

WholeContext AI

System Architecture and Engineering Design Specification

Document 3 in the technical review series

Kimberly Chulis PhD

Core Analytics

14 October 2026

Version 1.0

A traceable architecture for longitudinal evidence, source-grounded review and responsible AI development.

Review scope

Released synthetic MVP v39 • Shared observation model • Service and data contracts • Access controls • Public-data lineage • AI and image boundaries • Verification and delivery gates

This volume connects implemented software to its engineering evidence and the next development decisions. It documents how the system preserves patient identity, source meaning, chronology and human review across public context, scenario exploration and clinical packets.

Prepared for the HHS TOPx AI and Invisible Illness challenge and independent technical, clinical and research scrutiny. Institutional endorsement, regulatory clearance and clinical effectiveness are not claimed.

Public demonstration

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

Companion volumes

Technical Clinical and Impact Dossier

171 Resource Data and Engineering Matrix

1 Document control and evidence status

This specification describes the released synthetic MVP and defines the engineering work required for a controlled healthcare pilot. It is an implementation review and target design, not a certification or a clinical validation report.

Control	Baseline
Document	WCT ARCH 003 • Version 1.0 • 14 October 2026
Author and accountable product lead	Kimberly Chulis PhD • Core Analytics
Release inspected	Public release v39; source commit below. Later saved source is excluded from this baseline.
Source commit	4a85dc7a60a989e2c731cb9ed2ed29ffdba7cbe1
Companion evidence	Document 1: Technical Clinical and Impact Dossier. Document 2: 171 Resource Data and Engineering Matrix.
Audience	Software and data engineers, AI researchers, clinical reviewers, privacy and security reviewers, challenge judges.

Evidence vocabulary

Implemented means a behavior is visible in the pinned source. Verified means a named saved test or release record supports it within its stated scope. Required means an acceptance criterion proposed here. Planned means a future capability without a completed implementation claim.

Source inspection supports architecture statements; it does not prove that every deployment configuration, host control or user journey behaves as intended. Previous test results are identified as historical evidence. No new production penetration test, load test, clinical study or compliance audit was performed for this document.

The normative “shall” statements below are proposed pilot requirements. They are not assertions that v39 satisfies every requirement. Requirement ownership is assigned by role; staffing and approval signatures remain to be established. References [E1–E8] identify project evidence; [R1–R7] identify external review frameworks.

2 Reader navigation

The document follows an expert review path from scope and implemented architecture to evidence, gaps and an executable development sequence.

Review area	Sections	Primary question
Scope and structure	3–6	What exists, where does it run, and what is deliberately bounded?
Data and permissions	7–11	What is the source of truth and who can access each record?
Interfaces and concurrency	12–15	What does each service accept, return and persist?
Context and AI methods	16–21	How do public data, dated scenarios, text and images affect outputs?
Exports and client state	22–23	Can a reviewer trace a packet to the exact selected evidence?
Research and governance	24–27	Which models are research, and what protects a future pilot?
Operations and verification	28–32	How are failures, releases and acceptance assessed?
Decisions and delivery	33–36	What changes next, in what order, and how will success be measured?
Risks and references	37–40	What remains unresolved, and which artifacts support each claim?

How to use the companion matrix

The 171 Resource Data and Engineering Matrix is the source-level register. Its engineering contracts, derived-variable definitions, output definitions and join rules provide the detailed data specifications referenced here. Resource counts represent different inventories and must not be added together.

Review package boundaries

This architecture volume does not replace a model card, statistical analysis plan, clinical protocol, privacy impact assessment or operational runbook. It identifies the contracts those artifacts must satisfy and the evidence needed to approve each stage.

The public demonstration is available at <https://wholecontextai.kchulis.chatgpt.site/> . It uses fictional cases. The public package contains technical evidence and limitations without credentials, personal records or deployment secrets.

3 Intended use and system boundaries

WholeContext organizes longitudinal evidence, daily function, care questions and relevant public context so a person and a reviewer can understand what is known, what is missing and what needs review.

Implemented surface	Architectural boundary
Fictional patient exploration	65 synthetic identities; selected records and authored scenarios support repeatable demonstrations.
Personal workspace service	Authenticated record, membership and task operations exist. Their existence is not authorization to process real health data in this release.
Evidence and packets	Source-linked summaries, dated observations and selected exports; original attachments are not automatically embedded.
OpenDataDelta	Deterministic matching of stored public snapshots; county or source-native geography remains explicit.
Digital Twin candidate	Dated reconstruction and limited authored hypotheses; no physiological prediction or validated disease probability.
Optional text AI	Bounded synthetic-source categorization with exact spans, human review and deterministic fallback.
Photo workflows	Local preparation, simple capture heuristics, storage validation and human review; diagnostic image analysis is disabled.

Public demo availability: demo-shell.js permits the synthetic review and selected public retrieval routes, blocks other API calls in the sample UI, and redirects connected workspace, intake, source-review and upload pages to the demo guide. Backend code presence is therefore separate from public feature availability. The browser guard is not a server authorization control. [E2]

Explicit exclusions

The inspected release does not establish autonomous diagnosis, treatment selection, real-time clinical monitoring, EHR writeback, validated individual risk prediction, cross-account fixture synchronization, external email delivery or unrestricted patient-record AI. The capability endpoint declares personal-record AI, diagnostic image analysis and external messaging false. [E1]

Clinical responsibility remains with the treating professional. A source display or empty safety section cannot establish that a person is safe. Public prevalence, directory results and hypothetical observations must not become individual clinical facts. These are system invariants, not merely interface wording.

4 Architectural drivers and quality requirements

The most important engineering properties are identity integrity, provenance, explicit uncertainty and bounded failure. The pilot must demonstrate them before adding broader automation.


ID	Required quality attribute	Acceptance evidence
REQ 01	A record, observation and export shall retain one patient identity.	Cross-patient negative tests; adversarial mixed snapshots; export identity checks.
REQ 02	Every material output shall resolve to selected source records or a labeled public layer.	Lineage completeness report; source removal and correction tests.
REQ 03	Missing, excluded, invalid and hypothetical values shall remain distinct from measured zero.	Boundary fixtures and snapshot comparisons.
REQ 04	Access shall depend on current membership, role, recipient and record scope.	Multi-account allow and deny matrix on the hosted service.
REQ 05	Concurrent edits shall not silently overwrite reviewed work.	Two-editor conflict tests and explicit conflict recovery.
REQ 06	AI failure shall preserve the source and disclose fallback.	Timeout, invalid output, quota and provider outage tests.
REQ 07	A dated scenario shall exclude later patient facts and isolate assumptions.	Date sweep; reset; export exclusion; source eligibility tests.
REQ 08	A release shall be reproducible and reversible.	Pinned source, data manifest, dependency inventory and restore exercise.

These requirements prioritize safety of interpretation and recoverability over feature count. Responsiveness, accessibility and operating cost are also necessary, but their numeric service targets are proposed in Section 29 and require measurement. No production availability or latency achievement is claimed.

The design uses a modular application rather than independently deployed microservices. Shared domain logic reduces discrepancies between charts, scenario review and packets. Deployment simplicity trades against larger regression scope when a shared module changes. [E1, E2]

5 Runtime architecture and trust boundaries

The implemented system combines browser assets, a server-side Worker router, bound database services and selected external public services. Public synthetic review and authenticated persistence are separate paths.



Connected care and extraction are backend paths; the public sample UI disables these connections.

Figure 1. Observed v39 runtime boundaries. Dashed optional service is configuration-dependent; research pipelines are outside the request path.

The root route resolves to portal.html. Static HTML, CSS, JavaScript, public datasets and media are served through the asset binding. The Worker routes review, care, extraction, research, provider and education requests. D1-style prepared statements and batches are used through env.DB. [E1]

The host supplies authenticated visitor identifiers to care and extraction routes. The application relies on the host to authenticate and protect that identity channel. Header authenticity, audience enforcement, TLS, region, backup and encryption settings require host evidence; they cannot be established from the application source alone.

Optional synthetic extraction sends bounded source content to an external model service only after configuration and quota checks. Public directory and literature requests use their own scoped adapters. The runtime does not load the offline ecological or synthetic research models into a clinical inference service.

6 Component responsibilities

The source is organized as small JavaScript modules with shared domain logic. The table identifies responsibility and the contract that constrains each component.

Component	Responsibility	Contract or limit
index.js	Route dispatch and capability probe	Capability probe does not execute a model call.
demo-review.js	Hydrate allowed fixture selections; run review operations	Read-only; no care database writes.
care.js and circle-channel.js	Persist reviewed workspace and support activity	Current authorization and optimistic revisions.
observation-model.js	Canonical records, measurements and time series	Source, unit, date and missingness preserved.
packet-model.js	Selected evidence projection	Patient identity, current revisions and source references.
open-data-delta.js	Public-source matching and context changes	No alteration of patient findings; no numerical pooling.
digital-twin-core.js	Dated baseline and hypothetical comparison	Synthetic identity; assumptions excluded from clinical export.
source-extraction.js and extract.js	Exact-span indexing and optional category assignment	No rewritten source text or automatic record acceptance.
public-services.js, research.js, education.js	Scoped retrieval from external sources	Provenance, coverage and explicit failure state.
Browser renderers	Display, interaction, charts, media and PDF download	Shared model contracts; stale output invalidation.

These are logical service boundaries inside one deployment, not proof of network isolation between modules. Changes to shared observation or packet logic require regression across all consuming views. Naming a file as an architectural component does not imply a separate team, process or deployed container. [E1, E2]

7 Evidence flow and review lifecycle

The evidence flow separates source material, candidate structure, human selection and derived output. A generated result is not a new source record.




Figure 2. Logical evidence lifecycle. Human selection controls derived review; persistence requires a separately authorized action.

Transition	Implemented control	Required extension
Source to candidate	Exact text spans or authored fixture candidates	Formal schema version; parser test corpus.
Candidate to reviewed record	Review confirmation and source-span validation	Actor-attested approval ledger and durable provenance.
Record to output	Current record selection and shared observation model	Snapshot digest covering every output dependency.
Correction to refresh	Revision increments and stale-result guards in selected views	Uniform invalidation and conflict UX across every consumer.

Context Discovery demonstrates consent and accepted or dismissed candidate states using authored email, image and caregiver examples. Its source provenance explicitly marks `modelExecuted` false. It is not an email connector or a production multimodal extraction pipeline. Withdrawing a sample permission removes that sample's accepted decisions. [E1]

8 Canonical observation model

The canonical model is a shared read model derived from selected records, not a normalized clinical database or a certified terminology implementation.

Field group	Canonical fields and semantics
Identity	patientId, recordId, canonicalRecordId, observation id; identity and correction lineage remain separate.
Measurement	metric, label, value, unit; invalid or absent values use null with missingReason.
Time	date, originalDate, time, timezone, effectiveDateTime; source date is retained.
Provenance	sourceType, author, sourceId, sourceLabel, field, provenance, ref.
Revision and context	revision, updated, context, synthetic, epistemic; authored classifications are not independent clinical verification.

Read coverage and write coverage

The shared model defines 22 metrics, including symptom, function, sleep, activity and related reported measures. The care wellness write validator accepts six fields: function0to10, fatigue0to10, pain0to10, sleepHours, steps and restingPulse. These are different contracts. Extending persistence to additional metrics requires schema and validation work. [E1, E2]

Reconciliation rules

Current records are selected by patient, canonical identity, revision and update ordering. Excluded, deleted, superseded, rejected and hypothetical observations are withheld from ordinary measurement series. A corrected blank value replaces the previous reading; it is not permission to recover an older value silently.

Series remain separated by metric, unit, source type, context and timezone. Device-derived sleep is not silently merged with self-report; pounds and kilograms are not automatically converted. Daily and weekly aggregates summarize recorded readings with counts and source observation identifiers. Missing days are not imputed.

9 Time revision and lineage semantics

Clinical chronology, record revision time and public-source period answer different questions. The application must preserve all three.

Time or identity	Meaning	Rule
Event date	When the source reports an event	Retain original text when missing or uncertain.
Recorded or updated time	When the application stores or revises a record	Do not use it as a substitute event date.
Scenario asOf	Latest patient date admitted to a reconstruction	Exclude later dated records; count undated sources separately.
Source native period	Period represented by a public observation	Keep survey, surveillance and denominator periods explicit.
Retrieved or checked date	When a source was obtained or inspected	Does not establish publication date or historical availability.
Canonical record identity	Links projections and corrections	Retain source identifier and revision; prevent duplicate measurement.

Change comparison

The shared comparison routine uses equal calendar windows around a recorded change, with a configurable window bounded to 1–90 days and a 14-day default. Same-day readings are separated when ordering is unknown. It reports means, counts and recorded days by source; it does not estimate causal treatment effects. [E2]

Required temporal tests

REQ 09: the pilot shall cover leap days, invalid dates, daylight-saving transitions, timezones near midnight, missing event times, revision ties, correction cycles and late-arriving records. Weekly grouping currently uses date-based Monday bins; clinical timezone semantics need explicit acceptance tests.

A historical scenario reconstructs selected patient facts as of a date. Public snapshots are retrieved now; source observation cutoffs do not reconstruct what information was actually available to a clinician at that historical time.

10 Persistence model and schema ownership

The current schema stores flexible JSON record payloads with relational ownership and membership. Support conversations are defined in a separate SQL migration.




Figure 3. Implemented logical relations. Arrows run from parent to dependent tables through declared references; runtime foreign-key enforcement must be verified separately.

Entity	Key fields and purpose
care_patients	id, owner, profile JSON, created; owner is the primary workspace authority.
care_members	id, patient_id, email, role, created; application membership controls access.
care_records	id, patient_id, author, kind, payload JSON, support, clinical, revision, created, updated.
care_audit	id, patient_id, actor, action, record_id, created; selected application action history.
circle_threads and circle_messages	Recipient-scoped support conversation, status, revision and ordered messages.
ai_daily_usage	Composite actor and UTC day key; call reservations for actor and service budgets.

JSON flexibility supports the MVP, but schema drift, migration completeness and validation consistency need explicit ownership. care_audit is not proven immutable, comprehensive or transactionally coupled to every mutation. No general record erasure endpoint or object-storage tier is established in the inspected design. [E1, E3]

11 Authorization and recipient scope

Authorization is evaluated on the server for saved workspaces. A role alone does not grant every record; record flags, recipient restrictions and current membership also matter.

Action or data	Owner	Support member	Clinician member	Payer member
Workspace records	All owned	Support-scoped	Clinical-scoped	Approved financial projection
Owner entries and sharing	Create and revise	No	No	No
Assigned task status	Manage	Assigned task only	Assigned task only	No
Reviewed clinical note	Owner record controls	No	Own note; authorized sources	No
Clinical packet route	Clinical-scoped subset	No	Authorized clinical subset	No
Support conversation	All own threads	Recipient thread only	No	No
Membership changes	Yes	No	No	No

The payer response contains an explicitly approved snapshot of allowed financial fields. It excludes unrestricted clinical narrative. Later cost edits do not automatically replace a previously approved payer snapshot. The owner must confirm a new approval. Permission changes return a saved state; they do not send an email. [E1]

Revocation lifecycle

Removing a member prevents current membership-based access. Payer approvals for that recipient are removed. Existing support and clinical flags may remain on records, so re-adding the same recipient requires an explicit policy and regression test for prior grants. Previously downloaded files cannot be recalled by changing a server permission.

REQ 10: pilot grants shall be individually identifiable, reviewable, revocable and auditable, with explicit expiry and regrant semantics. Site audience access and application membership are distinct controls; both must pass hosted multi-account tests.

12 Public and optional service contracts

This catalog describes the routed interfaces inspected in v39. It is a human-readable contract inventory; a versioned OpenAPI artifact remains required before a pilot.

Route and method	Input and output	Important boundary
POST /api/demo-review	operation plus bounded fixture selection; returns registry, packet, evidence, guide, delta, twin or overview	Public synthetic review; server hydrates known source variants.
GET /api/capabilities	No clinical input; reports configuration and schema readiness	Probe performs no model call; readiness is not a validated AI result.
POST /api/extract	Labeled synthetic source, optional categories; exact spans, categories and trace	Signed-in visitor; bounded optional model use; review required.
GET /api/providers	ZIP, specialty and bounded directory paging; source-linked candidates	Directory presence does not establish access, coverage or availability.
GET /api/research	Predefined topic and kind pubmed or trials; up to five results	Rejects additional query keys; no patient narrative accepted.
GET /api/education	Supported general topic; MedlinePlus source results	General education; not personalized clinical advice.
/api/lab	Returns unavailable status	Conversational AI is not connected to Discovery Lab.

Research and provider caches are in-memory process caches, not durable materialized views or guaranteed cross-instance caches. A provider query covers one state and one 50-candidate page; distances use ZIP representative points. A successful source response is not a completeness guarantee. [E1]

13 Care service contract inventory

All routes below use the /api/care/patients prefix and require host-authenticated identity plus a connected database. Mutation requests require a matching Origin.

Suffix and method	Behavior	Control
Root GET and POST	List permitted workspaces or create a profile	Archived filtering; profile validation; owner attribution.
/intake POST	Create reviewed synthetic profile and 1–12 visit or evidence entries	Deterministic request-derived identity and duplicate detection.
/{{id}} GET	Return role-scoped workspace projection	Current member role and record recipient filters.
/{{id}}/profile PATCH	Revise or archive profile	Owner only; archive is not erasure.
/{{id}}/records POST or PATCH	Create or revise an allowed record; task update by assignee	Payload validation; recipient scope; expected revision.
/{{id}}/sharing POST	Grant selected records to a reviewed recipient	1–100 records; expected revisions; current recipient membership.
/{{id}}/members POST or PATCH	Set role or revoke membership	Owner only; no external notification.
/{{id}}/clinical-notes POST or PATCH	Write reviewed clinician entry or correction	Clinician attribution; authorized source IDs; own-note correction.
/{{id}}/packet GET	Return selected clinical evidence projection	Owner or clinician; only clinical-authorized records.
/{{id}}/circle GET POST PATCH	Read, create or update support threads	Owner or current support recipient; status revision check.

The retired importHandoff function is a no-op; the patient-list path does not implicitly import old seed records. Supported top-level methods are GET, POST and PATCH. Unsupported actions are rejected; no general DELETE contract exists. [E1, E3]

14 Public review request and response rules

The public review route accepts selection instructions for known fictional records. It does not accept an unrestricted patient history as authoritative source material.

Contract element	Implemented rule
HTTP envelope	POST with application/json. Streamed body limit: 180,000 bytes. If Origin is supplied, it must match.
Patient and records	Known synthetic patient; records array of at most 400 selections. Unknown source identities and unsupported patch fields fail.
Record patch	Allowlisted measurement, date, time, timezone, revision, selection and source-projection fields. Source text is reconstructed by the server.
Revision	Positive integer from 1 through 100,000 for the allowed synthetic patch path.
Context options	Five-digit county FIPS; up to 160 layer identifiers of 1–100 permitted characters; question up to 240 characters.
Twin assumptions	Up to seven hypothesis identifiers, further constrained by the supported case.
Context loading	State-scoped compressed context; decompressed stream capped at 3 MiB; load errors remain visible.
Response	No-store JSON; public projection removes internal-only fields; no record mutation.

Selection semantics

Omitting `selectedLayerIds` requests the relevant available context layers. Passing an empty array requests patient evidence without selected public context. A UI must preserve this distinction rather than treating both as a default. Unknown selections are reported rather than silently converted into valid context. [E1, E2]

REQ 11: freeze operation-specific request and response schemas, add explicit `schemaVersion`, and test real browser-produced payloads against them. A previous release repair demonstrated that hand-written service tests alone missed a valid underscore-containing layer identifier used by the client. [E4]

15 Mutation concurrency and idempotency

The care service combines expected-revision updates, selected batched writes and operation-specific duplicate detection. It does not implement a universal transaction or idempotency protocol.

Concern	Observed behavior	Pilot requirement
Lost updates	Record, note, task and thread status updates condition on revision; conflict returns 409.	REQ 05: retain user edits, fetch latest revision and require explicit reconciliation.
Multi-record sharing	All requested revisions are checked; guarded update prevents partial scope changes on mismatch.	REQ 12: test race conditions and bind audit event to the same transaction.
Duplicate intake	Request-derived identity supports duplicate detection for equivalent normalized intake.	Document canonicalization and retry behavior; test concurrent repeats.
Wellness import	importKey and deterministic record identity prevent duplicate measurements.	Distinguish ignored insert from successful creation in audit semantics.
General create	UUID creation for ordinary records; no generic request key.	Add operation-scoped idempotency keys before automated import retries.
Audit coupling	Several updates are followed by separate audit inserts.	Durable transactional audit or outbox; reconcile failed audit writes.

The record PATCH implementation sets support and clinical flags from explicit true values in the request. Omitting them can clear those flags, even when a payload is being edited. Clients must send the intended sharing state; the future contract should separate content edits from permission edits or make patch semantics unambiguous. [E1]

The care body limit is checked after reading text and uses a 200,000-character threshold. Extraction similarly checks text after reading. These differ from the streamed byte limit on public review. REQ 13 requires consistent streaming size enforcement, bounded JSON depth and parser tests across all write routes.

16 OpenDataDelta source matching

OpenDataDelta adds context and investigation questions around selected patient evidence. Its output explicitly states that patient findings remain unchanged.

Stage	Implemented transformation
Select patient evidence	Use selected source records and the current question; retain before findings, source IDs and date span.
Resolve geography	Match exact county FIPS. State estimates remain state estimates. Missing geography produces an explicit unavailable result.
Match relevant layers	Deterministic topic and source-measure matching; includeAllAvailable can expose broader source coverage.
Preserve observation meaning	Retain field, value, unit, geography, native period, population, citation, coverage and missingness.
Check time and population	Label historical or later context. Withhold adult-only survey context for documented pediatric ages.
Apply selected context	Return selected layers, unmatched identifiers, missing layers and changesPatientFinding false.
Return bounded interpretation	Context may change a review question; it does not change diagnosis, risk probability or treatment causation.

No cross-source numerical pooling is performed. Source years, populations and units remain distinct. An absent observation is not zero. Confidence intervals are retained where supplied by the source. Publication date may be unknown; a snapshot check date is not substituted for it. [E2]

Concrete reproducibility check

The saved Hennepin example reports annualized Lyme incidence of approximately 15.23 cases per 100,000 residents per year from 2019–2022 surveillance with a 2022 denominator. The same review retains the 2018–2022 ACS/SVI period for 8.9% of households without a vehicle. These are source-specific public observations, not individual exposure or transport status. [E4]

17 Data engineering and catalog lineage

The catalog, integration atlas and runtime source registry represent different engineering stages. Keeping their identities distinct is essential to an accurate submission.

Inventory	Count	What it establishes
Original resource catalog	171	Cataloged resource records preserved in Document 2.
Integration atlas	107	Planning and integration metadata; original README identifies version 1.3 dated 29 August 2026.
Runtime registry	120	12 stored snapshots, 90 catalog-only entries and 18 unavailable or restricted entries.
Live Delta adapters	0	No upstream live-refresh adapter is claimed by this registry. Other public services are separate.
Lineage mappings	198	Links among inventories; mappings are not additional datasets.
Engineering specifications	19 contracts	Alongside 54 derived definitions, 16 outputs and 21 join rules.

Eleven of the twelve stored runtime snapshot sources map to eleven entries in the exact 171-resource catalog. EPA annual county AQI is the additional runtime source outside that exact catalog. Counts must not be summed into an inflated integration total. [E5]

Required ingest contract

REQ 14: each refresh shall produce an acquisition manifest, source license or access basis, source URL, retrieval time, content hash, schema version, geography and period definitions, row counts, duplicate and missingness reports, transformation version and output digest. Failed validation quarantines the new snapshot; the last accepted snapshot remains identified.

The county assembly covers 3,143 county rows in the saved public context artifact. Field availability varies. A county row does not imply every source has an observation for that county. Refresh cadence and permissions must be specified per source rather than inferred from catalog inclusion.

18 Dated scenario and Digital Twin design

The Digital Twin candidate is a source-based reconstruction of a synthetic case. It exposes how selected facts, public context and hypothetical assumptions affect a review display.




Figure 4. Separate branches preserve a dated patient baseline while allowing labeled context and hypothetical comparisons.


Normalization requires a synthetic case identity, rejects mismatched patients and explicitly non-synthetic records, selects current accepted evidence and freezes the baseline object. Only dated records at or before the selected baseline enter the dated reconstruction. Partial fact selection withholds record-level derived flags and measurements. [E2]

The focused authored case supports up to seven hypothetical observations. Other cases use source-based review and do not inherit unsupported condition-specific assumptions. Scenario reset makes no record or database writes. Export includes the dated baseline and selected public context; hypothetical findings are excluded from clinical records.

Public layers with an observation period beyond the baseline are ineligible. Same-year sources without an exact cutoff are excluded from the dated scenario. This conservative rule prevents some temporal leakage, but does not reconstruct historical publication availability. No physiological simulation, causal treatment response or calibrated prognosis is claimed.

19 Optional AI execution architecture

The optional model path organizes exact synthetic source segments into allowed categories. It has no tools, record-write authority or sharing authority.



Disabled configuration, quota denial, timeout or invalid provider output selects fallback.

Figure 5. Bounded optional model execution. All outcomes return reviewable source spans; record acceptance is a later action.

Execution control	Inspected implementation
Input	8,000 source characters; at most 100 segments; labeled synthetic source and signed-in visitor.
Categories	Six defaults plus up to six valid custom category labels.
Budget	10 reserved attempts per actor per UTC day and 100 per service day. Reservations are not refunded.
Provider call	Configured model; one request; 30-second timeout; 2,000 output-token maximum; store false.
Validation	Strict structured category output; unique valid segment IDs; every source segment represented exactly once.
Fallback	Configuration, quota, schema, timeout and provider failures return explicit deterministic segmentation.

The provider output supplies category assignments, not replacement quotations. The server reattaches original spans. The response trace records mode, limits and review requirement; changedRecords remains zero. A store false request flag does not by itself establish contractual data retention, residency or eligibility for real health data. [E1]

20 NLP source contract and evaluation

The current source indexer is conservative exact-span processing with lexical cues. It must not be represented as a clinically validated NLP model.

Element	Current method	Required validation
Segmentation	Regex-based source segments with start, end and quote	Punctuation, abbreviations, lists, Unicode and empty segments.
Mention cues	Literal dates, test and medication terms, negation, uncertainty and attribution cues	Expert-labeled precision and recall; false associations and missed cues.
Date handling	Valid ISO dates recognized; incomplete dates remain literal	No invented year; no date carried across sentences without review.
Category assignment	Deterministic defaults or optional constrained model labels	Category confusion matrix; model and fallback reported separately.
Acceptance	Reviewed flag; exact source substring; valid offsets and unique span positions	Mutation and offset tampering; missing source; duplicate acceptance.

Span representation

Offsets originate from JavaScript string indices. The cross-language contract shall specify UTF-16 code units or provide a tested conversion. Emoji, combining marks and non-Latin text are required fixtures. Exact substring equality is a mechanical provenance check; it does not establish clinical correctness or correct attribution. [E1]

Proposed evaluation protocol

REQ 15: assemble a consented or synthetic, independently annotated corpus with patient-separated splits; define target labels before evaluation; adjudicate disagreements; report per-class metrics, uncertainty, failure examples and review burden. Include negation scope, historical versus current symptoms, caregiver attribution and conflicting dates.

A production parser or model change must preserve source spans and pass regression before promotion. Improvements in category accuracy must not be marketed as improved diagnosis or patient outcomes. Clinical usefulness requires a separate prospective study.

21 Image handling and computer vision boundary

Image handling is implemented; diagnostic computer vision is not. The architecture preserves that distinction in the capability response, stored record and user-facing review state.

Layer	Implemented behavior
Browser selection	PNG or JPEG under 10 MB; decode locally; resize longest side to at most 900 pixels.
Prepared copy	Draw onto a canvas with a white background and encode JPEG; reject a copy that remains above the client size limit. Original metadata is not carried by the canvas encoding path.
Technical sample	Small pixel sample measures brightness, variation, near-black and near-white fractions. Flags prompt human inspection.
Server validation	PNG or JPEG data URI; decoded payload 24–131,072 bytes; file signature and dimensions checked; at most 4,096 pixels per side and 16 million pixels.
Saved record	Upload-only analysis status and needs-human-review state; base64 image in JSON payload; authorization inherited from the record.

Low brightness variation is not a blur estimate. Passing technical thresholds does not establish diagnostic suitability. Header and dimension validation is not a complete server-side decode, safe transcoding service or malware scan. Browser preparation is useful but cannot substitute for server trust-boundary validation. [E1, E2]

Planned image architecture

REQ 16: introduce a quarantined object store, content digest, verified decoding and re-encoding, metadata policy, restricted retrieval, retention and deletion propagation before larger uploads or real-data use. Any OCR, lesion interpretation or multimodal model needs a distinct intended use, licensed dataset, patient and site separation, reference labels and subgroup evaluation.

A future CV report must include performance by device, lighting, skin tone, site and image quality; uncertainty and abstention; external validation; and failure examples. No sensitivity, specificity, tick identification or rash classification performance is claimed for v39.

22 Clinical packet and export contracts

A packet is a versioned view of selected evidence. Its trustworthiness depends on identity, selection, source references and reproducible rendering rather than the length of its narrative.

Contract	Current behavior or required evidence
Patient identity	packetModel rejects explicitly mismatched patient records; canonical reconciliation further constrains current records.
Selection and chronology	Excluded facts and future-dated records are withheld according to the selected view and cutoff. Rejected assumptions do not enter clinical export.
Source references	References such as S1 are assigned within the packet. They are local presentation labels, not globally persistent source identifiers.
Evidence structure	Source facts, medications, allergies, negatives, uncertainties and priorities are organized with original source links.
Human review	A clinical assessment is included only when its reviewed state is true. The flag alone is not an independently attested clinician credential.
Rendering	Browser PDF generation uses the shared packet model and a PDF library; selected Digital Twin outputs check for scenario changes during rendering.

Required export manifest

REQ 17: a pilot packet shall carry patient identity, generation time, record IDs and revisions, source hashes, selected public layer IDs and versions, time cutoff, renderer version and an export digest. It shall state omitted sources and whether original attachments are included. A complete machine-readable manifest is a target requirement, not a current universal feature.

Acceptance must compare rendered PDF text and charts to the exact selected snapshot, test long and multilingual source content, inspect pagination, and confirm that a revoked record is excluded from a newly generated packet. Previously downloaded packets remain outside server recall. [E1, E2, E4]

23 Browser state and collaboration boundaries

The synthetic demonstration and saved care workspace have different state models. Same-browser handoff must not be presented as durable cross-account synchronization.

State	Observed or required behavior
Fictional case interaction	Browser modules manage selected synthetic context and presentation state. These interactions do not establish a saved clinical record.
Saved care record	The care API is authoritative. The clinical workspace adapter refetches authorized packet data and checks that the selected patient and role remain appropriate.
Pending asynchronous work	Photo selection and workspace loading use generation counters; Digital Twin PDF generation checks result and revision identity before offering download.
Support conversation	Persisted in-app messages and status changes for current support recipients. Responses explicitly report no external notification.
Clinical handoff	The v39 release record describes the demonstration handoff as same-browser. Cross-account synchronization and background delivery are not claimed.

Required state discipline

REQ 18: maintain one explicit active-patient context; invalidate requests, selections, charts and downloads on patient or permission change. Never reuse an earlier patient's asynchronous response. All saved mutations must recheck authorization server-side, regardless of browser controls.

A pilot browser-state inventory shall classify each local key and cache by purpose, sensitivity, expiry and deletion behavior. Sensitive records must not be placed in persistent browser storage without an approved design. Test reload, back navigation, multiple tabs, expired sessions, role changes and interrupted downloads. [E1, E2, E4]

In-app task completion or message storage does not prove that a clinician saw the content. Delivery, acknowledgment and escalation require separate product contracts and staffed operating procedures. The current system is not a monitored clinical inbox.

The sample shell namespaces browser state by a session identifier; selected clinician handoff reads a patient-specific localStorage entry within that namespace. Reset attempts to clear the sample state. These mechanics support a fictional browser session, not a real-data tenancy or consent boundary. [E2]

24 Research models and production separation

The research program demonstrates data engineering and experimental rigor without turning ecological or synthetic results into patient-level product claims.

Research workstream	Recorded method	Deployment interpretation
County ecological modeling	3,143 counties; grouped geographic evaluation; nested training transforms; persistence and boosted-tree comparisons.	Research on county surveillance outcomes; no individual diagnostic endpoint.
Federal data assembly	County/year normalization, denominator matching, climate linkage and reproducibility checks.	Evidence for a repeatable research pipeline; separate vintages from runtime snapshots.
Synthetic uncertainty	Independent train, calibration, IID and shifted test sets; conformal prediction sets and review routing.	Method experiment with artificial administrative labels; no patient risk calibration claim.
CDC case and MEPS methods	Designed acquisition and modeling scaffold with distinct populations and survey considerations.	Planned analyses; no fabricated person-level linkage or unexecuted performance result.
NLP and multimodal development	Implemented exact-span text path and image handling; broader models planned.	Separate validation and intended-use approval required for each function.

Promotion boundary

REQ 19: no research model shall enter a user-facing decision path without a model card, licensed training provenance, frozen evaluation, independent review, known failure limits, subgroup results, versioned serving contract, rollback and monitoring plan. Comparison against a simple baseline is mandatory.

Document 2 records twelve AI and ML method tracks and nine October research assets. Training code, dataset versions, seeds, splits and metrics belong in a reproducibility package. The architecture uses those studies as evidence of completed work and a basis for next experiments, not as proof of clinical benefit. [E5, E6]

25 Security architecture and threat model

The threat model covers identity, record scope, untrusted sources, uploads, external services and exported copies. Existing controls are useful evidence; residual requirements remain explicit.

Threat	Observed control	Required evidence or change
Wrong identity or patient	Host identity dependency; server patient and role checks	Host assertion integrity; multi-account isolation; session expiry tests.
Unauthorized record access	Membership, sharing flags and recipient filters	Complete route matrix; revocation and regrant tests; least-privilege host access.
Source instruction injection	Model treats source as data; fixed schema; no tools or writes	Adversarial corpus; no hidden output-induced actions; category integrity.
Unsafe or oversized input	Allowlisted fields; size and file checks	Streaming limits on all routes; decode and quarantine; dependency review.
Stale or conflicting state	Expected revision and selected client generation checks	Uniform invalidation and recoverable conflict workflow.
Sensitive disclosure	Scoped projections and no-store responses on key APIs	Logs and telemetry review; export controls; browser cache inventory.
Supply-chain compromise	Pinned source and preserved release evidence	Locked dependencies, SBOM, provenance, secret scan and signed release manifest.

REQ 20: security review shall map chosen OWASP ASVS controls to code, infrastructure evidence and executable tests. NIST SSDF supplies the development-process baseline. Neither framework reference constitutes verification or certification. [R2, R3]

The synthetic public deployment must not accumulate real records through an accidental form submission. A pilot needs explicit environment separation, server-enforced data policy, rate and abuse controls, credential rotation, incident response and independent security assessment. [E1]

26 Privacy retention and data governance

The pilot design must establish the authority to collect, process, share and retain each data class before enabling real health information.

Data class	Governance requirement	Disposition
Public source snapshots	Document license, access limits, attribution, version and permitted derivatives	Refresh and archive by source policy; retain lineage.
Fictional case data	Maintain synthetic designation and generator or author provenance	Keep separate from real participant environments.
Personal health records	Purpose, consent or other applicable authority, access scope and approved processors	Minimum necessary fields; approved retention and erasure process.
Images and attachments	Consent, metadata policy, content access and derivative inventory	Delete original and derivative copies according to approved policy.
AI requests and traces	Approved input class; processor terms; content-minimizing trace design	No raw sensitive content in routine telemetry.
Exports and audit records	Recipient, purpose, retention, accountability and legal-hold handling	Document download limits; audit retention separate from record retention.

REQ 21: approve a data inventory, processing-purpose register, retention schedule and deletion map spanning browser state, database, object storage, exports, logs, backups and processors. Archive and membership revocation are not erasure. Exact periods require a documented policy; none are invented here.

HHS describes the HIPAA Security Rule as applying to electronic protected health information maintained by covered entities and business associates. Applicability and required agreements must be assessed for the actual pilot relationships. This document makes no HIPAA compliance or business associate agreement claim. [R5]

A privacy lead and clinical safety lead must review new purposes and external data flows. Public aggregate data do not justify identifying individuals or joining independent datasets at the person level without a valid linkage basis.

27 AI governance and clinical safety controls

Governance is implemented through function-specific decisions, evidence requirements and accountable review. A general “AI platform” label is too broad to establish safety.

Decision	Required owner	Evidence gate
Enable a new AI function	Product, ML and clinical leads	Intended use, prohibited use, evaluation and human review design.
Change model or prompt	ML lead and release reviewer	Versioned prompt/model, regression corpus, error analysis and rollback.
Process a new data class	Privacy and security leads	Authority, processor terms, retention and threat assessment.
Add a clinical interpretation	Clinical safety lead	Source validity, independent review, foreseeable misuse and escalation limits.
Launch a real-data pilot	Accountable sponsor and relevant review bodies	Protocol, approvals, operating support and stop criteria.

Framework alignment

NIST AI RMF provides a structure for governing, mapping, measuring and managing AI risk; the generative AI profile adds relevant risk considerations. WholeContext should maintain a function-level risk register, model inventory and evidence log using that structure. This is a proposed operating approach, not a completed conformity claim. [R1]

Clinical boundary review

FDA’s Clinical Decision Support Software guidance was reissued in January 2026. Before expanding into diagnosis, treatment or patient-specific recommendations, qualified reviewers must assess each function’s intended use and regulatory implications. Source transparency alone does not settle that assessment. [R4]

REQ 22: define stop criteria for wrong-patient output, unauthorized disclosure, unsupported clinical assertions and material source corruption. Record incidents, affected versions, containment and corrective actions. No automated reassurance or safety clearance should be inferred from missing data.

28 Failure modes and recovery behavior

A failure should preserve the user's work, explain what is unavailable and prevent a partial result from appearing complete.

Failure	Current behavior	Required recovery evidence
Database unavailable	Care returns 503 and indicates that the form was not saved or cleared.	Browser confirms retained edits; retry does not duplicate records.
Revision conflict	409 prompts refresh before edit or sharing.	Three-way review or explicit re-entry; no silent overwrite.
Public context load fails	Load errors and unavailable layers remain in result.	Last accepted version identified; no zero substitution or guessed county.
Model unavailable or invalid	Explicit deterministic exact-source fallback.	User can distinguish modes; no records accepted automatically.
External directory fails	502 with error and notZero marker.	Distinguish unavailable results from no matching provider.
Patient changes during work	Selected views reject stale response or download.	Uniform cross-view tests for all asynchronous operations.
Audit write fails after update	Some mutations and audit writes are separate operations.	Transactional event design and reconciliation; no assumed rollback.
Support recipient revoked	Current membership controls access; inactive thread cannot receive new replies.	Regrant and download-copy policy tested explicitly.

REQ 23: normalize API error categories for invalid input, unauthorized access, conflict, dependency failure, rate limit and internal error. Include a non-sensitive correlation identifier and retry guidance. The current public review handler maps many reconciliation failures to a generic 400; operational diagnosis needs finer internal classification without disclosing sensitive details. [E1]

29 Observability capacity and service targets

The following targets are proposed pilot acceptance criteria. They are not measured production achievements or contractual service levels.

Measure	Proposed initial target	Measurement design
Non-AI care response	95th percentile at or below 1 second	Measure server latency separately from rendering under an approved pilot workload.
Synthetic review response	95th percentile at or below 2 seconds	Warm and cold contexts; maximum supported record selection; report payload sizes.
Packet generation	95th percentile at or below 5 seconds	Representative browser and largest supported packet; verify content and download.
Availability	99.5% monthly for pilot service	Define observation window and planned maintenance; validate dependency behavior.
Recovery	RPO at most 24 hours; RTO at most 8 hours	Proposed recovery point and recovery time; demonstrate backup restoration.
Access integrity	Zero known cross-patient or unauthorized disclosures	Release-blocking negative tests and incident review.
Accessibility	Target WCAG 2.2 AA for critical journeys	Manual keyboard and assistive-technology testing plus automated checks.

REQ 24: record request count, latency, status, version, source load failures, schema failures, quota reservations, fallback rates, conflict frequency and export failures. Exclude raw clinical text, source images and direct identifiers from routine telemetry. Use restricted incident evidence when detailed investigation is necessary.

AI cost monitoring must separate reserved attempts, actual provider calls, tokens and invoiced spend. The current trace reports dollarCost null. Model choice and live configuration were not established by source inspection. Load and recovery tests must be completed before service targets are advertised. [E1, R7]

30 Build deployment and rollback design

The release unit includes static assets, server modules, public snapshots, media and schema compatibility. Source commit alone is insufficient to reproduce a deployment.




Figure 6. Required release workflow. Existing preserved checkpoints support recovery; a fully automated pipeline is not claimed.

REQ 08 requires a release manifest containing source commit, dependency lock and SBOM, data hashes, model and prompt configuration identifiers, migration version, build tool versions, test reports and deployment identifier. Build outputs must be traceable to inputs without storing secrets in the artifact. [R2]

The v39 release preserves v38 as the immediate rollback checkpoint. Media packaging encountered the hosting service's 256 MiB expanded archive limit; the release report notes that decompressed tar overhead counts toward this limit. A package-size gate must inspect the complete expanded archive, not just summed regular-file sizes. [E4]

A database migration must be backward compatible with the rollback release or carry a tested restore plan. A successful code rollback does not undo a schema or data migration. Environment configuration, audience controls and secret versions require separate change records. No release operation is performed by this document.

31 Verification baseline and its limits

The saved release evidence documents substantial regression work. Each result below is bounded by the environment and version in which it was recorded; counts are not additive coverage scores.

Evidence	Saved result	Limit
v36 client-request repair	Eight groups and 210 client requests exercised the repaired public review path.	A repaired request contract is not complete browser acceptance.
v38 context release	43 of 43 groups passed, including all 65 patient-county round trips and selected-source PDFs.	Controlled DOM/event adapters and isolated Worker; no real-browser navigation pass.
v39 prepublication	17 of 17 media, source, component and regression groups passed; 18 scene frames inspected.	Media decode and component evidence do not prove browser sound, seeking or accessibility.
v39 hosted asset checks	23 changed public assets matched the saved commit exactly.	Byte identity validates delivery, not all user interactions.
v39 hosted groups	Four of five passed; byte-range test received 200 where 206 was expected.	Seeking behavior remains unconfirmed.
Anonymous boundary probes	Care API returned 401; sampled private paths returned 404.	Bounded checks; not a penetration test or proof for every path.

The v38 report records selected-context Twin PDFs with seven pages and verified patient identities, county values and original citations. It also preserves initial harness failures and corrected results. Retaining failed attempts and corrections improves reproducibility; only final accepted results are counted as passing. [E4]

Open acceptance items include full browser journeys, narrow-screen layouts, keyboard navigation, operating-system downloads, two-account isolation and revocation, hosted photo workflow and a traced synthetic model request. This specification records those gaps rather than converting them into completed tests.

32 Requirement to evidence mapping

This matrix defines the next acceptance campaign. “Partial” means relevant implementation or historical evidence exists; it does not mean the pilot gate has passed.

Requirement	Named verification task	Status and accountable role
01–03 Identity provenance missingness	T01 Mixed-patient rejection; source ablation; null, zero and excluded-value fixtures	Partial • application and data leads
04 10 Access and grant lifecycle	T02 Hosted owner/support/clinician/payer matrix; revoke and regrant	Partial • security and application leads
05 12 13 Mutations and limits	T03 Concurrent updates; sharing conflict; retry; byte and depth bounds; audit failure	Partial • backend lead
06 15 Optional AI and NLP	T04 Quota, timeout and schema failures; exact-span corpus; independent labels	Partial • ML and backend leads
07 09 Dated scenarios	T05 Date sweep, timezone fixtures, hypothesis reset and export isolation	Partial • application and clinical leads
08 Build and recovery	T06 Clean rebuild, manifest comparison, migration rollback and restore	Partial • release lead
11 14 Interface and ingest	T07 Frozen schemas; browser contract corpus; quarantined source refresh	Partial • backend and data leads
16 Image pipeline	T08 Decode quarantine, metadata, authorization, deletion and subgroup protocol	Planned • security and ML leads
17 18 Export and browser state	T09 Snapshot manifest; stale-result races; multi-tab patient changes	Partial • frontend lead
19 22 Model and safety promotion	T10 Function-level review, model card, incident and stop-condition exercise	Planned • clinical safety and ML leads
20 21 Security and privacy	T11 ASVS mapping, independent review, retention and erasure exercise	Planned • security and privacy leads
23 24 Operations	T12 Error taxonomy, load, telemetry minimization and SLO measurement	Planned • operations lead

Every test report shall record version, environment, input fixtures, expected behavior, observed result and retained failure evidence. Release approval requires a named reviewer and disposition for every unresolved critical item. No staffing or approval signature is implied by this role assignment.

33 Architecture decision register

The register makes existing design choices and proposed changes reviewable. Observed decisions are reconstructed from implementation; they are not a claim that historical ADR files already existed.

ADR	Decision and status	Reason and consequence
01	Shared canonical observation model • observed	One reconciliation path reduces view disagreement; changes require broad regression.
02	Server-hydrated synthetic review • observed	Bound known records and patches; limits free-form public input.
03	Public context separate from patient facts • observed	Avoid ecological inference; keeps more metadata in every layer.
04	Synthetic dated scenario with isolated assumptions • observed	Supports transparent exploration; cannot claim predictive physiology.
05	Exact-span optional AI categorization • observed	Protect source wording; limits fluent synthesis and requires human review.
06	Modular Worker plus static browser application • observed	Simple deployment; shared runtime and packaging constraints.
07	JSON payloads with relational ownership • observed	Flexible MVP; schema migration and query governance become more important.
08	Transactional audit and explicit grant objects • proposed	Improve accountability and revocation; requires migration and compatibility work.
09	Quarantined object store for attachments • proposed	Separate binary lifecycle from records; adds processing and deletion obligations.
10	Standards mapping before EHR connectivity • proposed	Validate meaning and consent first; defer integration claims until conformance tests pass.

For each new decision, record alternatives considered, consequences, evidence, owner, approval date and superseding decision. Public documentation can explain architecture without exposing secrets or unnecessary attack detail.

34 Target architecture for a governed pilot

The next architecture retains the shared evidence model while adding stronger identity, provenance, attachment and operating controls. This is a proposed target, not a depiction of deployed services.




Figure 7. Proposed pilot architecture. Explicit grants, quarantine, audit and observability are additions to be built and verified.

The API layer should enforce versioned schemas, current actor claims, patient scope, request limits and idempotency. Record mutations and audit events should commit together. Versioned source snapshots and their manifests should be immutable inputs to derived review, with approved refreshes promoted only after data-quality checks.

Attachments require a separate lifecycle: upload authorization, quarantine, safe decoding, accepted derivative, restricted retrieval, retention and deletion. AI processing should read only approved source classes through a gateway that enforces purpose, model version, budget and trace policy.

Interoperability is a later adapter boundary. FHIR resources such as Patient, Observation, DocumentReference, Provenance and Consent are candidate mappings; no FHIR endpoint or conformance claim exists in this baseline. Select the release and implementation guide with the integration partner, then test profiles, terminology and round-trip fidelity. [R6]

35 Development sequence and delivery gates

The plan is dependency-based. Calendar dates, staffing and cost estimates require an approved delivery team; no elapsed-time commitment is invented.

Stage	Work package and owner	Exit gate
A Preserve the baseline	Release lead: archive manifests, source, datasets, test results and rollback checkpoint.	Reproduce selected outputs; close browser acceptance and media seeking disposition.
B Stabilize contracts	Backend/frontend leads: OpenAPI schemas, state invalidation, patch semantics and all write limits.	T01, T03, T05 and T09 pass against actual browser requests.
C Strengthen governance	Security/privacy leads: explicit grants, transactional audit, host evidence, retention and erasure.	T02 and T11 pass; critical findings resolved; data authority documented.
D Harden data and AI	Data/ML leads: source refresh contracts, corpus evaluation, attachment quarantine and model inventory.	T04, T07, T08 and T10 evidence reviewed; no unsupported model claims.
E Operate a controlled pilot	Clinical and operations leads: protocol, support model, telemetry, load and restore drills.	T06 and T12 pass; approvals and stop criteria established before enrollment.
F Evaluate and scale	Research/product leads: prespecified comparison and adoption analysis; independent review.	Measured benefit and acceptable burden; approved expansion decision.

Parallel work is possible after Stage A, but access, retention and safety gates cannot be bypassed by completing interface features first. EHR connectors, wearables, additional NLP and CV models remain separate work packages with their own permissions and evaluation.

The full Transformation Roadmap should expand this sequence into work breakdown, staffing, dependencies, budget, risks and measurable milestones. This specification supplies the technical gate definitions for that next deliverable.

36 Adoption and success measurement

Adoption should be measured as successful, understandable use of a reliable evidence workflow. Usage counts alone cannot establish clinical benefit.

Outcome	Operational definition	Study or telemetry requirement
Onboarding completion	Users completing a reviewed intake and opening the intended summary divided by eligible starters	Report role and device; distinguish abandonment from technical failure.
Evidence completeness	Prespecified required history elements correctly represented in the review packet	Blinded or independent adjudication against source records; omissions and false additions.
Reviewer effort	Time to complete a defined reconciliation task and number of corrections	Comparable tasks and baseline workflow; account for learning and case complexity.
Comprehension	Patient and clinician understanding of known facts, gaps and next questions	Predefined rubric; accessible language; no inferred comprehension from page views.
Coordination reliability	Assigned steps with acknowledged status and documented completion	Separate saved, seen, acknowledged and completed states.
Safety and equity	Wrong-patient events, unsupported assertions, access failures and subgroup task gaps	Stop criteria; adverse-event review; meaningful subgroup sample sizes.
Longitudinal benefit	Prespecified function and care-process outcomes at defined follow-up times	Approved protocol, comparator, missingness plan and uncertainty intervals.

The pilot should define primary outcomes and sample-size rationale before enrollment. Use within-person or controlled comparisons only where the design supports them; report uncertainty and attrition. Synthetic demonstrations and ecological experiments cannot establish clinical effectiveness. [E5, E6]

Adoption support should include role-specific orientation, a short evidence-review task, visible source boundaries and a feedback route with an accountable responder. The narrated guides, captions and transcripts already support orientation; accessibility and comprehension still require user testing. [E4]

37 Risks and unresolved engineering work

The most consequential risks are concrete and testable. Their presence defines the next work package rather than diminishing the value of the completed prototype and research.

Risk	Consequence	Mitigation and gate
Host controls not independently evidenced	Application checks depend on infrastructure assertions.	Document identity, audience, encryption and backup controls; T02 and T11.
Flexible payload and audit semantics	Schema drift or incomplete mutation history.	Versioned schemas, migration tests and transactional audit; T03.
Historical or incomplete public sources	Context could be mistaken for current personal evidence.	Native period and coverage labels; refresh manifests; source ablation; T07.
Research and clinical claims conflated	Reviewers or users overestimate model capability.	Function inventory; model cards; distinct research results; T10.
Unfinished browser acceptance	Real interaction, media or accessibility failures persist.	Critical-journey browser tests across roles and devices; T09.
Small deployment package headroom	A routine asset addition can prevent release.	Package-size budget; separate media delivery plan; T06.
Unproven deletion and attachment lifecycle	Copies may outlive intended retention or permission.	Lifecycle map, quarantine and deletion exercise; T08 and T11.
Clinical review capacity and burden	A technically correct tool may fail in practice.	Prospective workflow study, staffed support and measured correction effort.

The v39 publication receipt reports approximately 1.1 MiB of expanded-package headroom. Capacity planning should address media and artifact growth before adding another large asset. No hosting change is authorized or executed by this specification. [E4]

A release decision must state accepted residual risks, owner, rationale and review date. Unresolved identity, unauthorized-access or material clinical-misrepresentation failures block real-data pilot promotion.

38 Project evidence register

Project references below identify the evidence used to write this volume. Source artifacts and private implementation files are not implied to be publicly hosted merely because they are cited.

Reference	Evidence and scope
E1 Server implementation	Pinned v39 dist/server modules: index, demo-review, care, circle-channel, extract, source-extraction, photo-validation, research, education, public-services and import-handoff. Contracts and controls described in Sections 3–28.
E2 Shared and browser logic	Pinned observation-model, packet-model, open-data-delta, digital-twin-core, context-discovery-core; browser photo, packet and care-intelligence modules. Read models, source matching and selected stale-result controls.
E3 Persistence definitions	db/schema.ts; drizzle/0000_care_workspace.sql, 0001_ai_daily_usage.sql and 0002_circle_channel.sql. Schema declarations and migration definitions; deployed enforcement requires verification.
E4 Release and verification	v36 repair report; v38 Open Data release report; v39 onboarding release report; WholeContext v39 Publication Receipt dated 14 October 2026. Historical test scope and deployment baseline.
E5 Data engineering matrix	WholeContext AI 171 Resource Data and Engineering Matrix v1.0, 14 October 2026; original September catalog and August atlas preserved. Counts, mappings, contracts, joins and method inventory.
E6 Research evidence	Technical Clinical and Impact Dossier v1.0; October ecological and synthetic experiments; September federal-data pipeline; real-data and synthetic methodology source documents. Research-to-product boundary.
E7 Strategy and ontology	Technical Strategy Architecture Addendum; Evidence Ontology and Orchestration Plan; Data Preparation Selection and Management. Historical design intent, not proof of implementation.
E8 This review	Pinned source inspection, document reconciliation, selected source digests, rendered-document checks. No new application deployment or production acceptance run.

Evidence priority is implemented source plus scoped verification, then reproducible research outputs, then dated plans. When a plan and release differ, the released implementation governs the current-state description.

39 External engineering and review references

Primary sources were checked on 14 October 2026. They inform proposed review gates; citing them does not establish compliance, endorsement or certification.

R1 AI risk management

NIST AI Risk Management Framework 1.0 and Generative AI Profile, NIST AI 600-1. Governance structure and function-level risk review. <https://www.nist.gov/itl/ai-risk-management-framework>

R2 Secure development

NIST SP 800-218, Secure Software Development Framework version 1.1, final February 2022. This document uses that final baseline; the separately published version 1.2 initial public draft is not treated as final. <https://csrc.nist.gov/pubs/sp/800/218/final>

R3 Application security verification

OWASP Application Security Verification Standard 5.0.0. Use version-specific requirement identifiers in the future control-to-test register. <https://owasp.org/www-project-application-security-verification-standard/>

R4 Clinical decision support review

FDA, Clinical Decision Support Software, final guidance reissued January 2026. Function-specific intended-use and regulatory review before expansion. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>

R5 Health information safeguards

HHS, The Security Rule. Applicability depends on the actual entities, data and relationships. <https://www.hhs.gov/hipaa/for-professionals/security/index.html>

R6 Interoperability

HL7 FHIR R4 overview, version 4.0.1. Candidate reference for future mappings; select partner-specific version and profiles before implementation. <https://www.hl7.org/fhir/R4/overview.html>

R7 Accessibility

W3C Web Content Accessibility Guidelines 2.2. Proposed AA target requires criterion-level assessment and manual testing. <https://www.w3.org/TR/WCAG22/>

40 Reproducibility identifiers and review record

These selected SHA-256 digests identify the exact inspected source bytes. A digest proves content identity when compared with a trusted copy; it does not prove correctness, authorship or security.

Pinned source commit

4a85dc7a60a989e2c731cb9ed2ed29ffdba7cbe1

Inspected source	SHA-256
server/index.js	c625305f62bf5bd2e610d55e423771da f868591495f872cb544ab6033a30f0ec
server/demo-review.js	71eb80a4afdeab63bae2a31b9da697d3 b7c4094489fdca6e0e7db3b93b62b976
server/care.js	b7e06573d1a5c6fad9ecf34399747a56 63d57d20a41c88df6dc4570e236eec2a
server/observation-model.js	5dcfe87876eeced907575f1d2b6056cc ba89a7764836e1c87dd40452cd645552
server/open-data-delta.js	aad2f49691e3075acdb9946d79fe6fca c51f4c5ebbc7682970a5b45ca5511511
server/digital-twin-core.js	f1c1379d8467505d071b6dc075dc8fe2 70dee6f4ded8e11e45bc7ec2eec55994
server/extract.js	4d5de9b1d22a2feb717929dd1f2a6b58 46bb8713ff7aaeb83fb146abf7a73418
db/schema.ts	3517b9e1470a22bd9a146a97fd986d88 8bf91ba8db5b048fc2ccf0e564e00624

The server/ prefix above abbreviates dist/server/ in the pinned repository. Source inspection also covered browser modules, migrations and saved release reports listed in Section 38. Later saved code was not used to redefine this public-release baseline.

Review record: architecture reconciled against pinned implementation; historical results attributed to their release; catalog totals checked against the companion matrix; proposed controls and targets labeled; PDF and editable document rendered and checked for page integrity. Application tests were not rerun for this document.

Approval record: this volume is prepared for expert review. Clinical, security, privacy and pilot approvals are not asserted. A release or pilot decision requires named reviewers, dated findings and disposition of the acceptance items in Sections 31–37.