

WholeContext AI

Technical Clinical and Impact Dossier

Engineering evidence and methods for independent expert review

Kimberly Chulis PhD

Core Analytics

14 October 2026

Version 1.0

WholeContext connects a patient's longitudinal evidence with clinical review, daily function, care coordination and relevant U.S. public data. This dossier documents the released synthetic MVP, the research and data engineering behind it, and the evidence gates for responsible expansion.

Scope of this release

Public release v39 • 65 fictional patient identities • 171 cataloged resources • reproducible federal-data assembly • ecological and synthetic ML studies • source-linked workflows and exports

The technical contribution is a traceable path from fragmented information to reviewed evidence and a usable next question. The next evaluation must establish whether that path improves care workflows and outcomes.

Demonstration

<https://wholecontextai.kchulis.chatgpt.site/>

Prepared for the HHS TOPx AI and Invisible Illness challenge and independent technical, clinical and research scrutiny. No institutional endorsement is implied.

Document control and reviewer guide

This dossier presents the engineering, research, data assets and transformation program behind WholeContext AI. It is written for technical, clinical, research and public-sector review of the released synthetic MVP and the evidence supporting its next stage.

Field	Controlled value
Document	WCT TD 001 • Version 1.0 • 14 October 2026
Author and accountable lead	Kimberly Chulis PhD • Core Analytics
Release baseline	Public version 39 • published 14 October 2026 at 15:23 UTC
Source commit	4a85dc7a60a989e2c731cb9ed2ed29ffdba7cbe1
Public demonstration	https://wholecontextai.kchulis.chatgpt.site/
Scope	Public technical disclosure, synthetic care workflows, public-data research and explicitly identified future work
Evidence convention	Released = present in the pinned product. Tested = observed in the named engineering or research record. Designed = specified or scaffolded. Planned = requires execution and acceptance.

How to read the evidence

Sections 1–5 establish the problem, design basis and architecture. Sections 6–16 explain the data and analytical methods. Sections 17–21 document workflows, testing and release evidence. Sections 22–26 address safeguards and evaluation. Sections 27–30 specify adoption, transformation and reproducibility. Appendix A preserves the complete 171-resource index. References identify source records and primary methodological authorities.

Claims are tied to a dated source or release. Historical experiment counts, source registries and synthetic cohorts are not added together. A source entry, a field, a trained model, a passing software check and a patient outcome measure have different denominators. This document does not treat them as interchangeable evidence.

The baseline is deliberately pinned to version 39 even if later work advances the repository or demonstration. Engineering tests reported here were recovered from saved outputs; they were not rerun as a new clinical or scientific study for this publication. [E01–E14]

Contents

The report moves from the user problem and architecture to data and methods, then to evidence of delivery, safeguards and the transformation program. The complete catalog index is retained in Appendix A.

Reading path	Sections	Pages
Executive case and intended use	1–2	4–5
Research foundation and architecture	3–5	6–8
Data portfolio engineering and provenance	6–9	9–12
AI NLP and multimodal methods	10–11	13–14
Ecological and uncertainty modeling	12–14	15–17
Real-data and synthetic research programs	15–16	18–19
Public context Twin and collaboration	17–19	20–22
Software testing and release record	20–21	23–24
Security governance and clinical safety	22–24	25–27
Validation and success measurement	25–26	28–29
Adoption transformation and delivery	27–30	30–33
Evidence and method references	E01–E14 and R01–R07	34–35
Primary public sources and glossary	D01–D12 and terminology	36–37
Complete 171-resource catalog index	Appendix A1–A11	38–48

A short technical review

Read the architecture, canonical state and NLP contract, then compare the two modeling studies and their denominators. Inspect the v36 regression lesson and v39 release boundary. Use the threat model and validation protocol to assess the next deployment decision.

A short clinical or public value review

Read the intended use and research foundation, then follow Elliot, Marion and Caleb through context, chronology and selected handoff. Review the measurement dictionary and adoption plan. Separate current information-workflow evidence from future clinical outcome hypotheses.

1 Executive case

WholeContext AI makes a fragmented health history inspectable across the care journey. It connects patient reports, clinical records, caregiver observations, daily function and appropriately matched public context into a dated, source-linked account that a patient can review and a clinician can interrogate.

The contribution

The central engineering contribution is continuity of evidence. Source attribution, time, corrections, unknowns and sharing scope follow an observation through the journal, progress views, clinician review, Digital Twin and selected-record packet. OpenDataDelta makes the contribution of public information visible by comparing the same patient evidence with and without selected contextual sources. The system can also show that a source adds no useful change.

What has been built

The public MVP offers Patient, Support Circle, Clinician, Care and Costs, and Discovery Lab entry points. It contains 65 fictional patient identities, a protected computation service for synthetic review, reversible context comparisons and PDF outputs. The current evidence includes 43 passing v38 regression groups and 17 passing v39 media, component and affected-regression groups. These overlapping release checks are reported separately. [E01–E03]

What supports the engineering

The project includes a 171-resource research catalog; a 107-record scored integration atlas with 54 derived-variable designs and 16 output designs; reproducible federal-data assembly; a 3,143-county ecological modeling study; and an independently generated synthetic uncertainty experiment. These assets support different questions and have separate evidence boundaries. [E05–E12]

What the next stage must establish

The next decision is whether this evidence infrastructure reduces reconstruction effort, preserves clinically important information and supports completed care work in supervised use. Prospective patient and clinician evaluation, production identity and privacy controls, operational testing and function-specific regulatory assessment are required before broader real-patient deployment. Clinical improvement and savings remain hypotheses.

The intended public benefit is practical: patients spend less effort reconstructing their history; clinicians can locate the basis of a statement and see what is unresolved; caregivers can contribute without losing attribution; and care teams can follow whether a plan is being completed and whether function changes. The dossier supplies the mechanism, the present evidence and the tests needed to evaluate that benefit.

2 Challenge alignment and intended use

WholeContext addresses the Lyme Innovation pathway while developing reusable evidence and coordination services for complex illness. Its immediate intended use is preparation, evidence review, contextual exploration and continuity; its public demonstration uses fictional clinical records.

Review dimension	Evidence in this dossier	Next proof
Public benefit	Patient research, chronology, correction, function tracking and selected handoff	Observed comprehension, review effort and follow-up completion
Data and technical innovation	Federal-source lineage, same-case context comparison, ecological models and uncertainty methods	Independent reproduction and incremental workflow value
Execution and delivery	Public v39 deployment; protected review service; regression and PDF evidence	Hosted multi-user, mobile, accessibility and download acceptance
Scale and reuse	Shared observation and provenance contracts; condition expansion gates; pilot operating model	Measured service capacity, integration conformance and cross-condition evaluation

Users and decisions

Patients select and correct the history they want to share. Clinicians review source evidence, chronology, conflicts and missing information. Caregivers contribute separately attributed observations and practical support. Financial participants receive deliberately restricted cost information. Researchers explore public population data independently of a person's clinical evidence.

Functional boundary

A county reporting rate is contextual information. It does not become an individual infection probability. A symptom after a treatment is a temporal observation, not an established treatment effect. The Digital Twin represents recorded evidence over time; it does not simulate human physiology. The current workflow does not autonomously diagnose, prescribe, contact a care team or confirm appointments.

Review approach

Reviewers can evaluate a concrete chain: source → accepted observation → dated view → selected context → clinical question → packet. A capability can be useful before it is predictive, and a predictive research result does not establish patient benefit. The official Phase 2 framework emphasizes delivered function and public applicability; this dossier maps evidence to that framework without assigning itself a competition score. [R01]

3 Research foundation and design traceability

The product requirements were informed by a structured literature and evidence review, a small patient survey, clinician discovery, subsequent clinical and stakeholder feedback, expert input and iterative prototype work. The evidence supports design choices; it is not an effectiveness trial. [E05]

Evidence stream	Documented scope	Design consequence
Evidence workbook	50 records: 48 external research or guidance records, one local survey and one clinician interview	Chronology, uncertainty, function, access, burden and provenance
August survey	8 submissions; one test excluded; 7 substantive; primary U.S. analytic sample n = 6	Patient review and correction before sharing
Later survey reporting	12 respondents reported in beta materials; later export not reconciled for percentages	Retain the n = 6 analytic denominator for quantified findings
Expanded feedback	At least 4 physicians, 2 PhD expert contributors and one administrative stakeholder review reported; participation may overlap	Clinically meaningful timing, concise handoff, role clarity and billing visibility

From need to requirement

In the August analytic sample, chronology or memory needs were coded in 5/6 responses; privacy, accuracy, provenance or control in 5/6; and approval before sharing was rated very or extremely important by 6/6. These are convenience-sample design signals. They motivate a correctable canonical history, visible source links and patient-selected disclosure. They do not estimate national prevalence or adoption.


Requirements that can be tested

R01 preserves a coherent history and source identity. R02 preserves clinically meaningful timing. R03 makes uncertainty understandable. R04 supports permissioned contributions and corrections. R05 connects questions to feasible care and follow-through. R06 tracks function and treatment timing without automatic causal attribution. Each requirement has a corresponding task, expected behavior and future user measure.

The review is described as a structured evidence synthesis. A formal systematic-review claim would require a recoverable search protocol, screening log and review procedure. The original source directory retains the literature-level limitations. Individual research participants and private correspondence are excluded from this public report.

4 System architecture

The released application separates a public synthetic demonstration from authenticated record services. Server-side computation prepares clinical-review, context and packet results; the client handles interaction and presentation. The architecture below describes inspected code boundaries, not an independently audited production network. [E02, E04]



Inspected v39 component boundaries • production clinical validation remains a separate gate

Figure 1. The public demonstration and authenticated record path share presentation patterns but have distinct input and authorization boundaries. The optional model service is a separate, review-first synthetic-text pathway.

Runtime and persistence

The inspected deployment uses a Sites-hosted Worker module and static client assets. The server routes synthetic review, research, source extraction, education, provider lookup and authenticated care operations. Authenticated care storage is represented by a D1 binding and Drizzle schema for patient ownership, membership, records, audit events and model-use accounting.

External dependencies

Stored public snapshots drive OpenDataDelta. The optional text service calls an OpenAI model through a bounded server operation. Current source evidence does not establish an Azure production deployment, an autonomous clinical-agent system or a live EHR connection. Azure object storage and job execution belong to the transformation design.

The diagram intentionally omits credentials, private implementation rules and security-sensitive configuration. Public method transparency is provided through input/output contracts, provenance, test scope and measured results.

5 Component contracts and engineering decisions

The service boundaries are designed to keep each result attributable to a permitted input and to make a failed dependency visible. The release history shows a transition from client-side interpretation to a protected server computation path. [E02, E04]

Component	Contract	Failure behavior
Synthetic review	Known fixture identity, allowed record selection and bounded options produce a dated review	Reject foreign identities, arbitrary records and malformed selections
Canonical observation model	Stable observation identity, source, subject, time, value and unit	Preserve missingness; correction or reset recomputes dependent views
OpenDataDelta	Explicit geography, applicable time and selected stored measures	Report unavailable, excluded or no-change context
Packet generation	Selected records and context produce a dated patient-specific snapshot	Cancel stale results; preserve identity and source citations
Authenticated care	Server-enforced ownership, membership and record grants	Reject unauthorized access; isolate recipient-specific disclosure
Optional source categorization	Synthetic text becomes quoted sentence-level candidates	Validate returned labels and IDs; keep original text on failure

Architecture decisions

Use a shared observation model to prevent different screens from manufacturing competing histories. Keep interpretation on the server to control the execution boundary. Preserve original records and reversible scenarios so a context experiment cannot overwrite patient evidence. Prefer deterministic calculation for date, unit, source and permission rules. Use a model only where a bounded language task adds value.

Operational implications

Asynchronous callers need identity and request-version checks so a slow response for one patient cannot replace the selected patient's current view. Exports are immutable snapshots; a previously downloaded PDF does not update when the record changes. Durable cross-device synchronization, hosted recovery exercises and production service objectives require additional implementation and acceptance.

These decisions are architectural tradeoffs: the approach favors inspectability and bounded behavior over unrestricted automation. A graph database, vector store or general agent framework is not a prerequisite for the present evidence contract.

6 Data portfolio and the 171 resource catalog

The research portfolio was built in layers. Broad discovery identified possible sources; the integration atlas scored candidates; research pipelines acquired and normalized specific releases; the product selected a smaller set of stored measures. Counts describe those layers and are not additive. [E06–E08]

Asset and date	Count	Meaning
Source catalog • 22 September	171	Datasets, cohorts, portals, standards, methods, personal-source pathways and proposed resources
Health Data Atlas • version 4	107	Scored source records; 69 prioritized for integration now and 11 classified controlled or restricted
Atlas design outputs	54 / 16	Named derived variables / proposed outputs
Frozen backend • 15 September	3 / 4 / 5	Federal publishers / product-version selections / raw statistical variable files
OpenDataDelta inventory • 13 October	120	12 stored-snapshot, 90 catalog-only and 18 unavailable or restricted registry entries
Live Delta adapters in that inventory	0	Stored context is not a continuous upstream-refresh service
Public county-context asset	3,143	County records in the release snapshot; measure coverage differs by field

Selection logic

The atlas links sources to user decisions, join keys, derived variables, outputs and integration waves. “Integrate now” is a prioritization decision. It does not prove acquisition, legal permission, a connector or an evaluated product contribution. Likewise, an unavailable/restricted implementation label may describe an environment limitation rather than a publisher restriction; the exact release terms remain authoritative.


Catalog preservation

Appendix A includes every one of the 171 resource names and its recorded access label, organized into 11 domains. New reference IDs preserve the original within-domain order. These IDs are a navigation aid and do not replace the atlas source IDs or claim a completed row-level reconciliation. The full source catalog supplies units, access routes and limitations.

The next catalog release should map each candidate to an exact product version, acquisition receipt, transformation, output, permission and evaluation result. A single source can support several fields, while several catalog entries can describe the same originating evidence.

7 Data engineering and reproducible assembly

The September backend established a repeatable route from source bytes to normalized tables, compatible joins and an evidence bridge. This work provides a concrete engineering foundation independent of the interface. [E08]



Versioned execution · schema and join audits · source-family comparison · reproducible outputs

Figure 2. Source bytes and acquisition records remain separate from normalized observations, task-specific joins and product evidence. Evaluation checks both accepted contributions and rejected matches.

Acquired asset	Verified scale at the checkpoint	Use
CDC county Lyme • 2001–2023	72,519 county-year rows	Reported surveillance burden
Census Vintage 2023 • 2020–2023	12,576 rows	Frozen replay denominators
Census Vintage 2025 • 2020–2025	18,864 rows	Revised reference denominators
NOAA June 2023 precipitation and mean temperature	3,107 county-month rows; 3,097 matched and 10 unmatched	Bounded historical weather join

The 2023 CDC–Census union contained 3,154 rows and 3,143 available rates. Ten CDC-only keys and one Census-only key remained unmatched. The joined June weather output retained 3,164 analytical and audit rows, including 57 annual rows without NOAA context and 10 NOAA rows without a valid annual match. These rows were not repaired by invented values.

Two isolated offline rebuilds compared 76 outputs: 75 matched byte for byte and one matched after removal of a declared audit timestamp. The run passed 26 bridge checks and 38 contract, offline and registry checks. It produced no new predictive evaluation. Run: evidence-2026-09-15-64f15061ca88. [E08]

8 Data quality provenance and temporal validity

A dataset contributes only when its geography, time, unit, population, rights and intended interpretation fit the task. The engineering work preserves the reasons an observation is admitted, excluded or unavailable.

The source record

For each selected release, retain publisher and canonical source ID, original URL, product/version, observation period, release date, retrieval time, geography and grain, fields and units, denominator, access conditions, original hash, transform version, output hash and responsible steward. Derived fields must retain both their numerator and denominator lineage. A refreshed download does not make its observation period current.

Time is more than one date

The observation date, source publication date, acquisition time and application review time answer different questions. Retrospective reconstruction must use historically compatible data. Forecast evaluation must exclude information unavailable at the prediction cutoff. Current context can use a newer release if the mismatch is disclosed; it must not silently enter a historical benchmark.

Quality condition	Required handling	Example from the record
Missing or unmatched	Keep null and reason; retain rejected join	NOAA rows without valid annual FIPS remain audit records
Zero reported	Preserve as reported, not absence of disease	1,893 counties reported zero in the October ecological study
Suppressed or restricted	Retain status; do not infer a number	Sparse cells and governed collections remain distinct from zero
Different denominators	Label each period and population	2023 physician count divided by 2022 population
Duplicate provenance	Link common origin rather than count twice	Email confirmation of an existing diary observation
Changed reporting regime	Stratify or qualify comparison	CDC surveillance definition changed in 2022

The June 2023 Anne Arundel example combined 162 annual reported cases, a population of 594,582, a calculated annual rate of 27.25 per 100,000, 90.8 mm June precipitation and 21.99 °C June mean temperature. The annual outcome was not relabeled as a monthly outcome. This added eligible context; it did not demonstrate a predictive or clinical gain. [E08, R02]

9 Evidence ontology and canonical state

WholeContext organizes evidence around a subject, source and time. The design distinguishes what was reported, what was measured, what was requested, what was inferred and what remains unresolved. Terminology mapping is an interoperability objective, not a license to replace source meaning. [E09]

Contract element	Meaning and constraint
Identity and attribution	Stable source and observation IDs; patient identity; reporter and relationship; original artifact or quoted span
Temporal representation	Event time, precision and timezone where known; record availability; historical cutoff; explicit undated state
Clinical representation	Original text, candidate concept, value, unit, result polarity and review status; unmapped text remains inspectable
Relationships	Corroborates, conflicts, duplicates, supersedes, belongs to a test episode or supports a question
Context representation	Public geographic observation with period, denominator and applicability; separate from personal findings
Consent and disclosure	Owner, role, selected records, recipient scope and revision of the shared snapshot
Transformation provenance	Extractor or rule version, validation outcome, accepted correction and dependent output version

Corrections and propagation

The canonical observation work tests identities, dates, units, source links, missingness, corrections and reset across Marion, Elliot and Caleb. A correction must propagate to the derived views that use that observation. A hypothetical scenario stays in separate state and reset restores the recorded baseline. These are stronger invariants than visual agreement between independently computed charts.

Unknown does not mean negative

A documented lexical repair keeps “not reported as negative” unknown and “not negative” unresolved. Twelve targeted and eleven workflow tests passed for that correction. This is a bounded parser improvement, not a validated general clinical-negation model. It illustrates why source text, polarity and interpretation must remain distinguishable.

The September ontology plan identified incomplete runtime terminology binding and draft FHIR mappings. The released canonical observation model advances shared state, but a complete terminology service, full profile conformance and validated evidence graph remain separate acceptance work. [E02, E09]

10 AI and natural language processing methodology

AI contributes at different stages of WholeContext. Founder-directed AI-assisted engineering supports implementation and artifact production; an optional runtime service categorizes synthetic source text; and separate research programs evaluate numerical models. Those roles must remain distinguishable. [E04, E11–E12]

AI role	Present evidence	Validation boundary
Engineering assistance	Code, fixtures, documentation, debugging and regression development supported by AI tools	Changes reviewed through source diffs and executable checks; tool output is not authority
Synthetic-text categorization	One bounded model request returns sentence IDs and allowed labels	Mocked-provider contract tests; a hosted successful model run is not established by configuration alone
Evidence organization	Deterministic source, date, permission and display logic	No general clinical reasoning agent or autonomous diagnosis
Research modeling	Ecological county models and synthetic uncertainty experiments	Offline measured results; separate targets and populations
Narrated guides	Generated narration with captions, transcripts and synthetic product visuals	Media validation; no patient-specific clinical narration service

Runtime language contract

The optional source service accepts explicitly synthetic text from a signed-in visitor. It preserves original sentence offsets and quotes; the model proposes labels against permitted sentence IDs. Validation requires complete coverage of the submitted sentence set and rejects invalid identifiers or labels. Review and a separate authenticated acceptance action precede record creation. Imported text is treated as data, not as controller instructions.

Bounded operation and reproducibility

Input size, output tokens, timeouts and persistent per-user and shared request limits constrain the service. Traces record service/run identity, source attribution, model information and reported usage where available. Provider failure or invalid output preserves original text and labels the fallback. The exact runtime configuration and prompts are controlled implementation artifacts.

NLP evaluation sequence

The next benchmark should annotate entities, subject, negation, time, medication lifecycle, relations and contradictions against source spans. Compare a rules baseline, model-assisted extraction and clinician-adjudicated references. Freeze cases and prompts; group related case variants; report unsupported claims, omitted critical facts and correction effort. No general clinical NLP accuracy is claimed from the present sentence-label tests.

11 Computer vision and multimodal evidence

The current photo workflow documents an image and preserves uncertainty about its meaning. The multimodal research program is designed to propose reviewable evidence with source and time, while preventing a scene or signal from silently becoming a clinical fact.

Present image capability

The inspected implementation supports owner-selected images, date/context and sharing choices, local re-encoding, dimensions and limited brightness/contrast statistics. It includes stale-image or workspace-result checks. These operations do not establish diagnostic image quality, identify a tick, classify a rash or infer disease. Hosted browser upload acceptance remains distinct from local numerical tests. [E04]

Context Discovery example

Elliot's prepared photograph depicts an outdoor scene with a dog. The demonstration uses authored candidates. It does not establish dog ownership, the date or location of a hike, or tick exposure. The image can justify a clarification question. Undated evidence stays undated and is excluded from dated reasoning until supported. [E02, E13]

Planned modality	Method to evaluate	Required acceptance evidence
Documents and scans	OCR plus span-grounded extraction, document structure and table parsing	Character/field errors, dates and units; failed-OCR handling; patient review
Patient photographs	Quality triage and bounded multimodal description	Source-grounded descriptions; metadata consent; stratified image-quality and skin-tone review
Email and messages	Permissioned connector, deduplication, temporal and reporter extraction	Revocation, scope restriction, correspondence linkage and no unsupported fact creation
Wearables	Time alignment, device quality flags and within-person trends	Device validity, wear-time missingness, timezone handling and task-specific usefulness
Narrated visit brief	Patient-reviewed synthesis from selected accepted evidence	Every substantive statement traceable; correction and consent before sharing

Each modality needs a method card: intended task, permitted inputs, target labels, training or prompting procedure, benchmark composition, subgroup coverage, failure criteria and review workflow. A future diagnostic image function would require a new clinical and regulatory evaluation; it cannot inherit validation from a photo-documentation feature.

12 County modeling design

The October 6 ecological study asks where reported Lyme burden clusters and how geographic, environmental and population features relate to county reporting. It is an actual retrospective modeling study, separate from the synthetic patient-care demonstration. [E11]

Population target and features

The study matched 3,143 counties or county equivalents across 51 state/DC jurisdictions to 2023 outcomes, retaining 88,672 reported cases from the included county file. Ten unmatched CDC codes were excluded without guessed crosswalks; their 2023 counts were zero. A Census-only county with unknown outcome was not assigned zero. The source county-file total was not reconciled to a separate national headline.

Two main targets were evaluated: reported cases per 100,000 residents, modeled on a $\log(1 + \text{rate})$ scale, and whether a county reported any cases. The primary feature set contains 94 inputs spanning geography, population and density, land and habitat, social vulnerability and climate. Expanded historical snapshots are retained as sensitivity analyses. Zero reporting, surveillance practice and residence-versus-exposure differences constrain interpretation.

Model comparison and leakage controls

The study compares persistence and simple baselines with linear/generalized models, logistic regression, tree ensembles, support vectors, nearest neighbors, a neural network, XGBoost, LightGBM and a fixed equal-weight ensemble. Five outer folds hold out entire states. Declared tuning grids use three additional state-group folds within training. Imputation and scaling are fitted inside training folds. This design limits geographic leakage; it does not establish prospective validity.

Sensitivity and explanation

Source-family ablation, regional holdouts, temporal transport, residual spatial autocorrelation and surveillance-stratified associations probe transportability. State-cluster intervals are conditional on saved out-of-fold predictions rather than full model refits. TreeSHAP explanations apply to held-out XGBoost predictions on the log-rate scale; they describe model behavior, not causal effects. Correlated features may share attribution.

The fresh-source temporal comparison fits 2018 and evaluates 2023 largely in the same counties. Historical publication availability was not reconstructed. The 2022 surveillance change further limits trend interpretation. This is not a combined new-time and new-geography validation, and no county score is promoted to an individual disease probability. [E11, R02]

13 County modeling results and interpretation

The ecological results show both useful structure and important limits. A model can improve a log-scale metric while worsening absolute error in the original rate units. A competitive surveillance-history baseline is therefore central to the interpretation. [E11]




Figure 3. Common-county comparison on 3,133 counties. Lower is better in both panels. XGBoost with prior-year surveillance improves log-rate RMSE but has higher rate MAE than persistence. Source: common_county_history_comparison.csv.

Reported result	Estimate	Interpretation
Any-reporting fixed ensemble	Pooled AUC 0.9104; mean-fold AUC 0.9015; Brier 0.1154	County reporting discrimination and probability error, not patient accuracy
Any-reporting logistic model	Pooled AUC 0.9018; Brier 0.1224	A relatively simple comparator remains competitive
Fresh temporal transport	Log-rate RMSE 1.0228; R^2 0.6836; rate MAE 23.3909	2018 → 2023 in largely the same counties
Residual spatial dependence	Moran I 0.6263; permutation $p = 0.001$	Substantial remaining spatial structure limits naive independence assumptions

The ensemble's state-cluster interval for pooled AUC was 0.8725–0.9413 and for Brier score 0.0881–0.1415. These intervals condition on already-trained predictions. The prevalence baseline has AUC 0.5 within folds; its pooled AUC differs because fold training prevalences vary. Reporting both avoids a misleading baseline comparison.

Earlier September experiments are retained separately: prepared county replay reduced MAE from 7.6234 to 6.1491 per 100,000 while Brier worsened from 0.12993 to 0.13261. Masked forest reconstruction reduced random-mask MAE from 16.443 to 12.979 percentage points, but the state-block comparison had only two observed state clusters and uncertain transportability. Masked precipitation MAE worsened from 110.246 to 110.537 mm/year. These masked targets do not validate naturally missing values. Different targets, periods and inputs prohibit combining these into a single performance claim. [E08]

14 Uncertainty abstention and synthetic ML

An isolated October 10 experiment studies when a classifier should defer. Its labels represent artificial administrative document load. It uses no disease, patient or clinical-outcome target and supplies no clinical accuracy claim. [E12]

Design element	Recorded implementation
Independent partitions	3,000 training; 1,000 calibration; 2,000 same-distribution evaluation; 1,000 distribution-shift stress rows
Models	Logistic regression and histogram gradient boosting with fixed choices; majority-class baseline
Decision policy	Input missingness/out-of-distribution gate; split-conformal prediction sets; abstention and review counts
Reproducibility	Seeded generator; rerun produced identical results JSON and prediction CSV hashes
Analysis	Eligible-set metrics, retained singleton predictions, set coverage, artificial subgroup and shifted-sample results

Observed logistic results

Of 2,000 same-distribution evaluation rows, 1,862 were gate-eligible. Eligible accuracy was 0.7540 compared with a 0.5000 majority baseline; AUC was 0.8363 and Brier score 0.1652. Prediction-set coverage was 0.9071. The singleton policy retained 1,136 of 2,000 rows at 0.8477 accuracy, leaving 864 rows for review. These denominators describe different decisions and should not be collapsed into one “accuracy.”

Observed boosting results

Histogram boosting achieved prediction-set coverage 0.9066 and singleton accuracy 0.8472 at 1,139/2,000 retention in the same-distribution evaluation. The result is a software research comparison, not evidence that machine review reduces clinician workload by the corresponding amount.

Methodological implication

Conformal methods can quantify set coverage under their assumptions. Distribution shift, dependent samples or a changed selection policy can invalidate an assumed guarantee. A production adaptation must define the calibration population, eligible set, abstention policy and utility of human review before deployment. Artificial groups do not validate demographic fairness. [R05]

The demonstrated contribution is an executable uncertainty workflow with explicit retention and review denominators. A clinical adaptation would require appropriate labels, external validation, calibration monitoring and a clinically accountable escalation process.

15 Real data research program

The September real-data methodology specifies a broader empirical program connecting surveillance distributions, longitudinal care processes, ecology and biological time courses. Its implementation scaffold passed six package tests; the referenced record does not establish completed real-data model execution. [E10]

CDC population fusion

Public line-listed case files intentionally omit geography. The proposed method estimates phenotype dependencies within each surveillance era and calibrates newly generated synthetic population cells to separately released geographic margins using iterative proportional fitting. Sparse-cell hierarchical models are a sensitivity option. The method must not reconstruct the hidden county of an original case. The 2008–2021 and 2022–2023 eras are handled separately.

MEPS longitudinal care processes

The design uses HC-245 and HC-257 longitudinal people and public condition/event files to construct encounter, prescription, utilization, expenditure and access sequences. Public encrypted link keys connect the permitted files. Gaussian mixtures provide soft journey segmentation; agglomerative clustering and k-means serve as sensitivity comparators; cross-panel replication tests stability. An A69 sensitivity cohort is not a validated Lyme cohort.

Ecological contribution ladder

A surveillance-history baseline is augmented sequentially with climate, habitat, population/access and permitted One Health variables. Stacking is admitted only after out-of-fold comparison. The primary question is incremental value on geographic and temporal holdouts, measured with error, calibration and uncertainty. A source can also be retained for a distinct contextual use even if it does not improve prediction.

Program gate	Evidence needed before promotion
Acquisition and rights	Exact releases, source terms, hashes, schema audit and reproducible execution
Population fusion	Released margins reproduced without linking original people to geography
Care-process segmentation	Stable and interpretable segments replicated in an independent panel
Incremental context	Prespecified same-population comparison and source-family ablations
Clinical translation	Independent patient-level evaluation against an appropriate comparator

GEO biological series and field-ecology datasets can supply separate scientific evidence. They are not substitutes for care trajectories. Dummy-fixture model metrics validate execution only and are excluded from scientific findings in this dossier.

16 Synthetic data and controlled experiments

Synthetic records make failure states, sharing rules and counterfactuals testable without exposing private patient histories. WholeContext uses several cohorts created for distinct purposes; their sizes are not combined into a claimed patient sample. [E02, E10, E13]

Asset	Purpose	Boundary
Five-case pilot • September	Lyme-compatible, alternative, tick-borne and nonspecific-control journeys with geography counterfactuals	Hidden truth separated from visible inputs; not a clinical validation cohort
Development cohort • September	7 attributed published case backbones plus 14 entirely synthetic scenarios	21 development cases; publication date is not event date
Generated cohort • October	64 fictional cases across 8 authored families and variants	Engineering coverage; family variants are statistically related
Released fixture	65 fictional patient identities; Marion has 142 source records	Release-specific workflow and identity testing
Uncertainty experiment	7,000 separately generated administrative-load rows across partitions	No clinical labels and no training on the care cohort

Generation and leakage discipline

The longitudinal methodology separates latent truth, care-process events, noisy observations and actor-specific knowledge. Expected recoverable facts and unresolved facts are frozen before evaluation. Pilot corrections removed diagnosis-bearing case titles and visible gold-truth paths. Hidden annotations are read only after first-run outputs are frozen. Cases and their mutations must stay together when creating evaluation partitions.

Controlled changes

Geography-only counterfactuals hold the visible clinical record fixed. A weak nonspecific case moved into a high-reporting county must not become clinically supported Lyme. Other controlled variations should change one dimension at a time: source date, negation, reporter, unavailable layer, contradictory observation or sharing permission. Expected no-change behavior is part of the test contract.

Private family-derived histories are excluded from model development, benchmark construction, calibration, geographic feature work, public demonstrations and performance metrics. Abstract challenge motifs can inform design, but distinctive private chronology and identifying combinations are not reproduced. Clinical plausibility review and real-patient generalization remain separate gates.

17 OpenDataDelta and source contribution

OpenDataDelta exposes the effect of selecting public information while holding patient observations fixed. Version 38 connected populated service data to the patient, clinician and Twin views and tested all 65 patient-county lookups across nine distinct home counties. [E03]

Fictional case and county	Available measures	Question-focused	Twin at latest baseline
Elliot • Hennepin MN 27053	74	21	20
Caleb • King WA 53033	65	17	10
Marion • Worcester MD 24047	64	21	21

These are measures and topic matches, not independent source counts or clinical applicability scores. The Twin applies its own dated eligibility rules, which explains smaller selections. Selecting all available context is not evidence that every measure is useful to the clinical question.

Concrete contribution

For Elliot, stored Hennepin context includes an annualized reported Lyme rate of 15.23 cases per 100,000 residents per year for 2019–2022 using a 2022 denominator, physician supply of 124.96 per 100,000 using a 2023 count and 2022 population, and 8.9% of households without a vehicle from the 2018–2022 ACS/SVI context. These can inform questions about exposure history and feasible access. They do not change his laboratory findings.

Comparison protocol

Start with the selected patient record and a fixed cutoff. Compute the patient-only review. Add a defined source family and preserve the source period, geography, unit and denominator. Compare displayed context, questions and exclusions. Remove the family and recompute. Check that patient symptoms, measurements and hypothetical observations remain unchanged. Preserve explicit no-change outcomes and failed-source states.

What the release repaired

The earlier Local context panel read empty authored fields instead of the service; Twin selection was not inherited consistently; and a county report could lack a supported county. The repair connected the actual service, retained eligible selections, expanded question scope appropriately and rejected stale county responses. The 43-group release record includes actual component handlers and isolated Worker execution, not a real-browser acceptance pass.

18 Digital Twin and longitudinal deep dives

The WholeContext Digital Twin is a longitudinal evidence representation. It supports historical cutoffs, recorded body and function observations, public-context selection and reversible hypothetical scenarios. It makes the knowledge available at a point in time inspectable. [E02, E03, E13]

Elliot and context discovery

An email refers to an existing diary observation. Confirming the relationship adds corroboration without duplicating symptom ratings. A caregiver report remains separately attributed, including its own time and 4/10 function observation. An undated photograph is retained as selected context but excluded from dated evidence reasoning. The selected handoff preserves those distinctions.

Marion and function over time

Marion’s full fictional history contains 142 dated source records. Her personally meaningful goal is to dress and reach comfortably. Recorded symptom and function changes can be examined alongside interventions and clinical encounters. A historical cutoff prevents future records from appearing as information already known. A public-data toggle changes contextual information without changing her symptom values.

Caleb and question scope

Caleb’s respiratory-recovery question requires appropriate recovery and access context. The release widened the default question construction while keeping irrelevant mosquito context out of the focused view. This is question-specific source selection, not a diagnostic ranking. Different source counts across cases reflect available fields, topic matching and time eligibility.

Invariant	Demonstration expectation
Evidence and scenario separation	A hypothetical change is visibly distinct and reset restores the recorded baseline
Time consistency	Earlier body-history playback does not display context from a different unqualified cutoff
Source continuity	A statement in the packet can be traced to the selected original record or public source
Patient continuity	Patient switching clears stale results; identity is preserved in every export
Missingness	Unknown dates and absent observations remain visible rather than being interpolated

The tested Twin core includes 100 timepoint baselines and 128 hypothetical combinations in its release history. These are deterministic scenario checks, not simulated clinical trajectories or treatment-effect estimates. Downloaded packets are dated snapshots and require regeneration after a correction.

19 Collaboration care and cost workflows

WholeContext uses the same evidence foundations across five entry points, with role-specific actions and disclosure. The intended outcome is a completed next step in care, not simply a longer summary. [E04, E13]

Workspace	Core activity	Information boundary
Patient	Journal, review observations, track symptoms/function and prepare selected records	Patient controls corrections and disclosure
Support Circle	Contribute attributed observations and practical support	Caregiver statements remain distinct; grants are recipient specific
Clinician	Inspect chronology, evidence, uncertainty, context and Twin scenarios	Clinical judgment remains with the professional
Care and Costs	Review visits, travel, out-of-pocket and caregiver-time assumptions	Financial sharing uses a restricted snapshot rather than clinical free text
Discovery Lab	Explore public context and research evidence	Population research remains separate from personal clinical findings

Care navigation

The inspected service supports user-selected specialty and location, curated practice sources and recipient-scoped internal tasks with preview, confirmation, updates and revocation. It does not establish external communication, confirmed scheduling, network participation or appointment availability. The future integration should preserve a state trail from question through requested, ordered, scheduled, completed and reviewed actions, with evidence for each transition.

Cost interpretation

The financial workflow makes assumptions visible. Visits, transport, direct spending and caregiver time can be tracked separately. A modeled before/after scenario is not observed savings. Patient spending, payer expenditure and societal time costs must use different denominators and avoid double counting. Source years and any inflation adjustment must be stated.

Progress and response

Treatment timing and function trends help prepare a follow-up conversation. Concurrent treatments, natural history, missing observations and changes in measurement can explain apparent improvement. The initial evaluation should measure completeness and interpretability of the record before asking whether the software changes health outcomes.

The adoption pathway begins with a useful patient-reviewed packet and a manageable clinician review task. Further automation is admitted when authorization, source reliability and a measurable workflow benefit can be demonstrated.

20 Software testing and development sequence

Verification follows the information path from source to visible output. The saved reports distinguish local numerical checks, service contracts, component event adapters, actual PDF bytes, hosted HTTP checks and browser acceptance. Passing one layer does not substitute for another. [E01–E04]

Checkpoint	Recorded result	Scope
September reproducible backend	26 bridge + 38 contract/offline/registry checks; two rebuilds	Data assembly and contribution/rejection logic
Canonical observations	24 checks across three featured patients	Identity, source, dates, units, correction and reset
Protected service release v32	18/18 suites; service checks include all 65 cases	Worker, controlled components, isolated database and PDF bytes
Clinician and PDF repair v36	8/8 groups; 210 client-generated requests	Complete source-selection round trip and input validation
OpenDataDelta v38	43/43 groups	Context, body/recap visibility, service/model regressions and PDF integration
Narrated guides v39	17/17 groups before publication	Media, source, component and affected regressions
Hosted v39 checks	4/5 groups; 23 changed assets matched	HTTP/content checks; byte-range expectation failed

A regression that changed the test strategy

A valid source ID containing an underscore was rejected by the shared review validator. The clinician default selection and both PDF formats could therefore fail even though smaller manually selected test inputs passed. The repair validated the complete selection generated by the actual client code. The unchanged request failed on the old service and succeeded locally after correction. [E03]

Development sequence

Specify the user task and acceptance invariant; freeze representative source fixtures; implement the smallest complete data path; run component and service checks; inspect exported artifacts; exercise adversarial and permission cases; save the exact candidate; verify hosted assets and runtime; complete browser acceptance; retain a rollback checkpoint. New failures become durable regression cases.

Counts in this table overlap across releases and must not be summed. Existing checks establish software behavior on defined fixtures. No prospective clinical trial, external security audit or production load study is implied.

21 Release record and operational readiness

Version 39 was successfully published on 14 October 2026. Its release record includes the exact commit, deployment receipt, hosted asset comparisons, media checks and remaining acceptance items. This dossier preserves that release boundary. [E01]

Release fact	Evidence
Public URL	https://wholecontextai.kchulis.chatgpt.site/
Commit	4a85dc7a60a989e2c731cb9ed2ed29ffdba7cbe1
Published	14 October 2026 • 15:23:22 UTC • deployment succeeded
Narrated guides	Welcome 54 seconds; Patient 59 seconds; Clinician 68 seconds
Accessibility provisions in media	Burned-in captions, WebVTT, complete transcripts and no autoplay
Hosted preservation	23 changed public assets matched saved bytes; protected service request succeeded with 74 selections
Immediate recovery baseline	Version 38 • commit 0932f67a18484681de8823e36dbbb3d582d1ddb6

Acceptance still outstanding

The hosted byte-range test expected HTTP 206 and received HTTP 200. Complete MP4 downloads matched the saved files, but this does not establish seeking behavior. Browser playback, sound, mobile layout, keyboard navigation and operating-system downloads remain owner acceptance items in the recovered record. A supplementary HTTP probe that returned 403 added no acceptance evidence.

Boundary checks

The hosted private-data API returned 401 without authorization. Sampled private server and archived evaluation paths returned 404. These are useful bounded checks; they do not constitute a penetration test or prove every possible path is protected. Hosted two-account isolation, revocation, photo upload and a traced synthetic model call require separate evidence.

Operating constraints

The v39 package has little remaining size headroom. Future releases should separate archival media and engineering evidence from the runtime package while preserving durable provenance. Growth requires measured cold-start, memory, concurrency, request-size, response-time, storage and export behavior. One local resource smoke test does not establish production capacity.

The release can be publicly reviewed as a synthetic MVP. A production health-data service needs additional operational, privacy, clinical and contractual gates. Public availability and real-patient readiness are different release decisions.

22 Privacy security and threat model

The current public demonstration uses fictional records and stored public context. The inspected implementation contains ownership, membership, selected-record grants and a separate server boundary. The security strategy builds on those controls and identifies the evidence still needed for regulated-data operations. [E02, E04]

Threat	Present control or test	Additional acceptance
Cross-patient or cross-recipient disclosure	Known-case bounds; server membership/grants; isolated permission regressions	Hosted independent-account tests, revoke/re-add and audit review
Source text instructs the model	Bounded label task; no model tools; ID/label validation; injected-text tests	Adversarial corpus and independent prompt-injection evaluation
Wrong patient from asynchronous response	Stale-result rejection and patient-switch tests	Browser race and slow-network acceptance
Sensitive data in a public export	Selected records; restricted financial fields; bounded public-asset scan	Formal export policy, privacy assessment and exhaustive release inventory
Exposed implementation or secrets	Protected computation; sampled private paths inaccessible	Secret scanning, dependency inventory and independent application review
Excessive resource or model use	Bounded input, timeouts and persisted usage limits	Load tests, alerting, abuse response and cost monitoring
Loss or incorrect alteration	Source versions, canonical corrections and saved rollback	Hosted backup restore, retention enforcement and recovery exercise

Production safeguards to implement

Before real-patient intake, establish a data-flow inventory and risk assessment; least-privilege identities; suitable encryption and key management; secure development and dependency management; session and tenant isolation; audited administrative access; tested backup and restore; incident response and breach procedures; vendor agreements and data-processing terms appropriate to the deployment.

The October bounded privacy scan examined 929 text assets and 140 PDFs for three private-reference identifier patterns and found no matches. That result is narrow: it is neither a comprehensive de-identification assessment nor a certification. Synthetic public content does not remove the need to test authentication and private-path isolation. [E02]

23 Data governance and accountable AI

Governance links a source's permitted use to every derived product and claim. WholeContext's next operating model assigns stewardship, approval and monitoring responsibilities explicitly; these are proposed responsibilities, not claims that all roles are already staffed.

Accountability	Decision rights	Required record
Product and engineering lead	Approve scope, release criteria and rollback	Requirements, change history and signed release decision
Data steward	Admit exact source releases and authorize refresh	Rights, provenance, quality checks and retention schedule
Clinical safety lead	Approve intended use, evaluation and clinical wording	Hazard log, adjudication and unresolved safety issues
Privacy and security lead	Approve access, vendors and regulated-data boundary	Risk assessment, permissions review and incident record
Independent evaluator	Review benchmark, analysis and claim support	Frozen protocol, deviations, results and limitations

Consent and data lifecycle

Real-patient use should separate permission to process a record, permission to share it with a named role and permission for research or secondary use. Revocation should stop future access and processing where applicable, with a documented policy for already delivered snapshots and legally required retention. Deletion, export and retention must cover raw artifacts, derived fields, caches, logs and backups.

Public and controlled resources

Public availability is not equivalent to unrestricted reuse. Controlled cohorts, partner collections and participant-generated sources require their own access pathways. Geography intentionally removed from a public case record must not be reverse-inferred. A license or agreement attached to a raw source must remain discoverable in downstream exports and research releases.

AI change control

Maintain a model/service inventory, intended-use card, training or prompt provenance, benchmark version, approval history and monitored failure criteria. A provider or prompt change triggers affected evaluations. NIST AI RMF provides a useful organization for governance, context mapping, measurement and management; the crosswalk is a strategy, not a certification. [R03–R04]

Model-generated text, literature synthesis and stakeholder feedback are evidence inputs requiring review. They do not create institutional endorsement, establish clinical truth or authorize use of private records. Public reporting uses aggregate participation and excludes private correspondence.

24 Clinical safety and interoperability

Clinical accountability begins with the intended function and the information it exposes. A source-linked summary, a population model and a future diagnostic image function have different risks and require different validation. [R06]

Safety case structure

For each function, define the user, setting, decision supported, input population, credible failure, potential consequence, mitigation, residual uncertainty and evidence of mitigation. Priority hazards include omitted critical facts, wrong identity, reversed negation, incorrect time, unsupported inference, inappropriate reassurance from a zero county value, overreliance on a polished summary and disclosure to the wrong recipient.

Evidence and escalation

An uncertainty statement must remain attached to the output where it matters. Conflicting records should be shown rather than silently averaged or resolved. The system should preserve source statements and give reviewers a path to inspect them. A research score should never enter a clinical view without a new intended-use assessment and validation. The current demonstration is not a monitored clinical or emergency service.

Interoperability target	Proposed mapping	Acceptance gate
FHIR R4 exchange	Patient identity; Observation; DiagnosticReport; DocumentReference; Provenance and Consent as applicable	Resource/profile validation; round-trip fidelity and semantic review
Clinical terminology	LOINC, SNOMED CT, ICD and RxNorm where appropriate	Licensed/versioned terminology, mapping review and unmapped-state handling
USCDI and USCDI+	Named elements mapped to supported product fields	Explicit element crosswalk and documented gaps
Blue Button and EHR pathways	Permissioned claims or clinical-source connectors	Authorization, revocation, sandbox then partner conformance tests
TEFCA-related exchange	Future governed participation through a suitable arrangement	Applicable agreements, identity and exchange-policy evaluation

FHIR is an exchange specification; a draft mapping is not conformance certification. The current record supports schema design and draft mappings, not live production EHR interoperability. [E09, R07]

Regulatory assessment should be performed function by function against the applicable jurisdiction, intended user and actual claims. FDA's January 2026 clinical decision support guidance is a relevant U.S. reference. This dossier does not claim device exemption, FDA clearance, HIPAA certification or clinical effectiveness. [R06]

25 Scientific validation and evaluation design

The next evaluation should distinguish technical fidelity, clinical appropriateness, workflow utility and health outcomes. A strong result at one level cannot substitute for the next. The protocol below is proposed and requires independent clinical and statistical review before recruitment.

Study stage	Design	Primary evidence
Source fidelity benchmark	Frozen, independently annotated cases with blinded adjudication	Critical-fact recall; unsupported statements; date, negation, unit and attribution errors
Workflow comparison	Counterbalanced case review with usual records, structured packet and packet plus public context	Review time and missed critical information, with quality constraints
Supervised pilot	Consenting patients and clinicians in defined care settings	Completed handoffs, patient comprehension, correction burden and follow-through
External validation	New sites, periods, populations and relevant conditions	Transportability, subgroup error and calibration where applicable
Outcome evaluation	Appropriate prospective comparator and follow-up	Patient-reported function, care completion, utilization and cost with causal limits

Prevent avoidable bias

Freeze eligibility, comparator, primary outcome and analysis before examining results. Keep related patients or synthetic variants within one partition. Blind adjudicators where feasible and separate product development from final outcome review. Log deviations and all excluded cases. Account for clustering by clinician or site, repeated measures and learning effects. Report effect sizes and intervals, not only p values.

Test incremental value

Use the same case with patient-only evidence, a structured summary, and selected public context. Prespecify which missing facts and useful questions count. Include irrelevant sources, unavailable sources, contradictory evidence and no-change controls. More text, more sources or higher confidence alone is not an improvement.

Sampling and fairness

A feasibility pilot should estimate task variation and recruitment constraints; it should not be described as powered for clinical efficacy. The subsequent sample calculation must use a clinically meaningful difference, variance, clustering and expected attrition. Evaluate geography, age, sex/gender where appropriate, race/ethnicity, language, disability, access and condition-specific missingness, with privacy-aware suppression of small cells.

Register outcome definitions, adverse-event handling and stopping rules before use. Independently assess severe unsupported statements and privacy incidents. No diagnosis, recovery, time-to-care or savings effect is established by the current synthetic release tests.

26 Measurement of success

Success is a usable, accurate and appropriately shared care artifact followed by a completed action and measurable patient benefit. The proposed metric dictionary separates current software evidence from future pilot outcomes. Baselines and thresholds must be frozen before pilot analysis.

Metric and denominator	Measurement window	Owner and interpretation
Critical-fact recall: retained reference facts / eligible adjudicated facts	Per benchmark version	Clinical evaluator; distinguish critical from optional facts
Unsupported statement rate: unsupported material statements / all reviewed material statements	Per benchmark and release	Safety lead; report severity and absolute counts
Provenance completeness: substantive claims with valid source links / all substantive claims	Per generated packet	Engineering/data steward; inspect links, not just presence
Review effort: clinician active review minutes per completed case	Matched tasks; then pilot visits	Study lead; paired with omissions and comprehension
Patient comprehension: correctly explained key items / assessed items	After review and at follow-up	Patient research lead; assess uncertainty and sharing scope
Handoff completion: accepted usable packets / initiated packet workflows	Defined 30-day cohort	Product lead; include abandonment and correction work
Care follow-through: verified completed next steps / eligible agreed next steps	30/60/90 days as protocol specifies	Clinical operations; record access barriers and censoring
Function change: within-person change on prespecified valid measure	Baseline and defined follow-ups	Clinical evaluator; no automatic treatment attribution
Cost per usable case: attributable processing + support cost / accepted usable cases	Monthly pilot cohorts	Operations; include failed attempts and human review

Current evidence provides source-selection counts, correctness checks, release status and research metrics. It does not provide pilot baselines for review time, comprehension, completed referrals, health improvement or savings. Those fields remain unmeasured rather than assigned assumed values.

Proposed release blockers are any confirmed cross-patient disclosure, wrong-patient export or severe unsupported clinical assertion in the accepted test set. Lesser-error tolerances and service objectives should be agreed with the evaluation team before testing. A zero count in a finite test set is not proof of zero production risk.

27 Adoption economics and public value

The proposed adoption sequence starts with patient-led visit preparation and public exploration through Discovery Lab, then moves into supervised clinic pilots. Each step must justify the work it adds for patients and clinicians.

Patient and caregiver adoption

Offer a clear first task: select a fictional journey, inspect its evidence, understand the sharing model and see the resulting packet. In a governed pilot, assisted intake and correction should accommodate limited digital access, fatigue, cognitive difficulty and caregiver participation. Plain language, accessible navigation and a usable print artifact matter as much as feature breadth.

Clinical adoption

Pilot with a small defined workflow and an accountable clinical owner. Agree where the packet enters the visit, who reviews corrections and how unresolved actions are followed. Compare active review time and critical-information retention against the existing process. Avoid shifting unmeasured data-cleaning work onto clinicians. Training, onboarding and support effort belong in the economic assessment.

Commercial and operating hypotheses

Candidate routes include patient access supported by clinic or partner services, organizational workflow licensing and governed research collaborations. These are business-model hypotheses, not booked revenue or verified willingness to pay. Patient financial benefit, clinic efficiency, payer expenditure and societal burden require separate measurement.

Adoption stage	Evidence to advance	Practical risk
Public discovery	Users understand purpose and can complete a synthetic walkthrough	Confusing research outputs with medical advice
Supervised feasibility	Usable packets; acceptable correction and review effort; consent understood	Intake burden and incomplete records
Clinic integration	Workflow fit, accountable follow-through and reliable access controls	Integration work and clinician capacity
Organizational expansion	Sustainable cost per case, support capacity and independent validation	Site variation, inequity and vendor dependence
Cross-condition reuse	Condition-specific evidence, safety case and benchmark	Reusing infrastructure without validating new clinical logic

A credible economic model records volume, activation, retention, accepted cases, processing cost, support time and fixed operating costs. Savings scenarios should expose their assumptions and compare against observed utilization when available. Public engagement is useful for discovery and feedback; it is not evidence of clinical impact.

28 Transformation roadmap

The roadmap preserves the existing data engineering and product contracts while advancing through explicit evidence gates. Time horizons below are planning assumptions measured from an approved next-stage start, not funded commitments or promised delivery dates.

Phase	Deliverables	Exit gate
0 • Release acceptance First 2 weeks	Hosted browser and multi-account tests; known media issue disposition; release inventory and recoverability	Signed acceptance, severe defects resolved, reproducible rollback
1 • Pilot foundation Weeks 2–8	Production data-flow/privacy assessment; secure intake and storage; consent lifecycle; benchmark and pilot protocol	Independent security/clinical review; approved real-data boundary
2 • Supervised utility Months 2–4	Patient/clinician feasibility pilot; correction workflow; metric instrumentation; source-adaptor reliability	Usable cases, credible workflow comparison, documented failures and equity gaps
3 • Interoperability and operations Months 4–8	Selected EHR/claims connector; terminology mapping; durable jobs and versioned object storage; observability	Conformance tests, restore exercise, service objectives and measured unit cost
4 • Evaluated expansion Months 8–12 and beyond	New condition contracts; permissioned multimodal extraction; externally replicated research	Condition-specific validation and function-specific regulatory decision

Infrastructure transition

A proposed Azure destination uses versioned object storage, containerized bounded Python jobs and an application service consuming approved evidence outputs. The migration must reproduce frozen outputs, preserve identities and permissions, measure runtime and cost, and pass recovery tests. No Azure deployment is claimed by the current source evidence. A cloud move is justified by operating requirements, not by architecture branding.

Scientific transition

Complete the exact-source CDC/MEPS research program before promoting its conclusions. Replace exploratory comparisons with prespecified external benchmarks. A Condition Intelligence Engine would reuse evidence contracts while validating each condition's clinical interpretation. A Population Learning Network would require governed aggregation, representative denominators and applicable research oversight.

Prize funding could accelerate independent evaluation, secure patient-data infrastructure, standards-based connectors, clinical review and operational support. Resource allocation should follow the gates above, with feasibility evidence preceding broader recruitment and feature expansion.

29 Delivery plan dependencies and risks

The engineering program is managed through work packages with a named accountable role, a reviewable output and an exit decision. Kimberly Chulis is the present accountable product and engineering lead; specialist review roles below require assignment or contracted support.

Work package	Accountable role	Deliverable and dependency
Release and reliability	Engineering lead	Browser acceptance, regression evidence, recovery record; depends on hosted test capability
Data stewardship	Data steward	171-to-atlas-to-runtime crosswalk and exact-release ledger; depends on terms and source access
Clinical safety	Clinical lead	Function-specific hazards and adjudicated benchmark; depends on independent clinical reviewers
Pilot research	Study lead	Protocol, recruitment, consent and analysis plan; depends on privacy and ethics determination
Security and privacy	Security/privacy lead	Risk assessment, access tests, vendor review and incident procedures; gates real-data intake
Integration and scale	Engineering/data lead	Selected connector, queue/job semantics and measured service limits; depends on partner/API permissions
Adoption and economics	Product/operations lead	Training, support, usability and cost-per-case evidence; depends on actual pilot workflows

Material delivery risks

The main risks are incomplete raw lineage for inherited context, limited independent clinical evaluation, historical source vintages, unresolved hosted acceptance, package-size pressure, fragmented real-world records, clinical-review capacity and possible additional requirements for new functions. The mitigation is a bounded work package and measurable gate for each risk, rather than an assumed platform-wide readiness level.

Change and release control

Maintain a versioned backlog linking a user need to requirement, source change, code change, test, result and claim. Record decision rationale and superseded results. Changes to patient state, permissions, source eligibility, prompts, models and export wording receive risk-appropriate review. Keep rollback artifacts and avoid destructive schema changes without recovery evidence.

No fixed staffing, signed partnership, procurement, cloud budget or pilot enrollment is asserted here. The plan defines the evidence and capabilities needed to make those commitments responsibly and to measure whether the program is delivering useful results.

30 Evidence package and reproducibility

An expert should be able to trace a headline claim to a source record, an input population, a method and a saved result. This dossier is the public synthesis of that evidence. Supporting technical records preserve the execution detail and version history.

Artifact family	Contents	Review purpose
Product and architecture	Intended use, component/trust boundaries, observation contracts and decision records	Assess the mechanism and limits
Data atlas and catalog	171-resource index; 107-record scored atlas; derived variables and source admission rules	Assess breadth, relevance, provenance and rights
Data engineering	Acquisition manifests, normalized tables, join audits, hashes and offline rebuild comparison	Reproduce accepted and rejected contributions
Model and method cards	Targets, partitions, baselines, calibration, ablations, explanations and limitations	Evaluate scientific claims without mixing tasks
Software verification	Fixture inventory, contract/component tests, Worker requests and PDF outputs	Inspect release correctness and regressions
Release record	Pinned commit, publication receipt, asset checks, known gaps and rollback	Identify exactly what was published
Governance and safety	Threat model, hazard log, consent plan, access decisions and evaluation gates	Assess readiness for governed real-patient use
Adoption and outcomes	Pilot design, metric dictionary, economics assumptions and transformation roadmap	Test public value and sustainability

Reproduction standard

A reproducer needs the exact source commit, input hashes and permitted data, dependency versions, commands, split definitions, seeds where used, model or service settings, expected outputs and deviation log. Recompute metrics from saved predictions and confirm denominators before comparing models. Reproduction of a software result does not independently validate the clinical interpretation.

Public disclosure and independent scrutiny

Public materials should include the methods, source boundaries, selected synthetic examples, aggregate results and known limitations. Private patient records, credentials, participant correspondence and protected implementation assets are excluded. Additional controlled review can inspect reproducibility artifacts under appropriate access terms; this document does not imply that every supporting file is publicly hosted.

The next companion deliverables can expand the architecture, complete data matrix, detailed modeling protocol, validation and safety case, and transformation plan without changing the evidence standard established here.

Evidence references

Project records below form the evidence basis for the reported implementation and results. Dates identify the source checkpoint. Earlier records are used only for the historical claim stated; later release records supersede their deployment status.

E01. WholeContext v39 Publication Receipt, 14 October 2026. Successful publication, hosted results, narrated-guide verification and rollback. Source commit 4a85dc7a60a989e2c731cb9ed2ed29ffdba7cbe1; deployment appgdep_6acf9df9c5a88191b21ebe9a7ed93194.

E02. Protected-service release v32, 13 October 2026. evaluation/final-release-v32/RELEASE_READINESS.md and associated results; canonical observation, Twin, permissions and synthetic cohort records. Current source was inspected at the v39 baseline; later source-only changes were excluded.

E03. Clinician/PDF Repair v36 and OpenDataDelta v38, 14 October 2026. evaluation/release-repair-v36/REPAIR_REPORT.txt; evaluation/open-data-v38/RELEASE_REPORT.txt and saved JSON results. Publication status is supplied by E01 rather than the earlier candidate wording.

E04. Runtime contracts at v39. AGENT_DESIGN.md; dist/server/index.js and associated care, extract, demo-review, observation, Twin, source and context modules; db/schema.ts; dist/client/source-inventory.json. Code inspection supports architecture and contract descriptions, not an external security audit.

E05. WholeContext Research Foundation and Evidence to Design Synthesis, version 0.2, 4 October 2026. Aggregate research participation, August analytic denominator, evidence matrix and requirements traceability.

E06. Lyme and Related Health Data Source Catalog, 22 September 2026. 171 cataloged sources and candidate resources, access pathways and limitations. Appendix A preserves all recorded names and access labels.

E07. WholeContext Health Data Atlas Integration Matrix, version 4; WholeContext Federal and Open Data Catalog and Access Roadmap, version 0.1, 4 October 2026. 107 scored source records, 54 derived variables, 16 output designs and access gates.

E08. WholeContext Data Preparation Selection and Management, version 0.2.0, 15 September 2026. Frozen acquisition inventory, reproducible backend run evidence-2026-09-15-64f15061ca88, historical replay and masked-feature experiments.

E09. WholeContext Evidence Ontology and Orchestration Plan, version 1.0, 15 September 2026. Evidence contract, orchestration design, lexical correction and interoperability gaps.

Method and source references

The references distinguish project execution evidence from external method and standards guidance. Use of a published method or standard does not transfer its validation or imply endorsement of WholeContext.

E10. WholeContext Real-Data Modeling Methodology, version 0.2, and Synthetic Longitudinal Data Methodology, version 0.2, 20 September 2026. CDC/MEPS program, six-test scaffold, five-case pilot, leakage controls and claim gates.

E11. County ecological study, 6 October 2026. `dist/client/county-lyme-research/summary.json`, `common_county_history_comparison.csv`, `classification_models.csv`, `regression_models.csv`, `feature_block_ablation.csv`, `audit_checks.json` and `report.html`. Outputs preserve held-out predictions, inputs, sensitivity analyses and limitations.

E12. County and uncertainty checkpoint, 10 October 2026. `research/COUNTY_UNCERTAINTY_CHECKPOINT.md`; `research/run_uncertainty.py`; `dist/client/research/uncertainty/results.json`, `model-card.md` and `heldout_predictions.csv`. Independent synthetic administrative-load experiment.

E13. WholeContext Phase 2 Evidence and Methods Brief and Team and Proposed Solution, 13 October 2026. Feature narratives, synthetic case relationships, user workflow and adoption direction; release status updated using E01–E03.

E14. WholeContext Technical Strategy Architecture Addendum, 12 October 2026. Earlier architecture and source-count reconciliation. Used for historical design context, not as the final deployment record.

R01. NIH. TOPx HHS Tech Sprint for AI and Invisible Illness. Official challenge scope and Phase 2 evaluation framework. Accessed 14 October 2026. <https://www.nih.gov/challenges/topx-hhs-tech-sprint-ai-invisible-illness>

R02. CDC. Surveillance for Lyme Disease After Implementation of a Revised Case Definition — United States, 2022. MMWR, 2024. <https://www.cdc.gov/mmwr/volumes/73/wr/mm7306a1.htm>

R03. NIST. AI Risk Management Framework Core. Governance, mapping, measurement and management framework. <https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>

R04. Mitchell et al. Model Cards for Model Reporting, 2019; Breck et al. The ML Test Score, 2017. <https://research.google/pubs/model-cards-for-model-reporting/> and <https://research.google/pubs/the-ml-test-score-a-rubric-for-ml-production-readiness-and-technical-debt-reduction/>

R05. Angelopoulos and Bates. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. <https://arxiv.org/abs/2107.07511>

R06. FDA. Clinical Decision Support Software, final guidance, January 2026. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>

R07. HL7. FHIR R4 overview, version 4.0.1. <https://hl7.org/fhir/R4/overview.html>

Primary public data source register

The following publisher links identify the 12 stored-snapshot entries in the October 13 OpenDataDelta inventory. A link identifies the source family; the application metadata retains the selected observation period and field. These links do not assert a current upstream refresh or permission for every possible use. [E04]

D01. Reported tickborne disease cases by county, 2019–2022. County of residence and reporting-period constraints apply. <https://www.cdc.gov/ticks/data-research/facts-stats/geographic-distribution-of-tickborne-disease-cases.html>

D02–D03. Ixodes county status and pathogen county status through 2025. The two entries share a publisher page but represent different measures. Detection depends on surveillance effort. <https://www.cdc.gov/ticks/data-research/facts-stats/tick-surveillance-data-sets.html>

D04. NOAA nClimGrid-Daily. The product inventory records a selected May–July 2022 subset; the separate September backend used June 2023. These are distinct snapshots. <https://www.ncei.noaa.gov/products/land-based-station/ncimgrid-daily>

D05. CDC/ATSDR Social Vulnerability Index 2022. Community characteristics provide population context, not a personal social-risk diagnosis. <https://www.atsdr.cdc.gov/place-health/php/svi/svi-data-documentation-download.html>

D06. HRSA Area Health Resources Files 2024–2025. Selected fields have their own measurement years and denominators. <https://data.hrsa.gov/data/download>

D07. HRSA primary-care shortage designations. Facility, population and geographic designations require distinct interpretation. <https://data.hrsa.gov/topics/health-workforce/shortage-areas/dashboard>

D08. CDC BRFSS Long COVID supplementary table. State-level survey context remains separate from county or patient findings. <https://stacks.cdc.gov/view/cdc/174567>

D09. EPA annual county Air Quality Index. Period, monitor coverage and aggregation limit personal exposure interpretation. https://aqs.epa.gov/aqsweb/airdata/download_files.html

D10. CDC PLACES county data, 2025 release. Preserve measure, estimation type, source period and available uncertainty intervals. <https://data.cdc.gov/d/swc5-untb>

D11. CDC National Wastewater Surveillance System. Site coverage and population-level signal do not establish a patient infection. <https://www.cdc.gov/nwss/>

D12. CDC West Nile surveillance. Topic and date matching determine relevance; availability does not justify inclusion in every case. <https://www.cdc.gov/west-nile-virus/data-maps/>

Technical glossary

These terms describe different kinds of evidence and should retain their meaning across the product, research outputs and submission claims.

Term	Meaning in this dossier
Canonical observation	The shared, attributed representation used by multiple views; corrections preserve a coherent record.
OpenDataDelta	Recomputation comparing fixed patient evidence with and without selected public contextual layers.
Digital Twin	A dated evidence representation with reversible scenarios, not a physiological simulation.
FIPS	A standardized geographic code; county joins must retain five-digit identifiers and appropriate geographic vintages.
Source vintage	A particular release or revision of a dataset, distinct from its observation period and download date.
MAE and RMSE	Mean absolute error and root mean squared error. Their units and any target transformation must be specified.
ROC AUC	Ranking discrimination across thresholds. It is not the percentage of predictions that are correct.
Brier score	Mean squared error of predicted probabilities for a binary outcome; lower is better.
Calibration	Agreement between predicted probabilities and observed frequencies in a specified population.
Out of fold	A prediction produced for observations withheld from the model-fitting partition.
Ablation	A controlled comparison that removes a source family or feature block to assess its contribution.
Conformal prediction	A method producing prediction sets or intervals using a calibration sample, subject to assumptions.
Abstention	Withholding a single automated answer under a defined policy and routing the case for review.
TreeSHAP	Feature contributions explaining a tree-model prediction; attribution does not establish causation.
FHIR and USCDI	A health-data exchange specification and U.S. interoperability data-element frameworks; neither is itself a clinical model.
CDC and MEPS	Centers for Disease Control and Prevention; Medical Expenditure Panel Survey. Different sources and populations.
D1 and Worker	Database and application-runtime components used by the inspected Sites deployment.

Appendix A1 Lyme tickborne and related infectious surveillance

21 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G01-001	CDC Anaplasmosis surveillance	Public aggregate
G01-002	CDC Babesiosis surveillance	Public aggregate
G01-003	CDC COVID-19 case surveillance public-use data	Open
G01-004	CDC Ehrlichiosis surveillance	Public aggregate
G01-005	CDC FluView	Public aggregate
G01-006	CDC Lyme county case maps	Public view
G01-007	CDC Lyme Disease Dashboard CSV files	Open or public view
G01-008	CDC Lyme disease public-use aggregated data	Open or public view
G01-009	CDC Lyme public-use line-listed data 2022–2023	Open
G01-010	CDC National Wastewater Surveillance System	Public aggregate
G01-011	CDC Powassan virus historic data	Public aggregate
G01-012	CDC respiratory illness data channel	Public aggregate
G01-013	CDC Rocky Mountain spotted fever and spotted fever rickettsiosis data	Public aggregate
G01-014	CDC surveillance case definitions	Reference
G01-015	CDC West Nile virus surveillance	Public aggregate
G01-016	CDC WONDER Lyme disease tables	Verify current table
G01-017	Geographic distribution of tickborne disease cases	Open or public view
G01-018	Johns Hopkins Lyme and Tickborne Diseases Dashboard	Open/public view
G01-019	NNDSS weekly data	Open or public view
G01-020	State and local Lyme surveillance portals	Variable
G01-021	Tick Bite Data Tracker	Open or public view

Surveillance cases and aggregate reports describe reporting systems, definitions and geography of residence. They do not establish a complete care history or individual exposure. Original releases and case definitions must be checked before joins or trend comparisons.

Appendix A2 Vector animal and ecological data

22 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G02-001	Acorn mast and rodent ecology series	Study specific
G02-002	Blacklegged tick surveillance map	Open or public view
G02-003	CAPC canine parasite prevalence maps	Public view/partner
G02-004	Digitized historical deer densities	Open or public view
G02-005	GBIF occurrence downloads	Open or public view
G02-006	IDEXX veterinary testing data	Commercial/partner
G02-007	iNaturalist observations and images	Open or public view
G02-008	Ixodes county establishment/reporting table	Open or public view
G02-009	Ixodes pathogen county table	Open or public view
G02-010	National park deer-tick-pathogen integrated dataset	Open or public view
G02-011	NEON small mammal box trapping DP1.10072.001	Open
G02-012	NEON small mammal pathogen observations	Open, verify product
G02-013	NEON tick abundance and identification	Open or public view
G02-014	NEON tick pathogen status	Open or public view
G02-015	State deer harvest and population series	Variable
G02-016	State tick surveillance programs	Variable
G02-017	The Tick App research data	Partner
G02-018	TickReport passive surveillance statistics	Public view or terms
G02-019	TickTracker community reports	Verify or permission
G02-020	USGS white-tailed deer habitat and range layers	Open, select product
G02-021	Watauga County canine serology series 2017–2021	Public tables/author request
G02-022	White-tailed deer density estimates, 2008	Open or public view

Animal, vector, pathogen and ecological observations have different sampling designs and units. Presence is not abundance; a positive tick is not a human diagnosis; no record is not absence. Product IDs and reuse terms require release-specific verification.

Appendix A3 Lyme cohorts registries and biobanks

12 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G03-001	Geisinger Lyme EHR cohorts	Partner
G03-002	Johns Hopkins Lyme Disease Research Center repository	Partner
G03-003	LonghaulTracker	Verify or permission
G03-004	Lyme Disease Biobank	Partner
G03-005	Lyme Disease Biobank 2014–2023 three-month cohort	Author request
G03-006	Lyme Symptom Tracker	Verify or permission
G03-007	MyLymeData	Partner
G03-008	MyLymeData public chartbooks	Public aggregate
G03-009	National Disease Research Interchange Lyme tissue resource	Partner
G03-010	Pedi Lyme Net and pediatric biobank	Partner
G03-011	SLICE dbGaP study phs002793	Controlled
G03-012	SubLyme Network	Partner

A cohort, registry or biobank entry describes a possible research resource. It does not establish access, a partnership or permission to redistribute participant-level records. Governance and study-level eligibility precede any integration.

Appendix A4 Broader clinical cohorts claims and surveys

25 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G04-001	AHRQ social determinants database archived beta	Open archive
G04-002	All of Us Public Data Browser	Public aggregate
G04-003	All of Us Researcher Workbench	Access review
G04-004	BRFSS	Open
G04-005	CMS Master Beneficiary Summary File LDS	Controlled LDS
G04-006	CMS Medicare Limited Data Sets	Controlled LDS
G04-007	CMS Part D Prescription Drug Event RIF	Controlled RIF
G04-008	CMS Research Identifiable Files	Controlled RIF
G04-009	CURE ID treatment case reports	Public view/request
G04-010	HCUP state inpatient/ED/ambulatory databases	Controlled
G04-011	HCUPnet	Public aggregate
G04-012	Health and Retirement Study Medicare linked data	Controlled
G04-013	IDDO and SCCM VIRUS registry via CURE ID collaboration	Controlled
G04-014	Medicaid TAF/T-MSIS via CMS	Controlled
G04-015	Medical Expenditure Panel Survey	Open or public view
G04-016	Medicare claims via CMS CCW	Controlled
G04-017	Medicare Current Beneficiary Survey	Mixed
G04-018	MIMIC IV ED	Credentialed
G04-019	MIMIC IV hospital and ICU EHR	Credentialed
G04-020	National COVID Cohort Collaborative N3C	Controlled
G04-021	NHANES	Open
G04-022	NIAID AccessClinicalData@NIAID	Controlled
G04-023	OHDSI / OMOP observational data network	Partner
G04-024	PCORnet distributed research network	Partner
G04-025	PEDSnet pediatric EHR network	Controlled network

Clinical, claims and survey collections address different populations and observation processes. Public-use files may omit dates, geography or clinical detail. Link only through permitted identifiers and preserve survey design and denominator limits.

Appendix A5 Participant generated and digital data

12 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G05-001	Beiwe and Forest	Study platform
G05-002	Blue Button 2.0 beneficiary API	Participant authorized
G05-003	Consented patient timeline photos records and caregiver observations	Verify or permission
G05-004	ExtraSensory smartphone dataset	Open
G05-005	Google Data Portability API	Participant authorized
G05-006	Google Takeout individual exports	Participant authorized
G05-007	mPower mobile Parkinson study	Registered/controlled
G05-008	MyDataHelps EHR and sensor research collection	Study platform
G05-009	Open Humans participant authorized projects	Participant authorized
G05-010	Open Humans public data API	Open opt-in
G05-011	RADAR base	Study platform
G05-012	Sync for Science and patient directed FHIR exchange	Access method

Participant-generated sources require purpose-specific permission, minimization and attribution. Public visibility of a post does not authorize unrestricted collection or clinical inference. Connectors and consent paths require separate implementation.

Appendix A6 Climate land water and habitat

22 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G06-001	Annual National Land Cover Database	Open or public view
G06-002	Daymet	Open or public view
G06-003	EnviroAtlas national and community data	Open or public view
G06-004	Forest Inventory and Analysis DataMart	Open or public view
G06-005	GHCN-Daily	Open or public view
G06-006	LANDFIRE existing vegetation type	Open
G06-007	MODIS NDVI/EVI and land-surface temperature	Open or public view
G06-008	Monitoring Trends in Burn Severity	Open
G06-009	NASA MODIS MCD12Q2 land phenology	Open/registration
G06-010	NASA SMAP soil moisture	Open/registration
G06-011	NOAA 1991–2020 climate normals	Open
G06-012	NOAA Climate Data Online	Open or public view
G06-013	NOAA nClimGrid-Daily	Open or public view
G06-014	PRISM gridded climate data	Open or public view
G06-015	Protected Areas Database of the U.S. (PAD-US)	Open or public view
G06-016	US Fish and Wildlife National Wetlands Inventory	Open
G06-017	USDA Cropland Data Layer	Open
G06-018	USDA Forest Service tree canopy cover	Open, select product
G06-019	USDA NRCS gSSURGO soils	Open
G06-020	USGS 3D Hydrography Program	Open
G06-021	USGS LCMAP land change	Open
G06-022	USGS National Elevation Dataset/3DEP	Open or public view

Environmental products vary in grain, period, modeled versus observed status and spatial coverage. County aggregation requires documented geographic crosswalks and uncertainty. A contextual feature is not a patient measurement.

Appendix A7 Population health care access and geography

20 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G07-001	American Community Survey 5-year	Open or public view
G07-002	Area Health Resources Files	Open or public view
G07-003	BLS Quarterly Census of Employment and Wages	Open
G07-004	CDC Environmental Justice Index	Open
G07-005	CDC/ATSDR Social Vulnerability Index	Open or public view
G07-006	Census LEHD LODS	Open
G07-007	Census population estimates	Open
G07-008	Census TIGER/Line	Open or public view
G07-009	CMS Medicare Geographic Variation public use files	Open
G07-010	CMS Provider Data Catalog API	Open or public view
G07-011	Doctors and Clinicians National Downloadable File	Open or public view
G07-012	Doctors and Clinicians Utilization Data	Open or public view
G07-013	FCC National Broadband Map data	Open
G07-014	FHWA National Household Travel Survey	Open
G07-015	Health Professional Shortage Areas	Open or public view
G07-016	National Park Service visitor use statistics	Open
G07-017	NPPES NPI Files and Registry API	Open or public view
G07-018	PLACES county/tract estimates	Open or public view
G07-019	Rural–Urban Continuum Codes	Open or public view
G07-020	USDA Rural Urban Commuting Area codes	Open

Access, workforce and population data support contextual questions. Provider listing is not appointment availability, network participation or clinical quality. Denominators, geography and source years must remain visible.

Appendix A8 Trials literature omics and diagnostics

14 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G08-001	AACT clinical trials relational database	Open
G08-002	CDC Lyme disease rash image gallery	Open or public view
G08-003	ClinicalTrials.gov API	Open or public view
G08-004	FDA public IVD databases	Open or public view
G08-005	iNaturalist Open Data on AWS	Open or public view
G08-006	Kaggle erythema migrans rash collections	Verify or permission
G08-007	NCBI GenBank	Open
G08-008	NCBI GEO gene expression studies	Open/mixed
G08-009	NCBI Sequence Read Archive	Open/mixed
G08-010	NIAID Data Discovery Portal	Open or public view
G08-011	NIAID program collections and repository matcher	Public discovery
G08-012	NIH RePORTER API	Open or public view
G08-013	openFDA 510(k) data	Open or public view
G08-014	PubMed/PMC APIs	Open or public view

Literature, trials, omics and diagnostic resources are distinct evidence types. A reference or method is not a patient dataset. Accession-level rights, assay context and population applicability must be reviewed.

Appendix A9 Search media and public attention

11 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G09-001	Authorized community forum excerpts	Permission dependent
G09-002	GDELT	Public view or terms
G09-003	Google Health Trends research data	Verify or permission
G09-004	Google Trends API (alpha)	Verify or permission
G09-005	Google Trends public Explore	Public view
G09-006	Instagram hashtag content API	Verify or permission
G09-007	Lyme Twitter/X labeled corpus	Verify or permission
G09-008	OpenStreetMap / Overpass API	Open or public view
G09-009	Reddit Data API	Terms-dependent
G09-010	Wikipedia Pageviews API	Public view or terms
G09-011	X API v2 current platform	Terms/cost review

Search, media and attention data are vulnerable to publicity, sampling and platform changes. They may inform exploratory context but cannot be treated as independent clinical evidence or a diagnostic likelihood.

Appendix A10 Synthetic data standards and technical references

7 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G10-001	Annual NLCD derived edge and fragmentation metrics	Derived method
G10-002	Certified health IT API Service Base URL lists	Open or public view
G10-003	CMS CCW condition algorithms	Reference
G10-004	CMS synthetic Medicare claims resources	Open synthetic
G10-005	FHIR R4 and USCDI	Reference
G10-006	Lantern FHIR endpoint analytics	Open or public view
G10-007	Synthea and SyntheticMass	Open synthetic

Standards and synthetic-data tools provide structure and testing support. They are not real-world clinical outcome evidence, and use of a specification does not establish profile conformance.

Appendix A11 Local unpublished and proposed sources

5 catalog entries. Names and access labels are preserved from the 22 September 2026 source catalog. Access labels describe the catalog checkpoint, not current permission, acquisition or implementation. [E06]

Index	Resource or candidate	Recorded access
G11-001	Bee richness and flowering timing	Proposed
G11-002	Community tick collection workbook	Verify or permission
G11-003	FRED S&P 500 and VIX series	Open, peripheral
G11-004	Humidity and dew point series	Proposed
G11-005	Pollen counts	Proposed/mixed

These entries remain local, proposed, peripheral or incompletely verified at the catalog checkpoint. They are preserved to make the research inventory complete, not to imply integration or a scientific contribution.