

WholeContext AI

Transformation Roadmap and Engineering Delivery Plan

Document 4 in the technical review series

Kimberly Chulis PhD

Core Analytics

14 October 2026

Version 1.0

From a documented synthetic MVP to a reproducible, governed and measurable care collaboration pilot.

Review scope

Completed foundation • 18 work packages • 24 week reference plan • Six decision gates • Data and AI development • Privacy and security • Discovery Lab adoption • Capacity and evaluation

This volume turns the prior dossier, 171 resource matrix and architecture specification into concrete engineering work, accountable decisions and evidence required for the next release. Completed work, proposed delivery estimates and later research ambitions remain explicitly distinguished.

Prepared for the HHS TOPx AI and Invisible Illness challenge and independent engineering, clinical and research review. The roadmap does not claim clinical effectiveness, regulatory clearance or institutional endorsement.

Reader guide

Sections 1–7: baseline, dependencies and gates

Sections 8–16: detailed work packages

Sections 17–25: data, AI, governance, evaluation and adoption

Sections 26–30: ownership, resources, risks and release scope

Sections 31–35: traceability, sources and mobilization decisions

1 Document control and review contract

This roadmap converts the inspected MVP, data inventories and architecture requirements into a sequenced engineering program. It is a final planning artifact for expert review; the proposed program has not been funded, staffed, executed or approved for clinical deployment by this document.

Control	Declared baseline
Document	WCT ROAD 004 • Version 1.0 • 14 October 2026
Product lead	Kimberly Chulis PhD • Core Analytics
Software baseline	Public synthetic release v39; source commit 4a85dc7a60a989e2c731cb9ed2ed29ffdba7cbe1.
Evidence cutoff	The v39 publication receipt and the three companion volumes reviewed on 14 October 2026. Later saved source is outside this baseline.
Planning horizon	An illustrative 24 weeks from funded mobilization to a pilot readiness decision. Pilot observation and effectiveness studies occur afterward.
Status of estimates	Proposed person week ranges, role allocations, gates and budgets. They are neither recorded historical effort nor committed delivery dates.
Approval authority	Product, engineering, clinical safety, privacy and security roles must be explicitly assigned before gate approval. No signature is implied.

Evidence vocabulary

Established means supported by an identified artifact. Implemented means code exists in the inspected source; it does not automatically mean enabled or verified on the public site. Historical means evidence from an earlier checkpoint. Proposed means work or acceptance criteria introduced by this roadmap. Unverified means the necessary evidence has not been retained.

References E1–E9 identify project evidence; R1–R6 identify external frameworks used to structure review. The framework mappings are engineering choices, not certifications or institutional endorsements. All dates attached to future work are relative planning windows.

2 Executive delivery thesis

WholeContext AI has accumulated a substantial evidence and software foundation: longitudinal fictional journeys, a cataloged public data landscape, deterministic context comparison, shared care views, source grounded packets and bounded AI assistance. The next investment should make that foundation reproducible, secure, measurable and suitable for a governed pilot. [E1-E4]

The transformation to deliver

Move from a public synthetic demonstration to a consent governed collaboration workflow in which a person can review a source, approve selected evidence, share a bounded packet, receive a documented contribution and revisit function over time. Each transition must preserve identity, provenance, chronology, permissions and explicit uncertainty.

Outcome	Evidence needed for the next decision
Trustworthy evidence	Reviewed source to observation to chart to packet lineage; corrections and exclusions propagate reproducibly.
Reliable collaboration	Hosted multi account access, concurrent editing, revocation and audit tests pass for the intended roles.
Useful AI assistance	Source fidelity, abstention, failure behavior and review burden measured on a versioned evaluation corpus.
Responsible data expansion	A selected source has a permitted use, an admission contract, a manifest and a demonstrated decision contribution.
Repeatable adoption	Discovery Lab and care workflows reach meaningful value milestones with minimal telemetry and informed participation.
Pilot readiness	Independent findings, operational recovery and the prospective protocol receive a documented disposition.

What the 24 week reference plan buys

Eighteen work packages, twelve architecture test campaigns and six decision gates define the program. The reference capacity is 6.5 full time equivalent roles, producing 124.8 effective person weeks under an 80 percent planning availability assumption. The work package midpoint is 105 person weeks; the remaining 19.8 are capacity reserve, not spare scope.

The endpoint is a G3 pilot readiness decision. Clinical effectiveness, diagnostic accuracy, ROI and production scale remain separate hypotheses that require later evidence. A technically complete build can still receive a hold decision when permissions, safety, staffing or a partner workflow are unresolved.

3 What the completed work contributes

The work already produced is the starting capital of the program. Counts describe different units and cannot be added together as a single coverage or validation score. [E1–E4]

Asset established in the review series	Contribution and remaining conversion work
171 resource catalog records	Discovery breadth across datasets, portals, standards and methods. Select implementable sources; do not describe all 171 as connected datasets.
107 Health Data Atlas source records	Prioritization, proposed joins and use cases. The retrieved workbook README identifies version 1.3, dated 29 August; earlier prose calling it version 4 is not the authoritative file version.
120 runtime registry entries	12 stored snapshots, 90 catalog only and 18 unavailable or restricted; zero live OpenDataDelta adapters in this scope. Other public services are a separate category.
198 reconciled lineage mappings	Traceability between inventories in Document 2. Eleven stored snapshots map to eleven catalog entries; annual EPA county AQI is outside the exact 171 record catalog.
19 data engineering contracts	Candidate implementation boundaries, retained with 54 derived definitions, 16 proposed outputs and 21 join rules. None of these design counts proves deployment.
65 fictional identities in v39	Repeatable patient and role journeys. Synthetic demonstrations cannot estimate prevalence, clinical benefit or performance on real patients.
Research and AI methods	Separate ecological and synthetic modeling assets; bounded optional NLP implementation; image documentation and simple quality heuristics. These are different maturity levels.
Release and architecture evidence	Publication receipts, source inspection, 24 requirements and 12 named test campaigns. Close the remaining hosted browser and operational gaps.

Existing source engineering, scenario generation, feature definitions, clinical reasoning contracts, ontology work and repeated release corrections remain reusable. The plan preserves their source IDs, authorship and limitations rather than recasting planned work as already complete.

The September adoption telemetry receipt describes milestone tracking and collaborative utility, while explicitly recording that a full browser build was not run at that checkpoint. It is historical evidence for the measurement design, not proof of current v39 telemetry operation. [E7]

4 Transformation model and dependency structure

The program advances on evidence gates. Workstreams can overlap after their interface assumptions are reviewed; acceptance remains blocked until the required upstream evidence exists.




Figure 1. Proposed dependency structure. G2 joins foundation work; G3 requires integration, independent review and operational readiness. Later study gates are outside the 24 week delivery envelope.

Likely controlling dependencies

Identity and grant semantics, canonical revision handling and auditable mutations are likely to control patient workflow integration. Source rights and clinical partner decisions may control the schedule independently of software throughput. Until work package estimates are decomposed and role leveled, this is a dependency model rather than a computed critical path.

Parallel work with explicit boundaries

Data contracts, the pilot protocol and threat modeling begin early. NLP and image engineering can proceed with fictional or otherwise authorized data after source and interface contracts exist. They cannot use an unapproved patient corpus or promote a model into care because another workstream is ready.

Research continues in a segregated workspace. Only a reviewed promotion record may move a research artifact into a user facing decision workflow. Discovery Lab can expose public source exploration and synthetic education before patient pilot activation.

5 Indicative development sequence

Week 1 begins when the accountable owners, intended use and resource envelope are approved. The bands below are planning windows, not promises or a claim that the team is already assigned.




Figure 2. Proposed 24 week schedule. Numbers are work package IDs. G2 accepts the foundation subset; data and security work continue to final G3 review. Overlapping bars do not imply unlimited specialist capacity.

Sequencing rule

A work package may begin on a reviewed draft contract. It may not close against a moving contract. WP03 and WP04 freeze the pilot service and evidence interfaces before WP05, WP06 and WP07 acceptance; subsequent changes require migration and regression evidence.

Replanning triggers

Replan when a critical prerequisite slips more than one iteration, estimate to complete exceeds available capacity, an independent finding changes the intended use, or the pilot partner cannot supply the agreed workflow. Preserve safety and privacy gates; reduce optional sources, integrations or AI scope before compressing review.

Use two week iterations with a weekly dependency review. At weeks 2, 4, 12 and 24, publish a gate decision and evidence index. A calendar milestone without its evidence package is a hold, not a pass.

6 Gate register and decision authority

Every gate records the release, dataset and configuration under review; required evidence; unresolved issues; decision; approvers; and the next review date. Gate status in this document is proposed and not yet passed.

Gate and target	Minimum exit evidence	Accountable review
G0 Baseline Week 2	Pinned source and data inventory; claims register; owners; scope and estimate assumptions; evidence gaps accepted into the backlog.	Product and engineering; QA evidence review
G1 Public review Week 4	Critical synthetic browser journeys, packet identity and download checks; media defect disposition; accessible review path; publication manifest.	Product and QA; clinical wording review
G2 Foundation Week 12	T01-T03 foundation checks; frozen interfaces; source manifest skeleton; hosted grant lifecycle; retention design; threat model. Remaining T07 and T11 work is explicit.	Engineering, data, privacy and security
G3 Pilot readiness Week 24 target	T01-T12 complete for enabled pilot scope; recovery exercise; independent security review; approved protocol and data authority; support coverage; stop rules.	Product, engineering, clinical safety, privacy and security
G4 Feasibility After protocol window	Locked feasibility analysis; safety and withdrawal review; task completion, burden, fidelity and missingness; decision on iteration or comparative study.	Study lead with independent methods and safety review
G5 Expansion After suitable validation	Evidence appropriate to each intended claim; external or prospective validation where needed; operating capacity, economics and monitoring; renewed function review.	Accountable release board and applicable partner authorities

Hold and exception discipline

A gate can pass, pass with bounded noncritical follow up, hold, or stop. An exception requires an owner, expiry, compensating control, affected claim and retest. Cross patient disclosure, unauthorized sharing, untraceable material evidence or an unmet authority to process data cannot be waived into a patient pilot.

G1 permits a reviewed synthetic demonstration. G2 permits continued engineering and controlled testing. Neither authorizes real patient enrollment. G3 applies only to the enabled functions, specified participants, organization, environment and data uses in the approved protocol.

7 Work package estimates and ownership

Effort is total productive person weeks across contributing roles. It includes implementation, review and work package tests; WP17 covers cross system integration and independent campaign coordination without recounting the same unit work.

WP	Deliverable stream	Weeks	Effort	Primary roles
01	Baseline and evidence	1–2	2–3	Delivery / QA
02	Public demo acceptance	1–4	3–5	Frontend / QA
03	Service contracts	3–8	5–7	Backend
04	Canonical evidence model	3–10	5–8	Backend / Data
05	Identity and grants	5–12	6–9	Backend / Security
06	Reliable mutations	5–12	6–9	Backend
07	Data admission and lineage	3–16	7–10	Data
08	Dated context and scenarios	9–16	3–5	Frontend / Data
09	NLP evaluation and control	9–18	5–8	ML / Clinical
10	Image intake and review	9–18	4–6	Backend / ML
11	Reproducible care packets	11–18	4–6	Frontend / Backend
12	Role journeys and access	9–20	4–6	Frontend / UX
13	Operations and recovery	5–22	6–9	Platform / Operations
14	Privacy and security review	3–22	6–9	Privacy / Security
15	Discovery Lab adoption	13–22	4–6	Product / UX
16	Pilot protocol and co design	3–22	5–8	Clinical / Research
17	Integrated verification	17–24	6–9	QA / Independent review
18	Readiness and handoff	23–24	2–4	Delivery / Clinical

Total range: 83–127 productive person weeks; arithmetic midpoint: 105. The range is a planning interval, not a statistical confidence interval. Specialist availability, procurement and partner approvals are not represented by simply summing effort.

At G0, break each package into acceptance sized tasks, allocate work by role and week, identify reviewer independence and refine remaining effort. No named individual beyond the product lead is represented as hired or committed. The role mix in Section 27 is a proposed capacity pool.

8 Baseline preservation and public acceptance

First protect the existing accomplishment and make the reviewer experience reproducible. A release receipt and a successful local check are inputs to this work, not substitutes for it.

WP01 Baseline and evidence

Retain a recoverable baseline and an honest public claims register.

Delivery. Create the source, data and dependency manifest; archive publication and rollback receipts; reconcile document versions; record which functions are public, authenticated, configuration dependent, research only or planned. Link every material submission statement to evidence and a limitation.

Dependencies and ownership. Delivery owns the manifest; QA independently checks a sample of claims against source and publication evidence. Product resolves inconsistencies. Start from v39 and retain its rollback reference.

Exit evidence. An independent reviewer can identify the source and datasets for each enabled journey. Missing assets and unresolved findings have owners. The 171 catalog, 107 atlas and 120 runtime definitions remain distinct; proposed estimates are labeled.

Planning window 1–2 • Effort 2–3 person weeks • Gate G0 • Architecture tests T06.

WP02 Public demo acceptance

Prove the critical public synthetic journeys in actual browsers.

Delivery. Exercise entry, patient switching, dated context, source selection, clinician packet generation and download, navigation and mobile or keyboard use. Reproduce the hosted media range response issue and either fix it or publish a tested alternate viewing path.

Dependencies and ownership. Frontend and QA use WP01 fixtures and the exact release candidate. Clinical review checks that synthetic examples, context and hypotheses are clearly distinguished in visible copy.

Exit evidence. Retain browser and device versions, timestamps, fixture IDs, expected and observed results, screenshots where appropriate and downloadable packet hashes. Resolve identity leakage or misleading evidence. Historical 43/43 and 17/17 reports do not replace these hosted checks.

Planning window 1–4 • Effort 3–5 person weeks • Gate G1 • Architecture tests T05 T09.

9 Contracts and the canonical evidence model

A stable contract between the browser, services and derived outputs is the main protection against silent divergence as the platform grows.

WP03 Service contracts

Freeze versioned service interfaces and error semantics.

Delivery. Define request and response schemas, validation bounds, status and error codes, correlation IDs, pagination and compatibility policy. Capture real browser payloads as contract fixtures. Define protected route identity requirements independently of the public shell.

Dependencies and ownership. Backend owns schemas; frontend and security review them. WP01 supplies routes and implementation status. Drafts unblock WP04–WP07; accepted changes later require versioning and consumer tests.

Exit evidence. Malformed payloads, unknown fields, missing authority and stale revisions produce specified failures. Browser fixtures validate against the same contract. Error bodies contain no patient narrative, credentials or internal stack detail.

Planning window 3–8 • Effort 5–7 person weeks • Gate G2 • Architecture tests T07 T12.

WP04 Canonical evidence model

Preserve patient identity, source meaning and review state in one evidence model.

Delivery. Define stable observation IDs; patient, source and revision references; event time versus recorded time; unit and timezone; explicit missingness; author and reviewer; and correction lineage. Distinguish an asserted observation from a proposed extraction, public context and a hypothetical scenario.

Dependencies and ownership. Data and backend leads review WP03 contracts with a clinician and frontend lead. Map existing JSON records and the 22 readable metrics; reconcile the narrower wellness write contract rather than assuming every readable metric is writable.

Exit evidence. Mixed patient inputs fail; zero, absent, invalid and excluded states remain distinct; unit and temporal edge cases are tested. A correction creates a traceable revision and changes affected displays and packets deterministically.

Planning window 3–10 • Effort 5–8 person weeks • Gate G2 • Architecture tests T01 T05.

10 Identity grants and reliable mutations

These work packages turn collaboration into enforceable service behavior. A hidden page or a browser redirect is not an authorization control.

WP05 Identity and grants

Enforce current identity, membership, role, recipient and selected record scope.

Delivery. Document the trusted hosting identity boundary; implement or verify server authorization on every protected read and write. Specify invitation, grant, expiry, revocation and regrant semantics. Test ownership changes and whether an old grant may ever regain effect.

Dependencies and ownership. Backend owns enforcement; security reviews the threat model and hosted allow and deny matrix. WP03–WP04 provide stable record and principal contracts. Product and clinical owners decide intended sharing semantics before implementation.

Exit evidence. Distinct hosted accounts demonstrate allowed and denied access across patient, supporter, clinician and any approved payer role. Revocation applies to subsequent authorized retrieval and cached state handling; previously downloaded files are explicitly outside remote recall. T02 retains negative tests.

Planning window 5–12 • Effort 6–9 person weeks • Gate G2 • Architecture tests T02.

WP06 Reliable mutations

Make edits, sharing and audit changes reliable under races and retries.

Delivery. Require expected revisions; define conflict recovery; use atomic state and audit transactions where supported; implement bounded idempotency for selected mutation routes. Enforce streaming byte and JSON depth limits before expensive parsing. Preserve sharing flags during unrelated edits.

Dependencies and ownership. Backend owns mutation semantics. WP03–WP05 define schema, identity and grant rules. Data and QA review migration fixtures. Document any storage limit that prevents atomicity and redesign the operation or constrain the pilot scope.

Exit evidence. Two editors cannot silently overwrite each other. Duplicate retries do not create duplicate actions or grants. Failed audit persistence cannot leave an untracked privileged state change. Fault injection and malformed body tests retain both expected and observed states.

Planning window 5–12 • Effort 6–9 person weeks • Gate G2 • Architecture tests T03.

11 Data engineering and dated context

Source expansion is selective. Admission depends on provenance, permitted use and relevance to a user decision, not the size of the source catalog.

WP07 Data admission and lineage

Convert selected catalog resources into reproducible data products.

Delivery. Reconcile the 19 engineering contracts with runtime source IDs. For admitted releases capture retrieval, license, hash, schema, units, observation period, geography vintage, row counts and transform version. Add quality checks, quarantine, lineage and controlled refresh with rollback.

Dependencies and ownership. Data lead works from Document 2, WP03–WP04 and the rights review in WP14. At G2, accept the manifest and quarantine foundation for the selected core. Source expansion and output acceptance continue to G3.

Exit evidence. A changed or malformed upstream release cannot silently replace accepted context. Every material output resolves to an admitted release and transform. Required fields, duplicate keys, coverage and joins satisfy source specific thresholds approved before ingestion; failures are visible rather than imputed as zero.

Planning window 3–16 • Effort 7–10 person weeks • Gate G2 / G3 • Architecture tests T01 T07.

WP08 Dated context and scenarios

Keep patient history, external context and assumptions temporally separate.

Delivery. Implement the dated eligibility contract across charts, OpenDataDelta and the synthetic twin. Preserve source vintage and whether information was available at the selected encounter. Test geography mismatch, incompatible denominators, no useful change and removal of a previously selected source.

Dependencies and ownership. Frontend and data leads use WP04 and admitted WP07 snapshots. Clinical review approves explanatory language. Scenario hypotheses remain authored assumptions; no learned physiological or causal model is implied.

Exit evidence. Date sweeps, timezone boundaries, leap days and reset tests pass. Later patient facts are excluded at earlier dates. Hypotheses never enter clinical facts or exports as observations. Source removal updates deltas and lineage; unsupported geography yields explicit missing context.

Planning window 9–16 • Effort 3–5 person weeks • Gate G3 • Architecture tests T05.

12 NLP methodology and image intake

Two distinct programs are needed: evaluate bounded text assistance and make image handling trustworthy. Neither program authorizes diagnostic automation.

WP09 NLP evaluation and control

Measure extraction fidelity and review burden before broader AI promotion.

Delivery. Define annotation guidelines for exact spans, temporal cues, negation, uncertainty and attribution. Build an authorized, versioned corpus; independently label a subset and adjudicate disagreements. Compare the deterministic baseline with the bounded model assisted extractor; retain prompt, model, configuration and error traces.

Dependencies and ownership. ML owns evaluation; clinical reviewers own labeling guidance; privacy approves corpus use. WP03–WP04 supply source contracts. Corpus splits prevent the same source or synthetic template family leaking across development and test sets.

Exit evidence. Report span precision and recall, unsupported extractions, negation or date errors, abstention, review edits and latency by meaningful slices. Preserve UTF16 span semantics. Timeout, quota, malformed output and provider failure disclose fallback and leave source records unchanged. Thresholds are approved before the locked test.

Planning window 9–18 • Effort 5–8 person weeks • Gate G3 • Architecture tests T04 T10.

WP10 Image intake and review

Admit images through a controlled source documentation pipeline.

Delivery. Add server decode validation, type and size limits, pixel limits, quarantine, metadata minimization, isolated object storage and scoped retrieval. Preserve source, date, body location when authorized, review state and derivation history. Define deletion across originals, reencoded images, thumbnails and backups.

Dependencies and ownership. Backend, ML and security coordinate WP05, WP07 and WP14. The current brightness or contrast heuristic is documented as a basic quality check. A future classifier requires its own labeled data, representativeness and clinical review program.

Exit evidence. Malformed and oversized images, unauthorized fetches and deleted object requests behave as specified. Patient identity and image versions remain linked. Quality feedback does not imply diagnosis; no image classifier is promoted by this work package.

Planning window 9–18 • Effort 4–6 person weeks • Gate G3 • Architecture tests T08.

13 Care packets and role specific journeys

The output of the platform is a usable shared evidence workflow. Fidelity and comprehension must survive the last step into a packet, mobile screen or another person's view.

WP11 Reproducible care packets

Make every exported packet an identifiable evidence snapshot.

Delivery. Generate brief and extended packets from the same reviewed snapshot. Include patient identity, included record revisions, source hashes where available, creation time, selected public layers, material omissions and review status. Record generation failures and avoid mixing asynchronous results from different active patients.

Dependencies and ownership. Frontend and backend own export generation and state invalidation. WP04, WP05 and WP08 establish identity, scope and dated context. The clinical reviewer defines critical facts that must survive a brief summary.

Exit evidence. A manifest reproduces the packet content. Browser patient switches, multi tab changes and delayed responses cannot contaminate the active case. Excluded records stay excluded. Brief versus extended retention is measured on approved fixtures; no blanket claim of complete summary fidelity is made.

Planning window 11–18 • Effort 4–6 person weeks • Gate G3 • Architecture tests T09.

WP12 Role journeys and access

Complete patient, clinician, supporter and financial workflows with accessibility.

Delivery. Run task based walkthroughs covering source review, selected sharing, supportive contributions, clinician reconciliation, function tracking and follow up. For financial views, preserve estimate, bill, payment and adjudication distinctions. Create role specific onboarding and plain language limits.

Dependencies and ownership. Frontend and UX coordinate clinical and patient advisers. WP05 and WP11 provide enforceable access and packets. Payer or financial participation is enabled only for a justified workflow and authorized scope. WCAG 2.2 is the proposed accessibility target. [R6]

Exit evidence. Keyboard, focus, screen reader labels, zoom, contrast and error recovery receive manual and automated checks on critical paths. Users can explain what is shared, what remains uncertain and how to correct or stop sharing. Record task failures and fatigue burden, not only satisfaction.

Planning window 9–20 • Effort 4–6 person weeks • Gate G3 • Architecture tests T02 T09.

14 Operations privacy and security

The pilot environment needs a recoverable service and an accountable data lifecycle. Existing hosting and a provider configuration flag do not establish either condition.

WP13 Operations and recovery

Build observable, reproducible and reversible releases.

Delivery. Pin dependencies, preserve a software inventory, separate secrets and configuration, test migrations and restore, define release manifests and rollback. Measure response times, resource use, error classes, AI fallback and storage growth. Establish on call coverage and incident runbooks appropriate to pilot hours.

Dependencies and ownership. Platform and operations own execution; QA witnesses recovery. WP03–WP06 stabilize contracts and storage. Resolve the v39 package size constraint before adding large assets; evaluate external object delivery or packaging changes through an explicit architecture decision.

Exit evidence. A clean build and restore reach the intended functional state with reconciled records and grants. Recovery time and data loss targets are agreed and measured. A rollback does not reintroduce revoked access or incompatible data. Minimal logs support diagnosis without patient narrative leakage.

Planning window 5–22 • Effort 6–9 person weeks • Gate G3 • Architecture tests T06 T12.

WP14 Privacy and security review

Demonstrate the intended data authority and security controls.

Delivery. Maintain a data flow and retention map; identify organizational responsibilities, allowed purposes and provider terms; assess applicable obligations with qualified reviewers. Define access review, key and secret handling, incident notification, vulnerability triage and erasure. Map the selected ASVS controls to tests. [R2–R5]

Dependencies and ownership. Privacy and security own their reviews independently of feature authors. At G2, approve the architecture and lifecycle design; at G3, require exercised controls and independent findings. A contract requirement or partner approval is a delivery dependency, not a footnote.

Exit evidence. Retention and erasure tests include backups, logs, derived exports and image variants. Threat scenarios cover cross account access, injection, replay, exfiltration and malicious documents. Critical unresolved findings block pilot activation; a control map alone is not a compliance certification.

Planning window 3–22 • Effort 6–9 person weeks • Gate G2 / G3 • Architecture tests T11.

15 Discovery Lab and pilot preparation

Adoption and evaluation are designed alongside engineering. Discovery Lab offers a useful public entry point while the care pilot remains governed by its own permissions and protocol.

WP15 Discovery Lab adoption

Turn WholeContext Discovery Lab™ into an evidence based adoption path.

Delivery. Define public source exploration, an intelligible source limitation view, synthetic journey entry and a clear next action. Test whether a user reaches a meaningful result and can explain its limits. Build opt in routes to pilot interest and partner conversations without exposing care records in marketing analytics.

Dependencies and ownership. Product and UX own the funnel; data reviews source meaning and privacy approves telemetry. WP07 supplies source provenance, WP12 supplies comprehensible journeys and WP13 supplies minimal events. The historical adoption receipt informs event design but is reverified in the target release.

Exit evidence. A visitor can find a relevant source, inspect its period and geography, understand what it cannot say about an individual and enter a fictional journey. Consent to research or follow up is separate from browsing. Report task success, opt in conversion and usefulness; do not infer clinical benefit from traffic.

Planning window 13–22 • Effort 4–6 person weeks • Gate G3 • Architecture tests T12.

WP16 Pilot protocol and co design

Prepare a feasible prospective workflow study with the right oversight.

Delivery. Co design the pilot with patient, clinician and workflow owners. Specify intended use, eligible participants, data sources, consent and withdrawal, comparator tasks, measures, adjudication, missingness and stop rules. Determine whether ethics or institutional review is required and document that determination.

Dependencies and ownership. Clinical and research leads own protocol design with privacy and a qualified methods reviewer. WP01 defines claims; WP09 and WP12 inform burden and safety. Partner recruitment, approvals and study execution are not assumed to be completed simply because software is ready.

Exit evidence. Before G3, an approved protocol names primary feasibility outcomes, sample size rationale, observation period, analysis rules, responsible staff and escalation coverage. Training and consent materials are tested for comprehension. No real participant enrollment begins under an unsigned protocol.

Planning window 3–22 • Effort 5–8 person weeks • Gate G3 • Architecture tests T10.

16 Integrated verification and readiness handoff

The closing work creates a reviewable decision package. It does not turn unresolved risks into assumptions merely to reach a calendar date.

WP17 Integrated verification

Execute the architecture acceptance campaign against the hosted candidate.

Delivery. Run T01–T12 for enabled scope, including browser journeys, mixed identity rejection, grant lifecycle, mutation races, ingest quarantine, NLP fallback, image controls, packet reproduction, restore, model promotion controls, privacy and operations. Retain failure evidence and retest changed dependencies.

Dependencies and ownership. QA coordinates; work package owners remediate; independent specialists assess security and methods. WP03–WP16 supply stable candidate components. Test authors identify fixture provenance and do not substitute in memory simulations for hosted behaviors that depend on identity or infrastructure.

Exit evidence. Each requirement has an expected result, observed result, environment, version, fixture, reviewer and disposition. Disabled functions are verified as unreachable and are recorded as deferred rather than passed. All mandatory enabled scope criteria pass; critical failures remain release blockers.

Planning window 17–24 • Effort 6–9 person weeks • Gate G3 • Architecture tests T01–T12.

WP18 Readiness and handoff

Deliver a signed, bounded pilot readiness decision and operating handoff.

Delivery. Assemble release notes, software and dataset manifests, claims register, protocol, risk disposition, training, support model, recovery evidence, monitoring plan and next decision date. Run a rehearsal of onboarding, correction, withdrawal, incident response and rollback with fictional records.

Dependencies and ownership. Delivery coordinates product, engineering, clinical safety, privacy and security approvers. WP17 supplies acceptance evidence; WP16 supplies authority and protocol. Budget and staffing for actual pilot operation are separately approved before activation.

Exit evidence. The G3 record identifies exact permitted functions, environment, participant scope and stop conditions. Operators can locate the runbooks and escalation owners. A reviewer can reproduce a representative source to packet path. A hold decision includes a specific closure plan and does not authorize enrollment.

Planning window 23–24 • Effort 2–4 person weeks • Gate G3 • Architecture tests T06 T10 T11.

17 Source admission and data expansion policy

The 171 resource matrix is a portfolio of possible evidence. An integration is justified by a specific user decision and a complete engineering contract, not by an aggregate integration count. [E2, E5]

Admission stage	Required artifact	Rejection or defer condition
1 Purpose and fit	Named user question, output, unit of observation and incremental contribution.	No defensible decision contribution or ecological evidence presented as patient fact.
2 Rights and access	Release specific terms, allowed uses, steward, retrieval route and restrictions.	Unresolved rights, controlled access approval or redistribution authority.
3 Scientific compatibility	Observation period, geography vintage, denominator, coding scheme and confounders.	Incompatible population or time; unsupported linkage; no valid missingness policy.
4 Engineering admission	Manifest, schema, keys, units, hash, row counts, transform version and quarantine tests.	Unbounded failures, ambiguous keys or unreproducible transforms.
5 Output verification	Source ablation, no change case, changed output explanation and lineage.	Source cannot be traced or alters unrelated patient facts.
6 Operational ownership	Refresh policy, owner, stale state, rollback and cost envelope.	No steward or unsustainable maintenance burden.

Prioritization after hard gates

Among admissible candidates, score decision relevance and gap reduction at 30 percent, scientific fit at 25, access reliability at 15, engineering feasibility at 15 and reuse at 15. Use an ordinal 0–4 rubric with written reasons and two reviewers. These proposed weights rank planning candidates; they are not scientific evidence weights or model probabilities.

Wave A stabilizes the existing 12 snapshots and their lineage. Wave B selects a small permitted subset from the 19 contracts after scoring. Wave C evaluates partner or controlled sources only after access and use approval. The number in each expansion wave remains a scope decision; all 171 records remain cataloged whether selected or deferred.

Treat terminology standards, literature references, discovery portals, patient contributed sources and public statistical datasets as separate resource types. A standard is not a patient dataset; a portal is not evidence that every underlying release has been acquired.

18 AI and ML development sequence

WholeContext uses deterministic processing, optional language model assistance, authored reasoning contracts and separate research experiments. Method selection follows the task and available evidence rather than a requirement to insert a model into every step. [E1, E3, E6]

Method family	Development and evaluation sequence	Promotion boundary
Deterministic evidence and context	Freeze rules; construct counterexamples; test provenance, missingness, date and source ablation; review explanations.	Changes require regression evidence and clinical wording review. No individualized risk claim.
NLP assistance	Define labels; source grouped split; double annotation and adjudication; baseline comparison; locked test; failure and review burden analysis.	Promote only the evaluated extraction task, model and configuration. Human source review remains required.
Image and computer vision	First secure capture and review; then define a separate task, authorized dataset, label protocol and device or lighting slices.	Quality heuristics are not a classifier. Clinical image interpretation requires its own evaluation and function review.
Ecological research ML	Retain geographic unit, state grouped or temporal splits, preprocessing within training folds, baselines and leakage tests.	County associations cannot become personal diagnosis or exposure probability by relabeling the output.
Synthetic longitudinal research	Retain generator, seed, scenario family, dependencies and split boundaries; assess technical invariants and abstention.	Synthetic test performance demonstrates behavior on authored cases, not clinical validity.
Future clinical prediction	Specify intended use and target; authorize representative labels; predefine statistical plan; assess calibration, subgroups and external shift.	Separate G5 decision after appropriate prospective or external evidence.

Reproducibility and reporting

Each experiment retains dataset and split IDs, inclusion rules, training only preprocessing, feature definitions, model or prompt version, hyperparameter selection, seeds, compute, metrics and limitations. Report uncertainty where the sampling design supports it. Do not report confidence intervals that treat correlated rows as independent people.

Use the NIST AI RMF as a review structure for ownership, intended use, measurement and response. The project chooses to require a promotion record with permitted use, adverse cases, monitoring and rollback. This proposed governance process does not imply certification by an external framework. [R1]

19 Data governance and privacy work products

The governance strategy must be executable at the level of a source, record, purpose and recipient. Public data, synthetic cases, care records and research copies require different lifecycle decisions.

Work product	Required contents	Acceptance owner
Data inventory and flow map	Systems, subprocessors, categories, purpose, identifiers, movement, storage, jurisdiction and authorized recipients.	Privacy with engineering
Authority and consent register	Source of authority, notices, sharing scope, research permissions, withdrawal and contract dependencies.	Privacy with clinical and partner owner
Retention and erasure schedule	Duration and basis for each store; export and backup behavior; legal hold handling; verification and deletion evidence.	Privacy with operations
Role and access review	Named roles, least privilege, grant lifecycle, privileged access, review cadence and emergency procedure if applicable.	Security with workflow owner
Threat and abuse model	Cross account access, source injection, malicious uploads, replay, exfiltration, model misuse and unauthorized secondary use.	Security with ML and engineering
Incident and concern process	Detection, containment, evidence preservation, reporting responsibilities, participant communication and recovery.	Security and clinical safety
Vendor and function assessment	Data terms, retention, region, security evidence, service changes and intended function evaluation.	Privacy and clinical safety

Before processing real care data, qualified reviewers determine the applicable organizational and legal obligations. HHS Security Rule material informs that assessment when relevant; the presence of encryption or a model request flag is not a conclusion of HIPAA compliance. FDA clinical decision support guidance informs a function specific review; this roadmap does not declare every function exempt or cleared. [R4, R5]

Operational safeguards

Use data minimization, restricted logs, separated research workspaces, reviewable exports and explicit source permissions. Audit important actions without copying clinical narrative into analytics. Confirm deletion and revocation semantics with users, including the practical limits once an authorized recipient has downloaded a file.

At G3, retain the signed decisions and exercised controls, not just policy documents. Residual risks receive owners, review dates and stated boundaries. A change in data use, model function, recipient population or hosting requires renewed review.

20 Clinical and workflow evaluation program

Evaluate technical fidelity, workflow feasibility and patient outcomes as different questions. The first governed study should determine whether the intended workflow can be used safely and meaningfully with an acceptable burden.

Feasibility protocol before enrollment

Specify setting, intended users, eligibility, data sources, recruitment and consent, observation period, comparator tasks, primary measures, missingness handling and withdrawal. A qualified methods lead sets the sample size using the desired precision or decision threshold and expected attrition. No participant count or statistical power is invented in this plan.

Study question	Proposed measurement	Interpretation limit
Can evidence be reviewed faithfully?	Blinded assessment of source to packet retention, contradictions, omissions and correction effort.	Does not establish diagnostic accuracy.
Can people complete the workflow?	Observed completion of review, selected sharing, contribution and follow up; assistance required and reasons for failure.	Volunteer success may not generalize to broader users.
Does it reduce coordination burden?	Within task or suitable comparative timing, rework, clinician edits and patient reported burden.	Learning, selection and concurrent process changes require analysis.
Can longitudinal function be tracked?	Completion, missingness and interpretability of prespecified function measures and treatment timelines.	An observed change is not a causal treatment or app effect.
Is the workflow acceptable and safe?	Misinterpretation, inappropriate reliance, access incidents, distress, withdrawal and usability concerns.	Absence of observed harm in a small study is not proof of safety.

Comparison and analysis

Use the same task definitions across conditions and retain baseline workflow measurements. If feasible, counterbalance task order or use an appropriate comparison design. Report distributions and uncertainty, missing denominators, edits and adverse cases. Prespecify primary outcomes and distinguish exploratory subgroup analyses.

For subsequent effectiveness claims, design a study that supports the proposed causal interpretation, with appropriate comparator, duration and analysis. G4 may recommend refinement, stopping or a larger study. It does not automatically authorize broader clinical claims or automated decision functions.

21 Release engineering and operational readiness

A deployment is ready only when the team can identify what is running, recognize failure and recover without corrupting evidence or access. The following are proposed pilot operating requirements.

Control	Evidence retained at G3
Build and release identity	Source commit; dependency lock and inventory; configuration version; migration IDs; public asset and dataset manifests.
Environment separation	Distinct development, test and pilot configuration; fictional fixtures; controlled secrets; documented identity trust boundary.
Deployment rehearsal	Clean build; migration forward and recovery path; smoke and critical journey results; active version and rollback receipt.
Recovery exercise	Restore record counts and revisions; reconcile grants and audit records; measure actual recovery time and data loss against approved targets.
Observability	Error taxonomy, correlation IDs, endpoint latency, failed mutations, ingest freshness, AI fallback, resource and spend indicators.
Incident response	Named operational and clinical owners; severity criteria; contact coverage; containment and communication runbooks.
Capacity and asset budget	Load profile and concurrent users; response distributions; package size and storage budgets; explicit media and image delivery plan.

Service objectives are measured assumptions


At G0, define the pilot hours, expected concurrency and critical user tasks. At G2, set candidate latency, availability, restore time and recovery point objectives based on measurement and partner needs. At G3, retain load and restore evidence. Do not advertise a production availability percentage without a defined measurement window and actual observations.

The inspected v39 package is 267,274,240 bytes against a 268,435,456 byte limit, leaving roughly 1.1 MiB. Additional media or models cannot be assumed to fit. WP13 must create capacity headroom through a tested packaging or delivery decision before further asset growth. [E3, E4]

Use NIST SSDF and the selected OWASP ASVS requirements as explicit review references. Record their versions in the security plan and bind each chosen control to implementation evidence and a test; naming a framework does not complete its requirements. [R2, R3]

22 Discovery Lab and adoption strategy

WholeContext Discovery Lab™ is the fifth portal and the public entry to the platform's data and evidence work. Its role is to make source meaning, limitations and useful next questions understandable, then let interested users choose a deeper experience.



Public discovery and identifiable care measurement remain separate.

Figure 3. Proposed adoption paths. Public exploration can lead to a fictional care journey or an organizational conversation. A real care workflow requires separate consent, authority and G3 approval.

Value before expansion

The first useful result is a source grounded answer to a bounded question: what data exists, for which population, place and period, what it contributes, and what remains unknown. The next useful result is a fictional demonstration of how reviewed evidence becomes a packet and a collaborative action.

Distribution and partnership hypotheses

Test trusted discovery through patient advocacy organizations, complex care clinics, research or education communities and clinician champions. Each channel needs a relevant demonstration, an explicit invitation and a way to record useful feedback. An interested organization is a lead, not a contracted partner; no outreach or partnership is implied by this document.

The product led loop is collaboration: reviewed evidence creates a useful packet; a permitted recipient contributes or corrects it; the patient sees a more complete story and returns when another care event occurs. Success is accepted value and reduced coordination friction. Session duration and repeated clicks alone are weak evidence. [E5, E7]

23 Adoption experiments and stage decisions

Use small, consent appropriate experiments to learn which users reach value and why others stop. The experiments below are proposed; targets and sample sizes are set before data collection, after baseline usability work.

Experiment	Design and primary evidence	Decision supported
Discovery comprehension	Ask intended users to locate a source, identify its unit and period, and explain one limitation; compare alternative source views.	Improve labels or navigation before increasing traffic.
Synthetic first value	Observe a fictional journey from selected record to understandable packet; measure completion, time and assistance.	Choose the initial onboarding path and reduce avoidable steps.
Clinician usefulness	Review a blinded or counterbalanced set of packets against source material; record edits, omissions and next action clarity.	Refine packet scope and choose a bounded pilot use case.
Supporter collaboration	Test invitation, approved contribution and revocation in a controlled hosted setup; assess understanding and task closure.	Decide whether the sharing workflow is usable and enforceable.
Partner feasibility	Structured workflow interviews and a reviewed pilot brief covering ownership, data flow, support and procurement.	Determine whether a qualified pilot partner and sustainable workflow exist.
Return value	In an approved pilot, measure return at a subsequent meaningful care event and the useful action completed.	Evaluate longitudinal utility without equating screen time with benefit.

Commercial hypotheses

Patient subscription, clinician seat, care program and organization pilot models remain alternatives to test. Record who benefits, who operates the workflow, who pays, integration effort and cost to serve. Use actual willingness to pay and procurement evidence before choosing a price or presenting a market forecast. [E5]

A proposed experiment record contains hypothesis, target users, recruitment method, consent, task, primary measure, decision threshold, sample rationale, stopping rules, results and next action. Report unsuccessful attempts and drop off reasons. Do not select only favorable quotations or completed sessions.

An ROI analysis requires observed workflow time or resource changes, a defensible comparator and implementation and support costs. Public traffic, synthetic task completion and a generated packet are not financial savings.

24 Success measurement dictionary

The measures below are proposed definitions for consistent reporting. Baseline values and decision thresholds are not yet established. Report denominator, cohort, version and observation window with every rate.

Metric	Definition	Owner and limit
Review fidelity	Supported material statements / audited material statements; separately report critical omissions and contradictions.	Clinical / QA; fixed audit rubric.
Lineage completeness	Material output elements resolving to an accepted source or explicit assumption / all required output elements.	Data; assumptions remain a separate class.
Unauthorized disclosure	Count of confirmed disallowed reads, exports or writes; report tested deny cases and incidents separately.	Security; zero observed is not proof of absence.
Task completion	Participants completing the prespecified task without prohibited assistance / participants who attempted it.	UX; record exclusion and assistance rules.
Time to reviewed packet	Elapsed active task time from task start to an accepted packet, reported as a distribution with incomplete tasks.	UX / Clinical; no invented time savings.
Meaningful activation	Consenting eligible users reaching a reviewed record plus either an accepted packet or an approved contribution within the stated window / eligible users starting that workflow.	Product; synthetic and real cohorts separate.
Event based return	Activated participants completing a useful action at a subsequent eligible care event / activated participants observed to have such an event.	Research; report event ascertainment and missingness.
AI fallback and review burden	Requests returning declared fallback / valid requests for model assistance; edits and review time per accepted extraction.	ML; quota reservations and successful completions separate.
Source freshness	Selected source outputs within their freshness policy / all selected outputs subject to that policy.	Data; publisher lag and ingestion lag separate.
Cost to serve	Measured infrastructure, model and support costs / prespecified completed workflows in the same window.	Operations; setup and fixed costs reported separately.

Patient reported function and burden are analyzed as clinical or workflow outcomes according to the protocol, with missingness and confounding disclosed. They are not blended into a single engagement score. A proposed primary pilot success criterion must name its measure, threshold, uncertainty and safety conditions before enrollment.

25 Telemetry specification and measurement safeguards

Telemetry should explain reliability and accepted value without reproducing a care record. Separate public synthetic analytics, authenticated workflow measurement and research outcome data.

Event family	Minimal proposed fields	Collection boundary
Public source exploration	Event name and schema version; coarse time; release; source family; view outcome; synthetic mode.	No patient ID, precise location, search narrative or health text.
Workflow milestone	Pseudonymous study or workflow ID; approved role class; milestone; outcome; release; coarse duration.	Collect only under the approved purpose; keep the identity mapping in a restricted store.
Technical failure	Correlation ID; route class; error category; duration bucket; configuration version; fallback flag.	No request body, source quotation, credentials or raw model response.
Data pipeline	Source release ID; transform version; row counts; quality status; retrieval and completion time.	Publisher and engineering metadata; restricted source terms still apply.
AI operation	Task class; model configuration ID; outcome; token or cost totals where available; latency; fallback reason.	No free text; separate attempted, reserved, completed and accepted states.

Reliable counting

Assign an event ID and schema version; deduplicate retries; define event time and ingestion time; record consent state where required. Reconcile emitted events with observable application state on fictional fixtures. A page view cannot stand in for a completed review, and a generated packet cannot stand in for recipient acceptance.

Retention and access

At G2, privacy approves purpose, access and retention for each event family. Retain only as long as justified; aggregate where possible and suppress small groups before public reporting. Do not place advertising trackers in care flows or link public discovery behavior to identifiable care histories for marketing.

For missing telemetry, distinguish no event, offline or blocked collection, nonconsent and true task failure. Report how each affects denominators. Keep withdrawal and deletion behavior consistent with the study protocol and permitted audit retention.

The September receipt establishes that milestone telemetry was designed and tested in a historical checkpoint. WP15 and T12 must prove the chosen event schema and privacy behavior in the actual pilot candidate. [E7]

26 Accountabilities and review independence

The proposed roles describe responsibilities to staff, not a roster of people already committed. One person may hold several implementation roles, but should not be the sole independent approver of their own critical security or clinical evidence.

Decision or artifact	Accountable	Responsible	Required review
Intended use and scope	Product lead	Product / Clinical	Engineering, privacy, patient adviser
Service and evidence contracts	Engineering lead	Backend / Data	Frontend, QA, clinical
Source admission	Data lead	Data engineering	Privacy and domain reviewer
Access and security controls	Security lead	Backend / Platform	Independent security reviewer
Data lifecycle and authority	Privacy lead	Privacy / Operations	Partner data owner, legal as needed
AI evaluation and promotion	Clinical safety lead	ML / Clinical research	Independent methods reviewer
Accessibility and role journeys	Product lead	UX / Frontend	Users, accessibility reviewer
Release evidence	Engineering lead	QA / Operations	Package owners and independent reviewers
Pilot protocol and stop rules	Study lead	Clinical research	Required institutional and partner authorities
G3 activation decision	Product lead	Delivery coordinator	Engineering, clinical safety, privacy and security concurrence

Working cadence

Each iteration reviews accepted increments, risks, remaining effort and evidence quality. A weekly integration meeting resolves cross workstream dependencies. The data and model review examines any proposed source, feature, prompt or model promotion. Gate reviewers receive the evidence package before the decision meeting.

Maintain a decision register with date, owner, alternatives, rationale, evidence, affected requirements and revisit trigger. A participant can raise a safety or privacy concern outside the delivery chain. Assign a named deputy and escalation route before any real participant pilot.

Advisers review consent, fatigue burden, accessibility, clinical interpretation and collaboration. Compensate participation where appropriate and document conflicts of interest. Engagement does not imply institutional partnership, endorsement or recruitment approval.

27 Capacity model and budget assumptions

This is an illustrative planning model. No rates are market quotations, no staff are committed and no expenditure is authorized. Validate role availability and external costs at G0.

Proposed average capacity	FTE	Proposed average capacity	FTE
Product and delivery	0.5	Backend and platform	1.3
Frontend and UX	1.1	Data engineering	1.0
ML engineering	0.6	QA and test engineering	0.8
Security and privacy	0.5	Clinical and methods	0.5
Operations	0.2	Total	6.5

24 weeks × 6.5 FTE = 156 calendar person weeks. At 80 percent productive availability, capacity is 124.8 productive person weeks. The 105 person week work package midpoint leaves 19.8 productive person weeks of reserve. The upper estimate of 127 exceeds capacity by 2.2; even the midpoint can slip if scarce roles or approvals determine the schedule.

Scenario	Productive capacity in 24 weeks	Planning consequence
Lean 3.5 FTE	67.2 person weeks	Midpoint work requires at least 37.5 weeks by effort alone; reduce scope or extend the plan.
Reference 6.5 FTE	124.8 person weeks	Covers the midpoint plus 19.8 reserve; role leveling and dependency review still required.
Expanded 8.0 FTE	153.6 person weeks	Adds capacity, but cannot compress approvals, independent review or study observation by arithmetic.

Worked budget illustration

Assume a blended loaded cost of \$4,000 per calendar person week solely for this calculation. The 156 calendar person weeks cost \$624,000. Add an illustrative \$25,000 nonlabor allowance for tools, infrastructure and external expenses: total \$649,000. This includes the capacity reserve and nonproductive allowance; do not add them again as separate labor contingencies.

Actual rates, independent assessment fees, contracts, hosting and participant costs require quotes or agreements. The \$25,000 is a placeholder, not evidence of sufficiency. Post G3 pilot operation, participant observation, clinical effectiveness studies and large external integrations are excluded and need separate budgets.

Track committed cost, actual cost, estimate to complete and remaining capacity at each gate. Do not spend the reserve on additional features before the dependency and risk position has been reassessed.

28 Risk register and active responses

Risks below are prospective planning judgments. Likelihood and impact need refinement against the staffed environment and partner workflow; no numerical risk score is represented as measured.

Risk and trigger	Response and decision	Owner
R01 Evidence overclaim Cataloged or planned work described as live or validated.	Claims register, source citation and wording review; correct submission and release notes before publication.	Product / QA
R02 Identity or scope defect Disallowed read, stale patient state or ambiguous grant behavior.	Block patient activation; contain access; repair and rerun hosted T02 and T09.	Security / Backend
R03 Source incompatibility Changed schema, unsupported geography or missing denominator.	Quarantine source; preserve last accepted release with stale label or disable output; T07 retest.	Data
R04 AI or clinical misinterpretation Unsupported output, inappropriate reliance or safety concern.	Disable affected assistance; preserve evidence; clinical review; revise limits and evaluation before restart.	Clinical / ML
R05 Specialist bottleneck Critical tasks or independent review unavailable for an iteration.	Relevel resources; defer optional scope; update gate dates and estimate to complete.	Delivery
R06 Platform or package limit Capacity headroom or recovery objective not met.	Freeze asset growth; choose and test packaging or portability decision; hold release if unrecoverable.	Platform
R07 Authority or partner delay Data terms, protocol or workflow ownership unresolved.	Continue fictional engineering; no real data intake; rebaseline pilot window.	Privacy / Study lead
R08 Adoption without value Traffic grows while comprehension or accepted tasks remain poor.	Fix the value path; interview noncompleters; pause acquisition experiments.	Product / UX
R09 Study selection bias Low diversity, differential attrition or incomplete follow up.	Report limits; revise recruitment or analysis; avoid broad effectiveness claims.	Methods
R10 Operating cost growth Measured cost or support load exceeds approved envelope.	Apply quotas and scope controls; investigate drivers; revise service design or pilot scale.	Operations

Review active risks weekly and at every gate. Each activated risk receives evidence, a named owner, next action and review date. Stop rules outrank schedule or prize objectives. A mitigation is closed only when its supporting test or operational evidence is retained.

29 Submission and public release package

The submission should make the work easy to inspect: a clear claim, a reproducible example and an evidence trail. This roadmap prepares the engineering package; it does not certify competition eligibility or claim that any upload has been submitted.

Artifact	Reviewer question answered
Technical Clinical and Impact Dossier	What problem is addressed, what has been built, what evidence exists and what is still unproven?
171 Resource Data and Engineering Matrix	What is cataloged, how inventories overlap, which sources are implemented and how proposed outputs trace to inputs?
System Architecture and Engineering Design	How do identity, provenance, service contracts, AI boundaries and permissions work, and what must be verified?
This Transformation Roadmap	What work comes next, in what sequence, with which resources, gates, risks and measures?
Release evidence and demo guide	Which exact version should a reviewer use; which fictional journey demonstrates a claim; what are the known limitations?
Methods and evaluation package	Which datasets, labels, splits, baselines, metrics and failure cases support each model or workflow statement?
Governance and operations package	Who may use which data; how is access enforced; how are incidents, recovery, retention and withdrawal handled?
Claims and limitations register	Which statements are supported today, which are hypotheses and which require later evaluation?

Final assembly controls

Verify the authoritative submission requirements and attachment limits at assembly time. Check document titles, version references, links, captions, accessibility, fictional labels and redaction. Confirm that public files contain no credentials, private correspondence, patient identifiers or restricted source data.

Freeze the final source and data manifest, preserve the submitted files and retain the submission receipt when actually obtained. Do not replace the public demonstration with an unverified late change solely to match a new document. If a change is necessary, create and test a new candidate with its own receipt.

A concise reviewer path should begin with one synthetic care journey, show selected source provenance and packet parity, then point to the matrix, architecture and roadmap for deeper scrutiny. The package should be persuasive because its claims can be checked.

30 Transformation beyond pilot readiness

The longer horizon is conditional on evidence from the previous gate. These options remain outside the reference 24 week scope unless a controlled change explicitly replaces other work.

Horizon	Candidate capability	Evidence before expansion
After G3	Governed patient workflow pilot with review, selected sharing, support contributions and longitudinal follow up.	Approved protocol, operational coverage and authority; G4 feasibility analysis.
After G4	Broader organization adoption and workflow integration.	Demonstrated utility, acceptable burden, reliable support and measured cost to serve.
Partner dependent	FHIR or other healthcare system adapters; patient authorized wearable and document sources.	Agreed profiles and terminology, identity matching, permissions, round trip fidelity and partner testing.
Research dependent	Richer temporal NLP, reviewed literature retrieval and multimodal evidence assistance.	Authorized corpus, expert labels, retrieval quality, source fidelity, adversarial and subgroup evaluation.
Validation dependent	Condition specific clinical prediction or image interpretation.	Suitable intended use, representative labels, independent validation, clinical and regulatory review, monitoring and stop rules.
Scale dependent	Expanded public data federation and additional condition pathways.	Source admission, sustainable refresh, condition specific evidence contracts and appropriate evaluation.

Portability by contract

Keep domain logic, evidence snapshots and adapter interfaces separable from the hosting runtime. Export source manifests, schemas and configuration references without secrets. Before a migration, prove identity and grant equivalence, restore, audit continuity, latency and cost. Historical Azure discussions do not establish a current Azure deployment. [E5]

The digital twin remains a dated evidence reconstruction with isolated assumptions until a different model is explicitly specified and validated. Treatment and recovery views can organize observations without claiming causal attribution. Wearable values need device, sampling and unit provenance before they become comparable longitudinal measures.

Every new condition, user population, model function or data recipient can change the risk and evaluation problem. Reuse architecture where justified; do not reuse a validation claim beyond its population, task and environment.

31 Requirement traceability for core evidence

The following crosswalk carries the first half of Document 3's requirement set into this delivery plan. A mapped test is a planned evidence obligation, not a pass result. [E3]

REQ	Required behavior	Delivery packages	Acceptance and gate
01	Patient identity stays consistent across records and outputs.	04 11 17	T01; G2 foundation, G3 end to end
02	Material outputs resolve to sources or explicit assumptions.	04 07 11 17	T01; G2 foundation, G3 outputs
03	Missing, invalid, excluded, hypothetical and zero remain distinct.	04 07 08 17	T01; G2 foundation, G3 outputs
04	Access uses current membership, role, recipient and record scope.	05 12 14 17	T02; G2 hosted foundation, G3 regression
05	Concurrent edits cannot silently overwrite reviewed work.	04 06 17	T03; G2, retested G3
06	AI failure preserves sources and discloses fallback.	09 13 17	T04; G3
07	Dated scenarios exclude later facts and isolate assumptions.	04 08 11 17	T05; G3
08	Release build and recovery are reproducible and reversible.	01 13 18	T06; G0 baseline, G3 recovery
09	Temporal boundaries and timezone rules are consistent.	04 08 17	T05; G3
10	Revocation, expiry and regrant have explicit tested semantics.	05 14 17	T02; G2, retested G3
11	Public request and response schemas are versioned and tested.	03 07 17	T07; G2 contracts, G3 integration
12	Sharing and related audit changes maintain atomic integrity.	06 13 17	T03; G2, retested G3

For each requirement, the acceptance record identifies the implemented design, fixture set, positive and negative cases, results and reviewer. Where a feature is excluded from the pilot, record the exclusion and test that it remains disabled. Do not replace a mandatory requirement with a prose assurance.

WP17 coordinates integrated verification; it does not remove ownership from the implementing package. A source refresh, model change, migration or permission change can invalidate earlier evidence and trigger the affected campaign again.

32 Requirement traceability for advanced controls

This second half of the crosswalk links ingestion, AI, images, privacy and operations to explicit delivery owners and the final pilot readiness gate. [E3]

REQ	Required behavior	Delivery packages	Acceptance and gate
13	Requests obey streaming byte, structure and depth limits.	03 06 10 17	T03; G2 foundation, G3 all intake
14	Ingestions retain rights, hashes, counts, transforms and quarantine.	07 14 17	T07; G2 foundation, G3 admitted sources
15	NLP evaluation uses expert labels and exact span semantics.	09 16 17	T04; G3
16	Images have decode, quarantine, access and lifecycle controls.	10 14 17	T08; G3
17	Exports bind identity, record revisions and source provenance.	04 11 17	T09; G3
18	Patient changes invalidate stale asynchronous and browser state.	02 11 12 17	T09; G1 synthetic, G3 pilot
19	Model promotion has a card, rights, evaluation and rollback.	09 13 16 17	T10; G3 for enabled AI; G5 for new prediction
20	Selected security controls map to evidence and independent review.	05 06 13 14 17	T11; G2 design, G3 exercised controls
21	Retention and erasure cover stores, backups and derived artifacts.	10 13 14 17	T11; G2 design, G3 exercise
22	Clinical AI functions have reviewed limits and stop criteria.	09 10 16 18	T10; G3 and later promotion gates
23	Errors use a consistent taxonomy and safe correlation metadata.	03 13 17	T12; G3
24	Minimal telemetry measures load, latency, fallback and cost.	13 15 17	T12; G3

Test campaign status

The architecture baseline identifies relevant partial implementation for several campaigns and planned work for image pipeline, model promotion, security and operations coverage. This roadmap does not upgrade those statuses. Only a retained test report against a named release and environment may do so.

Public synthetic verification, controlled component tests, historical research results and prospective participant studies remain separate evidence classes. Each is useful for a different decision. Their counts must not be combined into an overall accuracy or readiness percentage.

33 Project evidence register

These project artifacts ground the completed work and the planning constraints. The three companion volumes reconcile earlier material against the inspected v39 baseline. They are review artifacts, not independent certifications.

ID	Artifact	Use in this roadmap
E1	WholeContext AI Technical Clinical and Impact Dossier v1.0, 14 October 2026, 48 pages.	Product scope, evidence classes, clinical limits, research history and full resource index.
E2	WholeContext AI 171 Resource Data and Engineering Matrix v1.0, 14 October 2026, 14 worksheets.	Catalog, atlas and runtime reconciliation; 198 lineage mappings; contracts, methods, joins and proposed outputs.
E3	WholeContext AI System Architecture and Engineering Design Specification v1.0, 14 October 2026, 41 pages.	Inspected implementation, 24 requirements, twelve acceptance campaigns and target controls.
E4	WholeContext v39 Publication Receipt and inspected source at commit 4a85dc7a60a989e2c731cb9ed2ed29ffdba7cbe1.	Public release boundary, validation limitations, package constraint and rollback context.
E5	WholeContext Technical Strategy Architecture Addendum, 12 October 2026; Federal and Open Data Catalog and Access Roadmap v0.1, 4 October 2026.	Commercial and multimodal hypotheses; source admission and expansion rationale. Earlier inventories and version labels are superseded by E2 where reconciled.
E6	Real Data Modeling Methodology v0.2; Synthetic Longitudinal Data Methodology v0.2; Research Foundation and Evidence to Design v0.1, as indexed in E1–E2.	Research method boundaries, synthetic evidence and separation from clinical validation.
E7	Adoption and Value Telemetry Receipt, 27 September 2026, checkpoint fbc0283fe31946485750daad9ff8318f1b8ee0c2.	Historical milestone design; 13/13 platform tests reported with full browser build not run and full typecheck blocked at that checkpoint.
E8	Evidence Ontology and Orchestration Plan; Data Preparation Selection and Management, as indexed in E1.	Source identity, normalization, provenance and review sequence.
E9	Health Data Atlas retrieved workbook and September source catalog, preserved in E2.	107 and 171 record source inventories, with the atlas README version 1.3 correction.

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

The website is mutable. Use the release identity in E4 for baseline claims. This roadmap does not assert a fresh production audit, hosted access test or active infrastructure inventory beyond the evidence described in the companion architecture volume.

34 External review references

Primary references below were checked on 14 October 2026. They inform proposed engineering and review practices; the project remains responsible for selecting applicable requirements and demonstrating them.

R1 NIST AI Risk Management Framework

AI RMF 1.0 and the Generative AI Profile provide a structure for risk governance and evaluation. This roadmap uses that structure for intended use, ownership, measurement and promotion decisions.

<https://www.nist.gov/itl/ai-risk-management-framework>

R2 NIST Secure Software Development Framework

SP 800 218, SSDF Version 1.1, is the pinned final reference for the proposed secure development process.

Record any later adopted revision explicitly.

<https://csrc.nist.gov/pubs/sp/800/218/final>

R3 OWASP Application Security Verification Standard

ASVS 5.0.0 is the proposed control mapping reference. Select relevant requirements with their versioned identifiers and verify them in the pilot environment.

<https://owasp.org/projects/asvs>

R4 FDA Clinical Decision Support Software

The January 2026 final guidance is a reference for qualified review of intended software functions. No device status or clearance conclusion is made by this roadmap.

<https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>

R5 HHS Security Rule

HHS material informs the assessment of applicable responsibilities and safeguards. Applicability, contracts and implementation require qualified review.

<https://www.hhs.gov/hipaa/for-professionals/security/index.html>

R6 W3C Web Content Accessibility Guidelines

WCAG 2.2 is the proposed accessibility reference for critical journeys. A conformance claim requires appropriate evaluation of the actual experience.

<https://www.w3.org/TR/WCAG22/>

Framework alignment is not a substitute for security testing, clinical validation, legal review or accessible user testing. The source references identify the standards used in planning; detailed acceptance criteria remain the project's responsibility.

35 Assumptions and decisions for mobilization

This final register makes the plan reviewable before anyone commits resources. It records what must be decided at G0 and what evidence would require the plan to change.

Planning assumption	Decision or verification needed	Change consequence
A1 v39 is the review baseline.	Confirm the candidate to develop; preserve its source and publication receipts.	Reconcile differences and rerun affected baseline checks.
A2 Initial pilot is review and collaboration.	Approve intended use, enabled functions, roles and prohibited automation.	New predictive or diagnostic functions require a separate evaluation path.
A3 Reference horizon is 24 weeks.	Approve funding, role availability and week 1 start; level the specialist workload.	Extend dates or reduce optional scope if capacity or dependencies fail.
A4 Effective availability is 80 percent.	Compare with actual throughput and review overhead after two iterations.	Revise capacity, reserves and cost forecast.
A5 Sources are selectively admitted.	Select the core from existing snapshots and contracts; verify rights and purpose.	Defer inaccessible sources; preserve explicit missing context.
A6 No clinical performance is established.	Approve evaluation questions and claims register; commission the methods review.	Do not expand claims until the corresponding evidence exists.
A7 Partner and oversight are unresolved dependencies.	Name the workflow owner, data authority and protocol approval route.	Keep development fictional and rebaseline enrollment timing.
A8 Budget is an illustration.	Replace assumed rates and allowance with staffing commitments and quotes.	Reapprove the resource envelope before spending.
A9 Telemetry is minimal and purpose bound.	Approve event schema, consent, access and retention; verify actual release behavior.	Disable collection that lacks authority or cannot meet the agreed safeguards.

Required G0 decision record

Approve or revise the intended use, delivery scope, role assignments, capacity envelope, source selection process, gate authorities and next review date. Record deferred work and the evidence required to reopen it. Do not interpret document acceptance as pilot activation.

The program succeeds when a reviewer can move from a public claim to its source, implementation, test and responsible decision, and when the people using the workflow can understand and control how their evidence is interpreted and shared.