



# CASCADE BRIDGE

Patient-Owned Health Data. Connected. Augmentable. AI-Ready.

## 1. Description of Solution and Problem Addressed

The 21st Century Cures Act and ONC's December 2023 EHI export mandate gave every patient a legal right to their complete health record. Over 45,000 certified providers now offer single-patient EHI exports, and TEFCA facilitated the exchange of more than 500 million records by early 2026. The data transport problem is largely solved, and FHIR is its centerpiece. However, access is not the same as usability. What arrives at the patient is a raw FHIR bundle, a sprawling C-CDA XML archive, or an unstructured PDF. These formats were engineered for institution-to-institution exchange, not for a patient managing a chronic condition across five providers. *FHIR is an exceptional shipping container, but it does not deliver usable, readable, or actionable information to patients.*

### The 3 C's

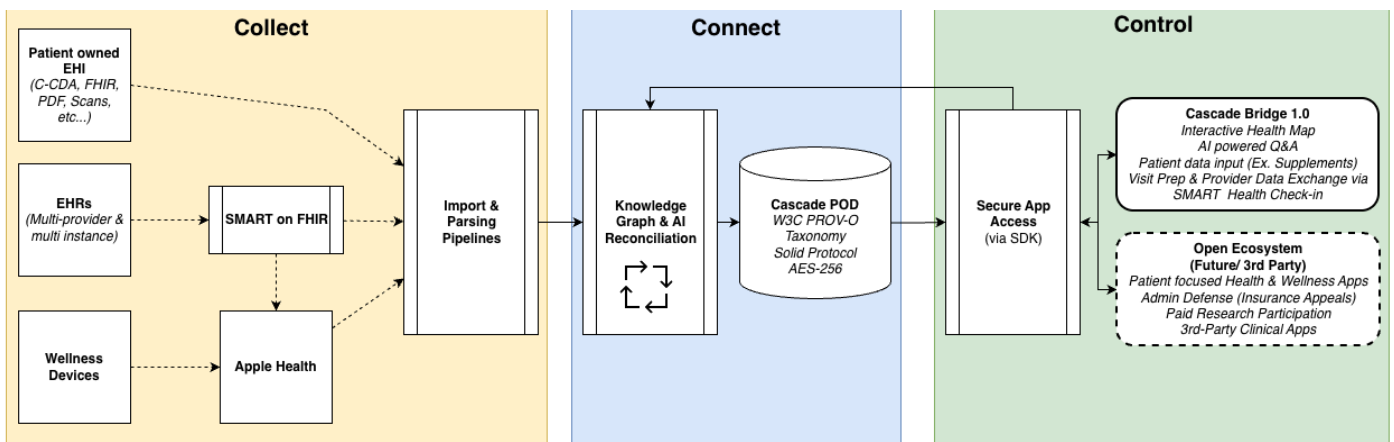
Patients face three distinct barriers: **Collecting** data scattered across providers, portals, and devices; **Connecting** those fragments into a single, deduplicated, computable picture; and **Controlling** that data (augmenting it, sharing it selectively, and reasoning over it with AI they can trust).

**Cascade Bridge** is a reference application, built on the open-source **Cascade Protocol** that closes these gaps. It ingests all types of clinical and wellness data including EHI exports from any FHIR or C-CDA compliant system, PDFs, wearable streams, and self-reported data. It then produces a navigable, semantically linked, locally encrypted health knowledge graph with complete provenance tracking. The result is a health record patients can actually use: explore, augment, query with AI, and share selectively, with every fact traceable to its source.

### EHIgnite Scenarios Addressed

Cascade Bridge directly addresses the challenge's baseline requirement to transform single-patient EHI exports into usable, readable, and actionable information. The architecture additionally covers four of the five scenarios:

Scenario	Cascade Bridge Coverage
1. Interactive Patient Tools	Health Map, 3D graph, AI Q&A with citations, visit preparation
2. Clinical Domain Customization	Conditions-first organization; domain-specific apps via Layer 3 vocabulary.
3. Integration Across Settings	No data left behind; multi-source ingestion and reconciliation
5. Participant-Defined Use Case	Open platform: SDKs, CLI, open protocol spec, enabling 3rd party apps.



**Figure 1: The Cascade Bridge Architecture.** Raw EHI exports are collected into a secure Solid pod, connected via an RDF knowledge graph with strict provenance. Open Source SDKs and tools enable developers to easily leverage the architecture.

## The 3 C's: (Collect, Connect, Control) Patient Experience

To make this concrete, meet *Priya*, a 45-year-old project manager. Her mother had ovarian cancer at 51. Two years ago, a consumer genetic test reported that Priya carries a *BRCA2 Variant of Uncertain Significance (VUS)*. Today Priya manages her cancer surveillance through her oncologist, sees a primary care doctor in a different health system for routine care, works with a behavioral health therapist after her mother's diagnosis, and tracks daily metrics on her Apple Watch. Her 18-year-old daughter Maya has never been tested. Priya's complete health story exists in seven different places. None of them talk to each other.

### Collect: *No data left behind.*

**The problem today:** Priya has tried to pull her records together. Her oncologist's portal supports SMART on FHIR, so she can download a clean FHIR R4 bundle. Her primary care system only offers C-CDA XML. The genetic counseling clinic mailed her a 12-page PDF. Her behavioral health practice will not share electronically at all and prints discharge summaries on paper. Her Apple Watch streams to HealthKit, her supplement diary lives in a third-party app, and her email contains a year of appointment confirmations she has never aggregated. The Cures Act gave her the legal right to all of this data, but the uneven implementation across providers makes assembly nearly impossible.

**Where FHIR falls short:** FHIR adoption is uneven. Some systems (like the one from her oncologist) offer rich, well-formed bundles, but others emit partial exports. Even when the bundle is complete, FHIR describes more than it actually delivers. Scans, faxes, and PDFs appear as (DocumentReference) resources, but the binary content lives behind the provider's authentication layer. When Apple Health pulled Priya's EHI export, the references arrived, but the content did not. Wearable data and patient-generated information (supplements, symptoms, daily metrics) can be expressed in FHIR, but EHRs rarely ingest them, so they live in separate silos that the bundle never touches.

**The Cascade Bridge solution.** Cascade Bridge is a universal receiver. One tap connects to 800+ health systems via Apple Health's SMART on FHIR integration. C-CDA archives import natively through an 11-section handler with Epic and Cerner normalization. PDFs and scans are parsed by on-device AI and tagged (`cascade:AIExtracted`) so that their provenance is visible at every later step. Apple Watch streams, supplements, and self-reports flow into the same encrypted, patient-owned pod. For patients who don't remember every provider, on-device AI can scan email for appointment confirmations and billing notices to surface sources they may have missed.

**Cascade Bridge accepts any health data, from any source. No format is rejected. No data is left behind.**

### Connect: *Collected data is now usable data.*

**The problem today:** Priya's records are finally in one place, and they immediately fight each other. Her primary care FHIR bundle lists her tamoxifen as "active." Her recent surgical discharge says it was "stopped." Her oncologist references a different dose. Two C-CDA exports from a single portal yield thousands of duplicated labs. Each individual file is internally consistent. Together, they are noise.

**Where FHIR falls short:** FHIR is a hierarchical document tree designed for institutional exchange. Its human-readable text is trapped inside isolated JSON objects, which is fine for transport but poor for cross-source reasoning. Any application that wants to merge Priya's seven records must implement custom reconciliation logic (typically dozens of nested SQL JOINS against shredded FHIR tables, or fragile GraphQL layers), and every downstream app must reinvent that logic. Local-first AI reasoning on a patient's device is computationally hostile to this architecture. There is no shared, computable representation of "Priya, the patient," only seven versions of her record, each stamped by a different system.

**The Cascade Bridge solution.** This is where the **RDF** (Resource Description Framework) knowledge graph enters the scene. RDF is a mature W3C standard built for **interoperability** and built for **reasoning**. Cascade Bridge translates every FHIR resource, C-CDA section, PDF extract, and wearable observation into RDF triples (subject, predicate, object) anchored to the same SNOMED CT, LOINC, and RxNorm codes that the source systems already use.

The (cascade reconcile) engine deduplicates labs by LOINC code, fuses overlapping observations, and validates the graph against SHACL shapes. When it encounters the tamoxifen conflict, it does not silently overwrite. It creates a (cascade:ConflictDetail) node that preserves both records, tags each with provenance (the oncologist, the primary care system, the discharge bundle), ranks them by trust, and surfaces the conflict for Priya to review. The trail of change is kept intact without compromising on authority.

Every fact says exactly where it came from and what it means. The result is one canonical, deduplicated, computable representation of Priya's health, ready for reasoning and AI.

## Control: *Dynamic ownership, selective sharing*

Control is where **Cascade Bridge** unlocks experiences that are fundamentally impossible with static EHI exports. This is where the patient's legal right to access their health data becomes truly actionable.

- **Correctable records, preserved authority:** Priya notices her primary care record still lists a medication she stopped a year ago. She marks it inactive, but Cascade Bridge does not erase the EHR's original assertion. It records her correction with (prov:wasDerivedFrom) linking back to the source fact, with full provenance. Authority is preserved; her record finally matches her life.
- **AI Q&A with citations:** Priya asks her on-device AI: "Why did my LDL spike this year?" The AI traverses the graph, sees that her statin was paused during her surgery, correlates the timing with the lab trend, and answers with explicit citations to the exact lab results, the medication-hold record, and the discharge note. Every claim is reconstructible down to its originating EHI export.
- **Living genomics:** Two years ago, Priya's BRCA2 variant was classified as a "Variant of Uncertain Significance". In a typical portal, that PDF would gather digital dust forever. This morning, an international clinical consortium published an (advisory:CascadeAdvisoryPatch) that reclassifies Priya's exact variant from VUS to Pathogenic. Because her data is stored in a living graph rather than in a static file, Cascade Bridge ingests the patch, traverses her (genomics:PedigreeMember) relationships, and her on-device AI surfaces a plain-language alert regarding the potential impact on her and her daughter Maya. Her agent automatically makes an appointment with a genetic counselor for next week.
- **Visit Preparation:** Prior to the appointment the clinic sends her a **SMART Health Check-in** request, a wallet-mediated request for specific items: insurance, demographics, current medications, family history, the BRCA2 variant report, and an intake questionnaire. Priya opens it on her phone. Cascade Bridge, acting as her health wallet, shows her exactly what each item will share. She approves family history and the BRCA2 variant. She *declines* to share her behavioral health record (not relevant to this visit). The intake questionnaire is pre-populated from her graph; she reviews and signs. The clinic receives a structured, signed, encrypted response with each item categorized. No PDFs, no paper clipboard, no transcription errors. The counselor sees Priya's pedigree and current variant interpretation before she even walks in.

## 2. Description of Submitting Individual, Team, or Entity

---

Jed Reinitz, Founder, Cascade Agentic Labs

After years of leading and scaling user-centric technology products at Amazon, LivePerson, ProQuest, and in New York City government, I founded Cascade Agentic Labs to directly harness the emergent potential of AI to tangibly improve people's lives.

In 2025, while working on a healthcare AI project, my eyes were opened to the tremendous scale of inefficiency, waste, and dysfunction within the industry. The problems are structural, they affect everyone, but as a product builder I saw that they were solvable.

I am a technical product leader from outside healthcare. What I bring to this space is my builder's conviction: patient-owned health data is an infrastructure problem, not a policy problem, and I have the ability to ship working software to prove it.

### Proof of Execution

My core thesis is that AI is not just a feature we build; it is a collaborative development partner. Working as a solo technical founder augmented by AI, I move at the velocity of a small engineering team. In eight months I built and deployed the foundation of the Cascade Bridge architecture:

- **Live App:** POTS Check, an iOS app built on the Cascade SDK, launched in December 2025.
- **Production tooling:** `cascade-cli` and `cascade-agent` live on npm, with lossless FHIR and C-CDA conversion, multi-source reconciliation, and on-device AI operating at 98% evaluation accuracy.
- **Open documentation:** The protocol specification is live at [cascadeprotocol.org](https://cascadeprotocol.org).

### Interdisciplinary & Community Engagement

This work is grounded in active collaboration with clinicians and patient advocacy groups. I have worked closely with the **Bateman Horne Center** on POTS Check v2 to ensure the application accurately captures the complex realities of dysautonomia. I am also collaborating with researchers at the **Ratner Early Detection Initiative (REDI)** on prototypes that apply the Cascade Protocol to early cancer detection, analyzing historical lab results alongside patient medical history.

## 3. Wireframes / Mockups

---

The annotated mockups on the following pages illustrate select screens from the *Collect, Connect, and Control* flow for the iOS version of Cascade Bridge.

- **Figure 2** includes one Collect screen and one Connect Screen. The first shows how **Multi-source import** pulls records from three EHR systems, C-CDA exports, Apple Watch, and scanned documents into a single encrypted pod. The second shows how the **Reconciliation pipeline** merges duplicates, resolves dosage conflicts with trust scores, and surfaces AI-extracted findings.
- **Figure 3** shows three Control views: 1. The **Health map** which transforms a patient's health data into a human readable format that is organized by meaning, not data type. 2. An example of **medication editing** that shows how full provenance of data is preserved. 3. A view of a conversation with the **AI Health Assistant** that includes clear citations to the underlying data sources.
- **Figure 4** is a screenshot from the interactive 3D health graph, which enables a patient to explore their personal health knowledge graph in a visually immersive environment.

The full UX specification for Phase 2 will cover authentication, onboarding, configuration, report generation, selective-sharing flows, and SMART Health Check-in wallet interactions.

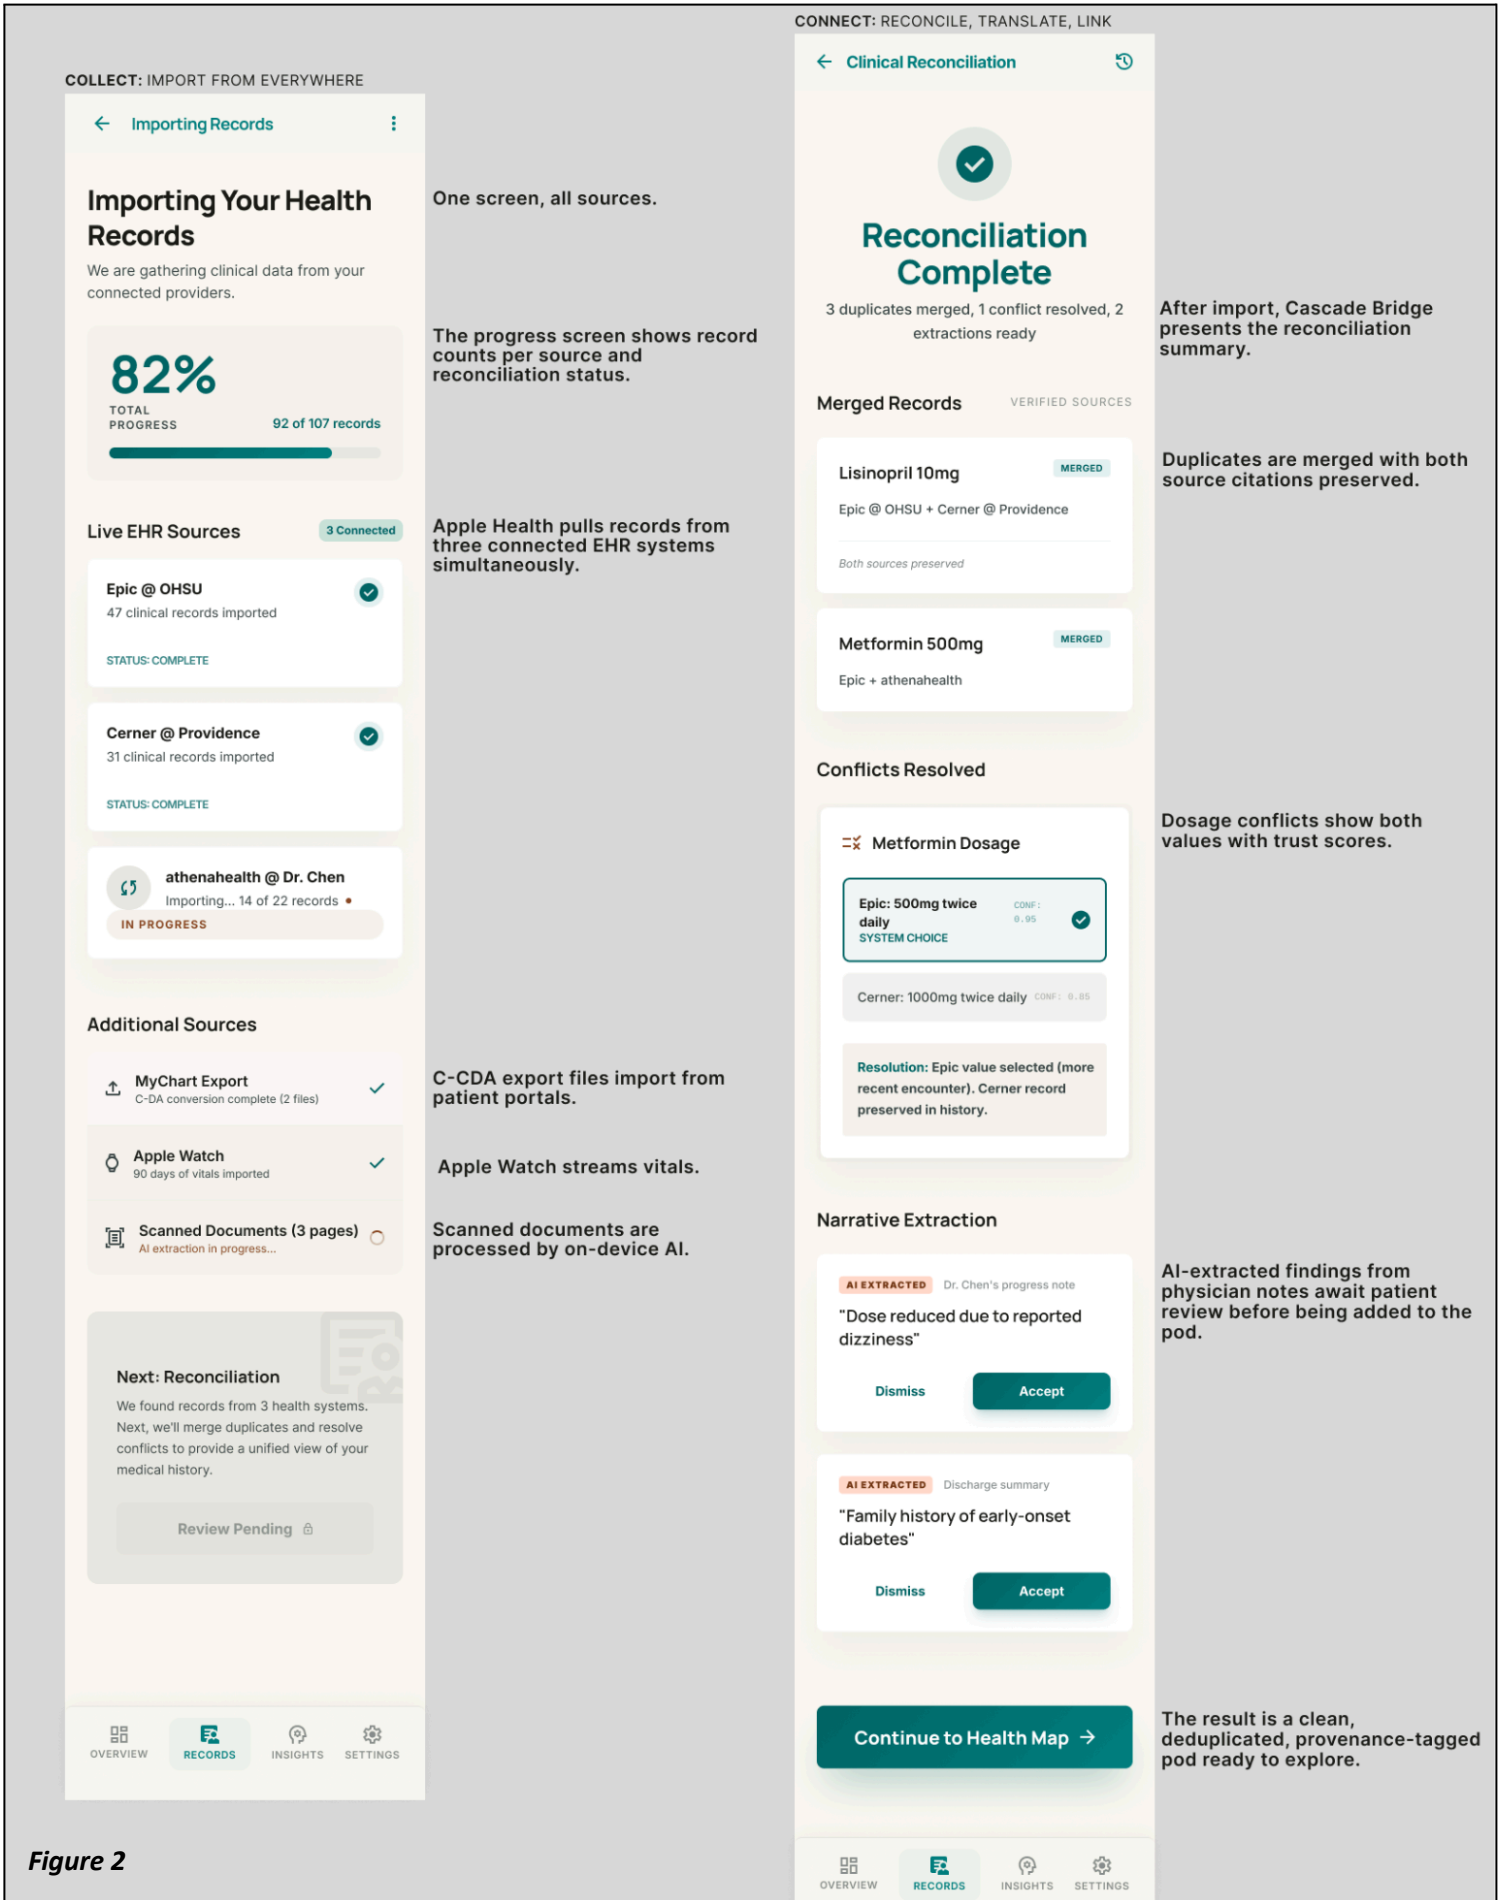



Figure 2




The Health Map organizes the patient's complete record by meaning, not data type.

Conditions cluster their related medications, labs, and family history.

"What I Take" shows all medications and supplements with provenance badges and condition links.

"My Numbers" surfaces flagged labs first.

For My Visit" lists items for patient to address at their next doctor's appointment.

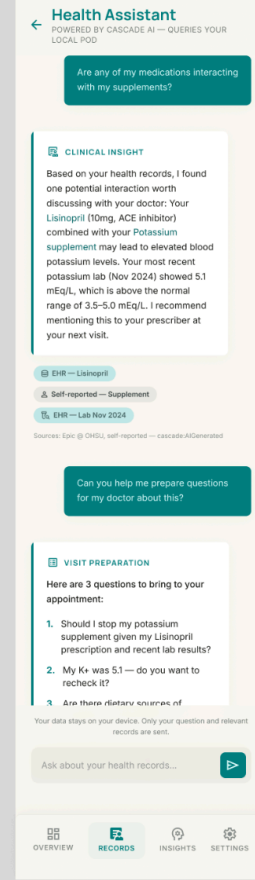


A discontinued medication still showing as active.

The patient marks it inactive and adds a note

The original EHR record is preserved. The correction is tagged as Patient-Corrected.

Both values are visible to any clinician who queries the pod.



Cascade Bridge connects three records from three sources: the lisinopril prescription (EHR-Verified, Epic), the potassium supplement (Self-Reported), and the elevated K+ lab (EHR-Verified, OHSU).

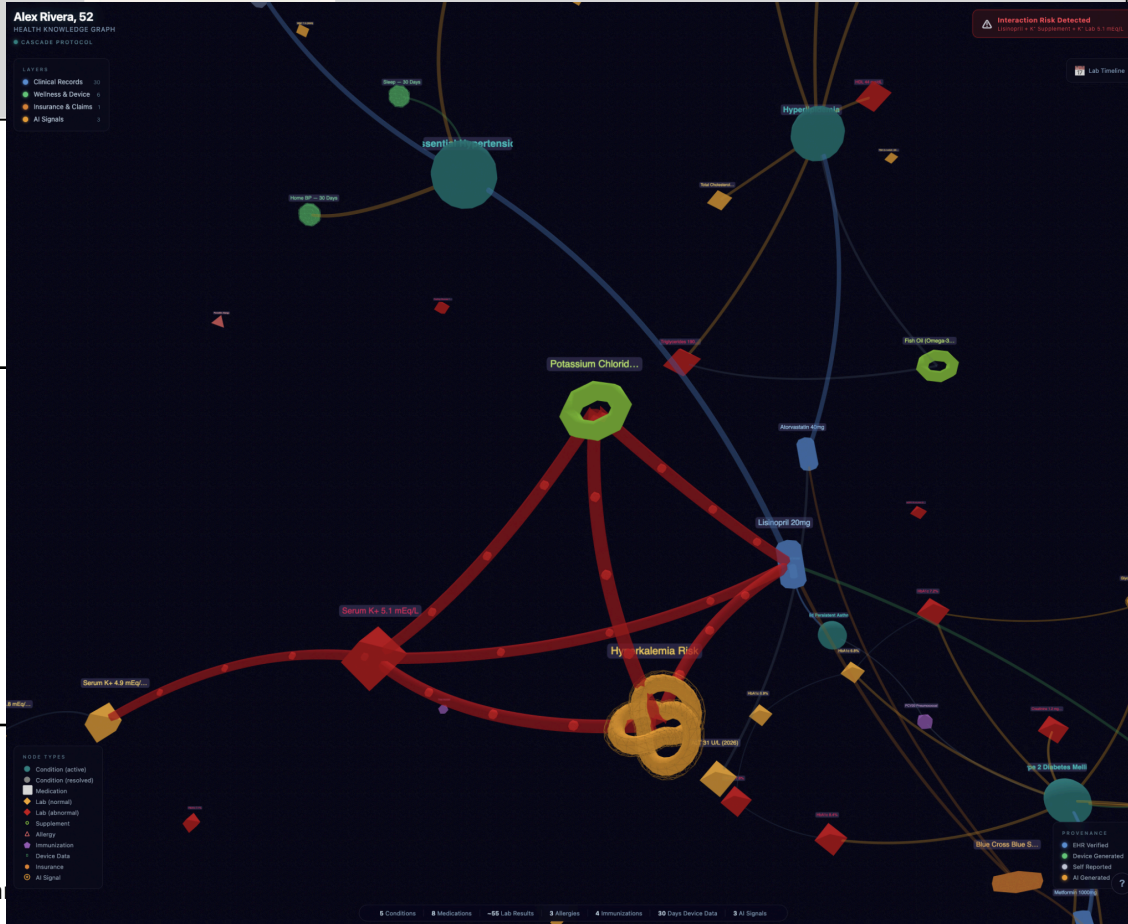
The answer cites each source with its provenance tag and explains the clinical mechanism in plain language.

A generic chatbot guesses from context. Cascade Bridge reasons over attributed facts in a linked graph.

Figure 3

Figure 4

Color encodes provenance.  
 Conditions (spheres)  
 Medications (cylinders)  
 Labs (octahedra)  
 Supplements (tori)  
 K+/ACE interaction shown as a pulsing amber triangle.



## 4. Technical Feasibility and Scalability

### Standards Alignment

Cascade Bridge is not a theoretical architecture; the core pipeline is already functioning in production. Its feasibility rests on a pragmatic, “established vocabularies first” design principle. Rather than asking providers to change their systems, Cascade ingests the formats they already export, and anchors every fact to FHIR, LOINC, SNOMED CT, RxNorm, and PROV-O before introducing any Cascade-specific term. When patients share data back with care teams, the outputs are grounded in canonical medical terminologies using formats and protocols that care teams already trust. The system is also USCDI v7 ready: new data elements map to the existing vocabulary with no schema revision required.

Standard	Ver.	Cascade Implementation
FHIR	R4	Lossless conversion; 0% silent drops (vs. 47.7% baseline)
HL7 C-CDA	v2.1	Native converter; 11 section handlers; Epic + Cerner normalization
US Core IG	v6	Full profile conformance; required terminologies (LOINC, SNOMED CT, RxNorm) preserved on import
SMART App Launch	v2	Apple Health bridge + direct EHR authorization
W3C PROV-O	1.0	Six-level provenance taxonomy; every record carries its lineage
W3C RDF (+ SHACL + OWL)	1.1 / 1.0 / 2.0	Knowledge graph foundation in Turtle; shape validation and formal semantics for AI safety guardrails
Solid Protocol	0.9	Patient-owned pod storage; WebID access control

### Security Architecture & Regulatory Alignment

Patient-owned data demands the highest standard of privacy and security. Cascade Bridge meets and exceeds HIPAA's technical safeguards by design, even where the law does not require it to.

- **Zero-Access Security:** Patient data is encrypted at rest (AES-256-GCM) with keys stored privately in the user's hardware-backed Secure Enclave. There is no centralized server to breach and no operator with administrative access to patient records. This exceeds HIPAA's addressable encryption-at-rest standard and eliminates the most common and damaging category of health data breach.
- **Patient-Directed Access:** Cascade Bridge receives data through patient-directed access, so it is neither a covered entity nor a business associate under HIPAA. The applicable rule is the FTC Health Breach Notification Rule (effective July 2024), and Cascade's NIST-standard encryption qualifies the data as "secured" under the rule.
- **HTI-5 Alignment:** The December 2025 ONC HTI-5 proposed rule continues the trajectory toward recognizing patient-directed AI agents as legitimate actors in health data access. Cascade Protocol's per-query consent and full audit provenance are designed for this direction.

### Scalability

Cascade Bridge's architecture is designed to scale efficiently along three distinct vectors:

- 1. Patients:** Each pod is self-contained on the patient's own device and computation runs on patient hardware, so adding users does not increase load on any shared system.
- 2. Providers:** (`cascade-cli`) processes any standards-compliant FHIR R4 or C-CDA export without per-EHR integration work. Every new certified provider (45,000+ today) is automatically in scope.
- 3. Developers:** The protocol scales by addition, not reimplementation. Third-party teams inherit Cascade's full Collect-Connect-Control stack and only author Layer 3 vocabulary for new domains. Each new application expands the ecosystem and compounds the value for everyone.

## 5. Innovation: Linked Health Data and Explainable AI

Most health apps treat interoperability as a data transport problem. *Cascade Bridge treats it as a knowledge representation problem.* The core innovation is the shift from flat, isolated FHIR documents to a semantically linked, AI native, knowledge graph, built for reasoning and ready for the *agentic age*.

**Why a new format if FHIR is winning?** Cascade Bridge does not replace FHIR. It is the AI-ready downstream consumer of it. We take the FHIR bundles a provider already exports, fuse them with C-CDA archives, PDFs, scans, and wearable data, and upgrade them into a patient-owned reasoning graph that local AI agents can actually understand. The hospital's investment in FHIR is preserved. The patient can finally participate as an active and educated member of their own health team.

### Standing on the Shoulders of W3C: RDF and the Three-Layer Vocabulary

Instead of forcing applications to parse massive JSON or XML trees on the fly, Cascade Bridge stores every fact as a linked statement (subject, predicate, object) in RDF, the same W3C open standard that powers knowledge graphs at Google, Wikipedia, and Wikidata. RDF was designed for reasoning and built for interoperability; it is a mature, ratified W3C foundation. Every concept in the Cascade vocabulary lives at a permanent, resolvable URL with its formal OWL definition, so any developer or AI agent can follow a term like (`cascade:dataProvenance`) directly to its meaning. There are no PDFs to download, and no proprietary schemas to license.

The Cascade Protocol composes a three-layer vocabulary on top of RDF:

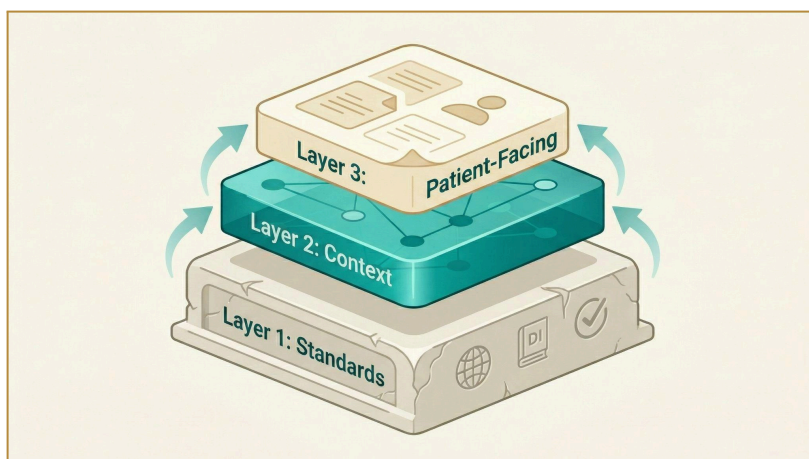
- **Layer 1, Established Standards.**

This layer imports raw, immutable clinical codes (SNOMED CT, LOINC, RxNorm, ICD-10) exactly as the source system exported them. Cascade never invents codes for concepts that standard terminologies already classify. Patient outputs share back in the same canonical codes that care teams already use.

- **Layer 2, Cascade Domain Context.**

This layer resolves conflicts and enforces provenance. Every raw code is wrapped with where it came from, when it arrived, and how trustworthy it is. Disparate data points are linked together (a medication, a lab, a symptom, a wearable reading) so the graph captures the patient's story, not just the institution's filing system.

- **Layer 3, Patient-Facing Translation.** This layer produces plain-language summaries, condition-specific dashboards, visit-prep templates, and AI-readable explanations. It turns opaque codes into something Priya can actually understand and act on, without ever severing the chain back to Layer 1.



The three-layer design is what enables a domain-specific app like *POTS Check to exist* on the same foundation that supports cancer surveillance, billing appeals, or genomic counseling. Each app inherits Layers 1 and 2 and authors Layer 3 vocabulary only when needed.

## Provenance as a First-Class Primitive

Consider how dangerous it would be for an AI agent to confuse a heart-rate reading from a Fitbit with one from an ICU monitor. In raw FHIR, the (Provenance) resource is optional and frequently omitted in real-world exports, leaving the AI to guess. In the Cascade Protocol, provenance is a non-negotiable, mathematical primitive. Every record must carry a (cascade:dataProvenance) tag or it fails SHACL validation. Apple Watch readings are (cascade:DeviceGenerated). EHR labs are (cascade:ClinicalGenerated). The patient's symptom diary is (cascade:SelfReported). An advisory from a clinical consortium is (cascade:ConsortiumAdvisory).

## Validation at Write-Time

Every record entering a Cascade pod is checked against a SHACL shape: required fields, allowed value ranges, mandatory provenance, cardinality, datatypes. A malformed or unattributed record cannot be silently added and discovered later. It is rejected at the boundary. The graph is therefore trustworthy by construction, not by hope.

Because OWL provides formal semantics on top of those shapes, we can write strict safety guardrails and have them enforced by the schema. Data is never destructively overwritten; both the patient and the AI can dynamically reason about the validity, recency, and weight of conflicting facts to establish a singular ground truth.

## Cross-Domain Reasoning

Consider an interaction that no single system catches today: a patient prescribed lisinopril (an ACE inhibitor) who self-reports taking an over-the-counter potassium supplement, with a recent EHR lab showing potassium at 5.1 mEq/L. Three sources, three silos, one severe hyperkalemia risk, invisible unless something stitches them together. In Cascade's graph this is natively computable:

Subject	Predicate	Object	Layer
Lisinopril Rx	clinical:rxNormCode	rxnorm:314076	1
Lisinopril Rx	clinical:therapeuticClass	"ACE Inhibitor"	2
Lisinopril Rx	cascade:dataProvenance	cascade:EHRVerified	2
K+ supplement	cascade:dataProvenance	cascade:SelfReported	2
K+ lab result	health:resultValue	"5.1 mEq/L"	1
K+ lab result	health:interpretation	"HIGH"	2
Visit prep summary	checkup:flaggedRisk	"K+/ACE interaction"	3

Our software traverses these connections directly. No SQL JOIN gymnastics, no per-application reconciliation logic. The 3D health graph in Figure 4 visualizes this exact interaction.

## Explainable AI: The Cure for Healthcare's Hallucination Problem

Where most health AI must infer relationships across disconnected FHIR resources, Cascade Bridge provides the model with a precise subgraph of attributed facts. When the on-device AI generates an alert (whether it is Priya's BRCA2 reclassification or the K+/ACE interaction above), It cites the source EHR for the prescription, the supplement diary entry, the lab result, and the consortium advisory. Every citation is a live URI pointing back to the source triple. A reviewing clinician can follow the chain back to the original FHIR or C-CDA record in one click. This is the explainability that healthcare demands and the type of AI-enabled interoperability solution that the ONC HTI-5 rule intends to advance.

## 6. Potential Impact

---

The Cures Act made health data available, but availability without usability is a dead end. Cascade Bridge's ultimate impact is to convert the legal right to access into the practical reality of data agency. The benefits compound across patients, providers, and developers.

### 1. Patients: Privacy, Engagement, and Data Agency

Today, patients are forced into a passive role when it comes to understanding and managing their health, overwhelmed by fragmented portals and an opaque view into their own data. Cascade Bridge turns passive data into an active asset:

- **AI assisted care.** AI can transform healthcare, but only if applications are grounded in secure, unambiguous patient data. The Cascade Protocol's decentralized architecture lets patients benefit from the AI revolution without surrendering their data to corporate interests.
- **Health literacy and informed decisions.** By translating opaque clinical codes into condition-organized, plain-language summaries, Cascade Bridge radically improves how patients engage with their own care. Priya's experience illustrates exactly this use case.
- **Reducing administrative friction.** Linked, provenance-tagged records paired with systems like SMART Health Check-in will help kill the clipboard and make it more efficient and accurate for patients to communicate with providers. Access to well organized medical facts will reduce unnecessary testing, and will finally make navigating billing disputes and insurance denials tractable for normal people.
- **Data agency.** Patients reclaim ownership of their data and decide how it can be used. For example, they can selectively share data with research studies and be compensated directly, bypassing third-party data brokers entirely.

### 2. Providers: Accelerating Intake and Reducing Noise

Providers are overwhelmed by demand and drowning in erroneous data. Cascade Bridge reverses this burden with canonical, structured pre-visit summaries that focus limited provider time on the highest-value clinical decisions. When patients are prepared, visits run more smoothly. For example, POTS Check lets patients use their Apple Watch to take NASA Lean tests at home and share structured longitudinal results formatted for provider review, meaningfully accelerating the famously difficult dysautonomia diagnosis. When Priya arrives for her genomic counseling appointment she is armed with pertinent questions, and her provider is already briefed with the latest relevant variant interpretation.

### 3. Developers: The Compounding Impact of an Open Ecosystem

The *Cascade Protocol* is not a walled garden. A single app solves a single problem; an open infrastructure solves an industry-wide bottleneck. Phase 2 will deliver Cascade Bridge v1.0, the open Cascade Protocol SDKs, and three condition-specific reference applications within nine months, alongside a production SMART Health Check-in wallet and conformance tooling for third-party developers. Each new application built on the protocol inherits the full *Collect-Connect-Control* infrastructure, so the architecture grows organically as new domains, EHR integrations, and patient use cases stress-test it. The result is a compounding, patient-owned data ecosystem that will accelerate the development of hundreds of specialized health tools, and that gives every American the practical ability to use the data that, by law, is already theirs.

A standard built for the whole industry has to be shaped by the whole industry. Recognition under EHIgnite would do more than fund engineering: it would convene the conversation with ONC, USCDI maintainers, clinical specialty societies, EHR vendors, and patient advocacy organizations whose perspectives this protocol needs to incorporate as it matures. Rewiring healthcare from the outside is hard; the most valuable thing this challenge offers is a seat at the table where the future of patient-owned health data is being shaped.