies.
Corridor profile	Business, property, permit, transit or service context connected to a policy opportunity.
Policy Outcome Engine	At least two measurable objectives with separate implementation and outcome status.
Hyper-local participation	Ward, division, polling place, committee verification, turnout and condition overlap.
Research and polling	At least one civic source, one Pew or equivalent survey source, sponsor and methodology display.
District AI	Three approved source-grounded questions and visible limitations.
Claim Audit	One predetermined claim evaluated with equivalent standards and contrary evidence.
Product configuration	A clear demonstration that the same core can support another jurisdiction or client.

Go / no-go conditions

The demonstration should not proceed when data integrity cannot be confirmed, party coverage is materially misleading, tenant boundaries are uncertain, a critical source changed unexpectedly, AI exposes restricted information, access cannot be revoked, or a material limitation would have to be concealed. A noncritical feature may be shown as simulated or replaced with an approved cached result when clearly disclosed.

15. Productization and Sustainability

The Philadelphia instance is intended to become a reusable Troubadour Solutions product rather than a one-time campaign website. The product separates a reusable core, jurisdiction templates, client configurations, and deployment services.

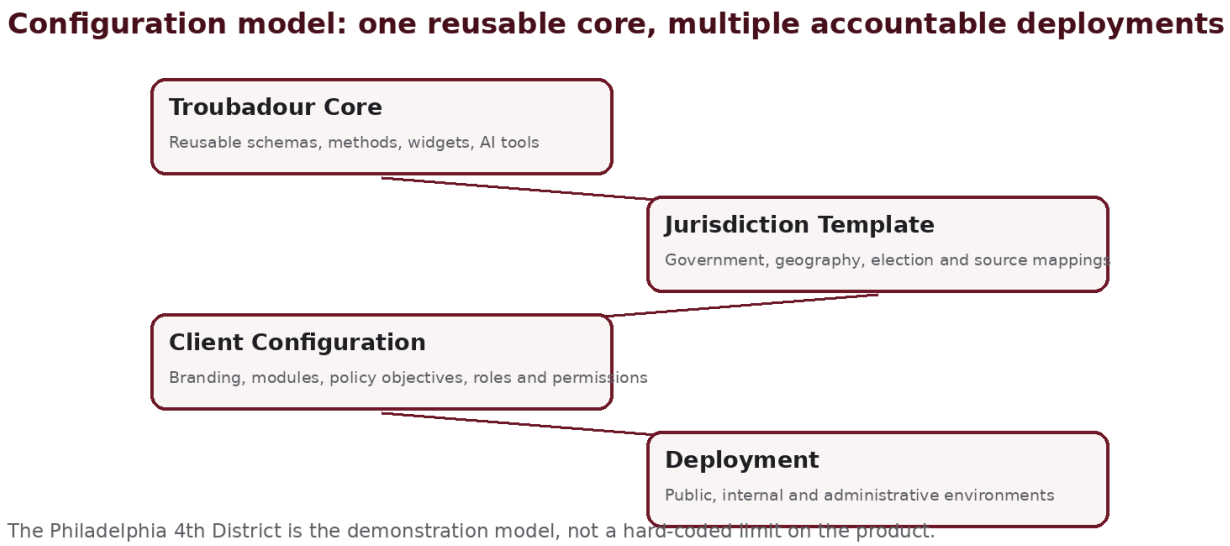

Figure 4. Configuration hierarchy for reuse without client-specific code branches.

Reusable core

- Geographic object and crosswalk model
- Indicator and methodology registry
- Source and research registry
- Policy Outcome Engine and Accountability Tracker
- Political geography, committee and participation framework
- Controlled AI and Claim Audit tools
- Bias, integrity, publication, correction and audit workflows
- Reusable widgets, deployment templates and evaluation suites

Configuration rather than hard-coding

Jurisdictions define their government, geography, political units, election authority, source portals, terminology, and data vintages. Clients define branding, policy objectives, selected modules, users, private data, AI capabilities, reporting needs, and publication rules. Philadelphia terms such as ward, division, committeeperson, and polling place are configuration values, allowing a future deployment to use precincts, county committee districts, vote centers, or other local structures.

Commercial posture

Troubadour Solutions retains ownership of the reusable software, schemas, methodologies, widgets, AI tools, evaluation suites, and generalized improvements. Clients retain control of their confidential data, branding, private notes, and client-specific content subject to contract. Public data remains subject to source ownership, licensing, attribution, and redistribution terms.

The proposed business model combines platform configuration, managed hosting, data integration, AI usage, support, professional policy and research services, dedicated environments, and white-label or enterprise options. A future open-core strategy may publish selected schemas, public methods, connectors, and accessible widgets while retaining sensitive operations, security, client intelligence, and enterprise administration as proprietary services.

16. Development and Evaluation Roadmap

Founder-led sequence

Stage	Completion focus
0. Ownership and governance lock	Confirm Troubadour ownership, accounts, repositories, credentials, decision authority, change control and publication authority.
1. Architecture validation	Have a senior engineer validate the stack, risks, cost assumptions, cloud boundaries and deployment sequence.
2. Geographic foundation	Load the district, Census, neighborhood, ward, division, corridor and polling-place geographies; build persistent identifiers and crosswalks.
3. Data and methodology	Implement source registry, ACS and selected City pipelines, indicator definitions, quality tests and publication schemas.
4. Public prototype	Build the district overview, neighborhood profile, five KPI cards, map selection, Source Drawer, responsive layout and accessibility alternatives.
5. Policy and participation	Add objectives, milestones, outcome ladder, wards, divisions, committee structures and participation comparisons.
6. AI and ChatGPT App	Expose a small set of read-only tools, structured result envelopes, source-grounded answers and one district intelligence widget.
7. Integrity and security	Implement the independent gate, audit fingerprint, quarantine, tenant controls, prompt-injection testing and access revocation.
8. Controlled demonstration	Run executive, client, public-office and technical scenarios with backup modes and a structured evaluation instrument.
9. Anchor-client pilot	Configure a second client or jurisdiction to demonstrate reuse and refine onboarding, licensing and support.

Evaluation dimensions

Dimension	Review question
Comprehension	Can a viewer explain what the platform does and how it differs from a dashboard or chatbot?
Credibility	Are the sources, methods, limitations, bias controls, corrections and security posture trustworthy?
Usefulness	Can reviewers identify concrete campaign, public-office, civic, policy and research applications?
Technical feasibility	Are the architecture, data pipelines, geographic methods, access boundaries and failure states implementable?
Public value	Does the platform improve accountability or merely add another layer of political technology?
Commercial sustainability	Can the reusable core and managed services support adoption without compromising public standards?

17. Questions for Peer Review

Reviewers need not agree with the project's political context or commercial strategy. The most useful comments will identify where the design is technically weak, methodologically ambiguous, insufficiently transparent, vulnerable to manipulation, unnecessarily complex, or missing an important public-interest safeguard.

Does the product thesis solve a genuine decision-support and accountability problem that existing public dashboards do not adequately address?

Is the distinction between implementation, output, intermediate outcome, and long-term outcome clear and operationally useful?

Are the geographic allocation and comparison standards sufficiently rigorous for custom council-district, neighborhood, corridor, ward, and division analysis?

Does the proposed public source record disclose enough information about quality, uncertainty, sponsor, rights, and corrections?

Are the rules for using polls and civic research appropriately conservative, especially when geography or subgroup size is limited?

Can the nonpartisan analytical standard be tested consistently, or does it rely too heavily on subjective human judgment?

Does the Secondary Bias and Integrity Gate add meaningful structural protection, or could it become ceremonial, duplicative, or too costly?

Which gate checks should be automated, independently reviewed, or exposed publicly?

Are public, internal, restricted, and prohibited data boundaries clear enough for campaigns and civic institutions?

What security, privacy, or misuse risks remain under-addressed?

Is the proposed ChatGPT App / MCP layer appropriately constrained, and which tools should remain outside conversational interfaces?

Does the Philadelphia 4th District demonstration contain enough scope to prove the product without overstating readiness?

Which parts should be open methodology, open source, proprietary, or independently audited?

What evidence would be required before a campaign, government office, foundation, university, or civic organization should pilot the platform?

Requested review format

Please identify the section, concern, expected consequence, and proposed correction. Reviewers are also encouraged to flag hidden assumptions and suggest comparable projects, standards, datasets, or prior art.

A. Appendix A - Core Indicator Set

Population and household conditions

IND-001 Total Population

IND-002 Total Households

IND-003 Youth Population Share

IND-004 Senior Population Share

IND-005 Disability Rate

Economic opportunity

IND-006 Median Household Income

IND-007 Poverty Rate

IND-008 Unemployment Rate

IND-009 Bachelor’s Degree or Higher

IND-010 Broadband Access Gap

Mobility and housing tenure

IND-011 Households Without a Vehicle

IND-012 Renter Household Share

IND-013 Housing Vacancy Rate

Housing affordability

IND-014 Renter Housing-Cost Burden

IND-015 Owner Housing-Cost Burden

IND-016 Median Gross Rent

IND-017 Median Owner-Occupied Home Value

Property, services and environment

IND-018 Likely Vacant Property Rate

IND-019 Open L&I Code-Violation Rate

IND-020 Dangerous Building Rate

IND-021 311 Service-Request Rate

IND-022 Illegal-Dumping Request Rate

IND-023 Flooding and Water-Related Request Rate

Safety and economic activity

IND-024 Reported Crime-Incident Rate

IND-025 Active Place-Based Business-License Density

The public interface displays a limited set of headline cards at one time. Other measures appear in topic, geography, diagnostic, and policy views. Every indicator retains a definition, numerator, denominator, unit, preferred geography, refresh cadence, quality status, source identifiers, and methodology version.

B. Appendix B - Initial AI and MCP Tool Catalog

Tool	Access	Function
get_district_overview	Public	Headline conditions, findings, objectives, sources and limitations.
get_geography_profile	Public	Approved profile for a district, neighborhood, ward, division, tract or corridor.
compare_geographies	Public	Comparable values, differences, uncertainty and methodology warnings.
get_indicator_history	Public	Historical observations and comparability status.
get_corridor_profile	Public	Corridor indicators, projects, assets and limitations.
map_service_gaps	Public / controlled	Service access and underserved-area analysis using an approved distance rule.
track_objective_progress	Public for published objectives	Baseline, target, current value, implementation status and outcome status.
audit_claim_against_data	Constrained public / enhanced internal	Extracted assertions, evidence classification, contrary evidence and limitations.
get_ward_profile	Public	Ward geography, conditions, participation and sources.
get_division_profile	Public	Division-level approved conditions and participation.

Tool	Access	Function
get_committee_structure	Public	Official structure, selection method and verification status.
get_committee_vacancies	Public with verification warning	Verified vacancies, contested, pending and unverified records.
get_polling_place	Public	Election-specific polling location and available accessibility attributes.
compare_division_participation	Public	Comparable participation measures and denominator disclosure.
compare_participation_with_condition	Public	One participation measure compared with one community condition, with correlation warning.
generate_district_brief	Internal	Structured evidence packet for approved internal drafting.
generate_public_narrative	Internal and reviewed	Draft narrative linked to approved evidence and publication review.

Administrative configuration, root or break-glass controls, quarantine override, arbitrary SQL, raw database access, unrestricted exports, and role administration are explicitly excluded from conversational tools.

C. Appendix C - Public Acceptance Framework

Area	Minimum public-review standard
Geography	District, neighborhood, ward and division selection returns the correct configured geography and boundary vintage.
Evidence	Every public metric identifies source, reporting period, method, quality status and limitations.
Comparison	Only compatible geography, period, denominator and election types are compared.
Policy	At least two objectives contain a baseline, target, deadline, responsible entity and separate implementation/outcome status.
Participation	Polling places, committee records and turnout are election- or verification-dated and party coverage is explicit.
Research	Sponsor, fielding organization, geography, sample and rights status are displayed where applicable.
AI	Approved questions return source-grounded answers; unanswerable, causal, restricted and geographically invalid requests are handled correctly.
Integrity	Every generated result passes the independent gate and retains a fingerprint; blocked data generates no public value.
Security	Public tools cannot retrieve internal information; sessions and affected portal functions can be suspended when necessary.
Accessibility	Keyboard, screen-reader, focus, non-color status and text/table alternatives are available.
Corrections	Material errors can be reported, corrected, versioned and publicly explained.
Product reuse	A second jurisdiction or client can be configured without rewriting the core architecture.

D. Appendix D - Artifact Provenance and Reduction Map

Subject	Master artifacts	Peer-review location
Concept and product requirements	FDCI-A00, A01	Sections 1-4
Information architecture and wireframe	FDCI-A02, A03	Section 4
Data and indicators	FDCI-A04, A04.1, A10	Sections 6, 8, 9; Appendix A
Technical system	FDCI-A05	Sections 5, 10, 13
Implementation and founder model	FDCI-A06 v1.0 and v2.0, A13	Sections 15-16
Data, AI and publication governance	FDCI-A07	Sections 6, 9-13

Subject	Master artifacts	Peer-review location
Demonstration scope and overlays	FDCI-A08 v1.0 and v1.1	Sections 8 and 14
Productization	FDCI-A09	Section 15
Nonpartisan standard	FDCI-XA01	Section 11
Secondary integrity gate	FDCI-XA02	Section 12
Security and public notices	FDCI-A11 and XA03	Section 13
Demonstration runbook	FDCI-A12	Sections 14 and 16

The complete 308-page Master Artifact Compendium remains the source of record. This peer-review edition reorganizes and reduces that material for public review; it does not supersede the full technical, governance, security, or implementation artifacts.

E. Appendix E - Reviewer Response Guide

A useful review may be submitted in prose, annotated PDF, memorandum, issue list, or structured response. The following fields are recommended:

Field	Reviewer entry
Reviewer name / organization	
Area of expertise	
Section reviewed	
Finding or concern	
Why it matters	
Recommended change	
Evidence, standard, precedent or comparable project	
Severity	Advisory / material / critical
May the comment be attributed publicly?	Yes / No / With conditions

Suggested submission statement: “I reviewed the public peer-review edition of the District Civic Intellegence Dashboard. My comments address the published design and do not certify the security, accuracy, legality, or production readiness of an implementation.”

R. Selected Technical and Civic References

[1] Esri, Overview of ArcGIS Experience Builder Developer Guide.
https://developers.arcgis.com/experience-builder/guide/overview-of-experience-builder/

[2] Esri, Widget and Theme Deployment.
https://developers.arcgis.com/experience-builder/guide/widget-theme-deployment/

[3] Esri, Experience Deployment.
https://developers.arcgis.com/experience-builder/guide/experience-deployment/

[4] AWS, Managing Spatial Data with the PostGIS Extension.
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.PostgreSQL.CommonDBATasks.PostGIS.html

[5] AWS, What Is Amazon Elastic Container Service?.
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/Welcome.html

[6] AWS, Restrict Access to an Amazon S3 Origin.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html

[7] U.S. Census Bureau, 2024 American Community Survey Five-Year Data.
https://www.census.gov/data/developers/data-sets/acs-5year/2024.html

[8] Philadelphia City Commissioners, Political Maps.
https://vote.phila.gov/resources-data/election-resources/political-maps/

[9] Philadelphia City Commissioners, Reports and Data.
https://vote.phila.gov/resources-data/election-resources/reports-and-data/

[10] Philadelphia City Commissioners, Archived Election Data Sets.
https://vote.phila.gov/resources-data/past-election-results/archived-data-sets/

[11] Committee of Seventy, City Wards and Divisions.
https://seventy.org/political-maps-of-philadelphia/city-wards-and-divisions

[12] Committee of Seventy, Philadelphia Ward Leaders and Committeepeople.
https://seventy.org/philadelphia-ward-leaders-committeepeople

[13] The Pew Charitable Trusts, 2025 Philadelphia Resident Survey reporting and methodology.
https://www.pew.org/en/research-and-analysis/issue-briefs/2025/06/pew-poll-philadelphians-concerned-about-public-safety-and-financial-well-being

[14] OpenAI Developers, Apps SDK and MCP resources. https://developers.openai.com/

[15] OpenAI API Quickstart, Responses API and tools.
https://platform.openai.com/docs/quickstart/make-your-first-api-request

[16] Troubadour Solutions, District Civic Intellegence Dashboard Master Artifact Compendium, Version 1.0.
Private source-of-record compilation, July 2026

END OF PUBLIC REVIEW DRAFT