KITZU

Wellness Intelligence Platform

Architecture & Implementation Specification

Version 0.2 | April 2026

A local-first, privacy-first intelligence layer that unifies genetic, blood, microbiome, wearable, and glucose data into a single adaptive wellness protocol — powered by AI, measured by objective data, and delivered through human coaching.

Founded by
Shelley Syrros | Founder & Visionary

Platform built by
Bill Syrros | Technology & AI
with Claude (AI Technical Partner)

Contents

1. The Problem

The wellness industry is fragmented by design. A person today can wear a smart ring that tracks their sleep, draw blood every quarter to measure 50+ biomarkers, sequence their DNA, map their gut microbiome, and strap on a continuous glucose monitor — yet none of these systems talk to each other. Each platform operates in isolation, generating recommendations based on a single data source while ignoring everything else that's happening in that person's body.

The blood panel says "take more Vitamin D" without knowing the person carries a genetic variant that makes their receptors inefficient. The microbiome test says "eat these superfoods" without knowing those same foods spike the person's glucose. The wearable says "you're recovered" without knowing the liver enzymes are trending in the wrong direction. The genetic report says "variant detected" without connecting it to three years of abnormal blood markers.

The result: conflicting advice from five different dashboards, no unified protocol, and no feedback loop to confirm whether any of it is actually working. People invest thousands of dollars in testing and devices, receive fragmented guidance, and default to the same generalized wellness advice they could get for free.

Meanwhile, the most personal data a human being possesses — their DNA, their blood chemistry, their gut biology — sits on corporate servers across five different companies, each with their own privacy policies and breach exposure.

2. The Vision

A local-first, privacy-first intelligence layer that unifies all personal health data sources, cross-references them through AI, and produces a single adaptive wellness protocol — updated daily from wearable signals, recalibrated quarterly from bloodwork, and grounded in genetics that never change.

Combined with measurable physical wellness modalities and human coaching, the platform closes the loop between prescription and follow-through — solving the compliance problem that renders most wellness programs ineffective.

No data leaves the client's device. No corporate cloud stores their DNA alongside millions of others. The analysis happens locally, the protocol adapts continuously, and every intervention is objectively measured.

Principle	What It Means
Local-First Privacy	All data lives on the client's own device. No cloud health database. No data custody risk.
AI Cross-Referencing	Claude analyzes genetics + blood + microbiome + wearable data in a single session — connections no single app can make.
Adaptive Protocols	Health plans update daily from wearable data, not just quarterly from blood draws.
Measurable Modalities	Every wellness intervention (sauna, cold plunge, meditation) is tracked and its impact proven through objective data.
Passive Compliance	The wearable data tells the truth — no self-reporting needed. Coaches see what actually happened.
Human Coaching	Shelley delivers the high-value human moments. AI handles the 90% that can be automated.

3. Data Sources

Tier 1: Foundation (static or infrequent)

Source	Data Type	Frequency
Genetic Testing (e.g. 23andMe)	Raw SNP genotypes (~600K markers)	Once (genetics don't change)
Microbiome (e.g. Viome)	Gut microbiome analysis, food recommendations	1-2x / year
Blood Biomarkers (e.g. InsideTracker)	50+ biomarkers, biological age, food recs	Every 3-6 months

Tier 2: Dynamic (daily/weekly)

Source	Data Type	Frequency
Wearable (e.g. Oura Ring)	Sleep, HRV, readiness, heart rate, activity	Daily (API)
Vitals (e.g. iHealth)	Blood pressure, weight, body composition	Daily-weekly

Tier 3: Recommended Addition

Source	Data Type	Frequency
CGM (Freestyle Libre 3)	Continuous glucose readings (1-min intervals)	2-week stints, quarterly

4. Architecture

The platform is organized into five layers, each with a distinct responsibility. Data flows upward from collection through analysis to adaptive output.

Layer 1: Data Lake — *Local Storage*

All data lives on the client's machine. No cloud. No third-party servers. Every record has a source, timestamp, and confidence field. Units are normalized across all sources. Historical data is append-only.

Layer 2: Ingestion Pipeline — *Data Collection*

Six independent collectors, each on its own schedule. Genetics parsed once. Microbiome after each test. Blood after each draw. Oura daily via API. Vitals daily. CGM during sensor stints. No external dependencies — uses scheduled tasks only.

Layer 3: Inference Engine — *AI Cross-Referencing*

Claude reads the unified data lake and produces cross-referenced insights no single platform can generate. Genetics + blood reveals receptor inefficiencies. Blood + wearable confirms overtraining vs. stress. Genetics + blood + microbiome tells the full iron or cholesterol story. Five analysis modes: daily brief, weekly review, blood draw analysis, ad hoc queries, and protocol generation.

Layer 4: Protocol Engine — *Adaptive Output*

The health plan becomes a living system. It adapts daily from wearable signals — readiness drops trigger training reduction, HRV trends flag cortisol risk, sleep deficits activate recovery protocols. Includes passive compliance detection: the data reveals whether the client followed the protocol without requiring self-reporting.

Layer 5: Presentation — *Dashboard + Conversational*

Two interfaces: a visual dashboard showing all metrics, trends, and protocol status (accessible from any device via secure network), and conversational AI sessions where clients can ask any question against their full health dataset.

5. Implementation Plan

Phase 1: Foundation (Week 1-2) — "Wire the Data"

All data sources parsed and stored in the unified data lake. Collectors built for genetics, blood, microbiome, wearable, and vitals. Unified schema designed. Daily Oura pull scheduled.

Phase 2: Intelligence (Week 2-3) — "Connect the Dots"

Cross-referencing engine operational. SNP annotation engine maps genetic variants to health implications. Cross-reference prompt templates built for all source combinations. Daily brief and weekly review automated.

Phase 3: Protocol Engine (Week 3-4) — "Make It Adaptive"

Living protocol that updates daily from incoming data. Adaptation rules engine processes wearable signals and adjusts recommendations. Compliance detection active. Dashboard auto-updates.

Phase 4: CGM Integration (Month 2) — "Close the Glucose Loop"

Freestyle Libre 3 via doctor prescription. Glucose spikes cross-referenced with microbiome food list and genetic insulin sensitivity markers. Personal glucose response map generated.

6. Wellness Modalities (13 SKUs)

The platform doesn't just measure — it prescribes and tracks physical wellness interventions. Each modality is an individual SKU that clients can add to their program. Every modality's effectiveness is measurable through the data platform.

Physical Wellness Modalities

SKU	Modality	How We Measure Impact
KZ-RLT	Red Light Therapy	HRV recovery delta, hsCRP trend, sleep quality post-session
KZ-BT	BrainTap	HRV during/after, sleep architecture changes, cortisol trend
KZ-O2	Oxygen Therapy (ENG3)	Readiness score delta, exercise recovery time, HRV baseline
KZ-SAU	Dry Sauna	HR during session, HRV recovery curve, ferritin response, sleep onset
KZ-CP	Cold Plunge	HRV spike post-session, next-day readiness, hsCRP trend, glucose response
KZ-MT	Music Therapy	HRV during/after, sleep score comparison, cortisol trend
KZ-HM	HeartMath	Coherence scores, HRV correlation, cortisol and testosterone response

Growth & Development Modalities

SKU	Modality	How We Measure Impact
KZ-MED	Meditation	HRV during/after, sleep onset latency, cortisol trend, coherence scores
KZ-YGA	Yoga	HRV recovery, readiness trends, cortisol/inflammation markers, sleep quality
KZ-PC	Personal Consulting (Shelley)	Protocol adherence, biomarker trajectory, InnerAge trend, wellbeing
KZ-BC	Breathwork	HRV real-time response, SpO2, sleep architecture changes, cortisol
KZ-JRN	Journaling / Reflection	Correlation with sleep/readiness scores, longitudinal mood tracking
KZ-NUT	Nutrition Coaching	Bloodwork trajectory, CGM response, weight/comp trends, energy scores

Modality Stacking Protocols

Stack	Sequence	Measured By
Recovery	Cold Plunge > Red Light > BrainTap	Next-day Oura readiness vs. training-only days
Stress Reset	HeartMath > Sauna > Music Therapy	7-day HRV trend reversal, cortisol at next blood draw
Sleep Optimization	Red Light (evening) > BrainTap > Music Therapy	Deep sleep % and sleep efficiency over 14 days

Stack	Sequence	Measured By
Metabolic	Sauna > Cold Plunge > fasted walk (next AM)	CGM glucose baseline shift, weight trend
Mind-Body	Yoga > Breathwork > Meditation > Journaling	HRV baseline over 30 days, cortisol, sleep quality
Onboarding	Personal Consulting > Meditation > Nutrition > 1 recovery modality	Protocol adherence, first Oura trends, engagement

7. The Compliance Problem & Solution

The traditional wellness model is broken at the follow-through layer. A physiotherapist gives you exercises, a nutritionist gives you a meal plan, a trainer gives you a program — then you go home and nobody knows if you did any of it. When you show up 6 weeks later with the same problem, they give you another sheet and the cycle repeats. There's no feedback loop, no objective measurement of compliance, and no way to distinguish between "the protocol isn't working" and "the client isn't following the protocol."

Kitzu eliminates this entirely through passive compliance detection. The wearable data tells the truth without the client having to self-report anything.

The Data Doesn't Lie

- Cold plunge sessions show up as HRV spikes at specific times of day
- Red light therapy evenings correlate with measurable deep sleep increases
- Meditation and breathwork sessions show up as HRV coherence patterns
- Exercise compliance shows up in activity scores, training load, and calorie burn
- Nutrition compliance shows up in CGM glucose patterns and weight trends
- Sleep protocol compliance shows up in sleep onset time, efficiency, and regularity

What This Enables for Shelley

Shelley walks into a quarterly session already knowing what the client actually did versus what was prescribed. She can show objective proof that the protocol works when followed — "the weeks you did your recovery stack, your readiness averaged 82; the weeks you skipped, it averaged 64." She can distinguish between non-compliance and protocol failure — two completely different conversations. And she can identify partial compliance patterns that explain why results are mixed.

The key principle: Kitzu doesn't rely on willpower or self-reporting. The feedback loop is built into the data layer. If a client stops doing the work, the platform detects it within days — not months — and Shelley can intervene with a supportive conversation before the client falls off entirely.

8. Client-Facing Platform

Client Tiers

Tier	Name	What's Included	Monthly (CAD)
1	Foundation	Wearable data analysis + daily briefs + basic protocol	$149
2	Optimize	Foundation + blood biomarker analysis + quarterly protocol updates	$299
3	Precision	Optimize + genetic analysis + microbiome + CGM + full cross-referencing	$499
4	Transform	Precision + modality access (SKU menu) + stacking protocols + weekly coaching	$749

Monthly pricing covers platform access, AI analysis, and consulting. Third-party test kits and devices are additional costs borne by the client.

Client Investment: Beyond Kitzu Fees

Item	Approx. Cost (CAD)	Frequency
Oura Ring (Gen 3)	$400–450	One-time
Oura membership	$8/mo	Monthly
InsideTracker blood test	$250–600	Quarterly (Tier 2+)
23andMe Health + Ancestry	$250	One-time (Tier 3+)
Viome Gut Intelligence	$250–400	Annual (Tier 3+)
Freestyle Libre 3 sensor	$50–90/sensor	Per stint (Tier 3+)

Total first-year investment for a Precision client: ~$10,000–13,000 CAD (Kitzu fees + tests + devices). This positions well below concierge medicine ($25K–100K+) while delivering comparable depth of insight.

Client Onboarding: Zero-Friction, Privacy-First

Step 1: Discovery Call *(30 min, free)*

Understand goals, inventory existing devices and data, assess comfort level with data sharing, recommend starting tier.

Step 2: Device Setup *(Self-service, guided)*

Step-by-step guide walks client through API tokens, test ordering, and data export. Client runs a local setup on their own device. No data uploads to Kitzu servers. Ever.

Step 3: First Analysis Session *(60-90 min, $500 onboarding fee)*

Shelley joins a session with the client. AI reads the local data lake and produces the initial cross-referenced health report. Client sees their unified profile for the first time. First protocol generated. This is the wow moment.

Step 4: Ongoing Monitoring *(Automated + quarterly)*

Daily automated briefs from wearable data. Weekly trend summaries. Monthly modality reviews. Quarterly deep-dive with Shelley after each blood draw — the highest-value touchpoint.

Step 5: Modality Integration *(Transform tier)*

Client selects modalities from the SKU menu. Some are home-based (red light, cold plunge, BrainTap). Some are facility-based (sauna, oxygen). Every session logged and measured. Protocol adapts based on response.

9. Vendor Resilience

The platform must never depend on any single vendor. The architecture is built on data formats, not brand names. Every data source slot uses an adapter module — swap one vendor for another by writing a new adapter. No changes to the inference engine or protocol engine.

Genetics Alternatives

Vendor	Strengths	Notes
23andMe	650K+ SNPs, health reports	Post-bankruptcy. Download data while available.
AncestryDNA	700K+ SNPs, largest genealogy DB	Less health detail than 23andMe
Nebula Genomics	30x whole genome (~6B data points)	Future-proof. Prices approaching $200.
SelfDecode	SNP + AI health reports	Largest report library

Blood Biomarker Alternatives

Vendor	Strengths	Notes
InsideTracker	InnerAge 2.0, food recs, 43-50 markers	Best actionable guidance
Function Health	160+ markers, $365/yr	More markers, less interpretation
Outlive.bio / OneTwenty	Blood + wearable integration	Closest to Kitzu's model. Watch this.
Tally Health / SystemAge	Epigenetic (DNA methylation) age	Complementary to blood-based bio age

Microbiome, Wearable & Vitals Alternatives

Microbiome: Viome (RNA, best food recs), Thorne (metagenomics), ZOE (bundled approach — competitor to watch), Ombre (budget), Biomesight (best data sovereignty). Wearables: Oura (best sleep API), Ultrahuman Ring (no subscription), WHOOP (best training load), Apple Watch (most sensors), Garmin (best for athletes). Vitals: iHealth (current), Withings (best API, Lingo CGM integration), Omron (most clinically validated BP). CGM (Canada): Freestyle Libre 3 via prescription. Dexcom G7 as alternative.

> *Review cadence: quarterly vendor assessment tracking new entrants, data export changes, pricing shifts, Canada availability, API stability, and privacy policy changes.*

10. Privacy & Security Architecture

- **Local-first:** All data stored on the client's machine. No cloud health database.
- **No outbound transmission:** APIs pull data IN. No personal health data goes OUT.
- **Ephemeral analysis:** AI processes data in session, doesn't persist it.
- **Encrypted at rest:** Device-level encryption covers the disk.
- **No third-party uploads:** No uploading raw DNA to external analysis sites.
- **Secure access only:** Dashboard accessible only via client's private network.
- **Client isolation:** Each client's data is physically separate. No shared database.

The hard line: client health data and Kitzu infrastructure never mix. Kitzu provides the methodology, the modalities, and the expertise — not data storage. This eliminates every privacy objection before it's raised.

11. Team Model

Role	Person	Responsibilities
Founder / Visionary	Shelley	Kitzu's creator and driving force. All client-facing wellness consulting: onboarding, quarterly deep-dives, personal coaching, protocol interpretation, accountability
Technology & Platform	Bill	Builds the technical platform with Claude/Luke. AI/data integration, architecture, conference speaking, business development
AI Technical Partner	Claude / Luke	Data analysis at scale, cross-referencing engine, daily/weekly automated briefs, protocol generation

Shelley is the founder, visionary, and face of Kitzu. She created the concept, leads the wellness methodology, and delivers the client experience. Bill pulls the technical side together — building the data platform, wiring the integrations, and working with Claude/Luke to make the AI layer real. This model lets Kitzu serve many clients without Shelley being the bottleneck — automated daily/weekly briefs handle 90% of touchpoints, freeing Shelley for the high-value human moments.

12. Competitive Positioning

Competitor	What They Do	Kitzu Differentiator
InsideTracker	Blood analysis + recs	Single-source. No genetics, gut, or wearable. Cloud-stored.
Viome	Microbiome + food recs	Single-source. No blood or genetic cross-reference. Cloud-stored.
Levels / Lingo	Glucose monitoring	Single-source. No genetic or microbiome context. Cloud-stored.
Function Health	100+ biomarker testing	Data without interpretation depth. No wearable integration.
Peter Attia's Practice	Elite longevity medicine	$100K+/year. Not accessible. Not scalable.
ZOE	Microbiome + blood + CGM	Closest model. But cloud-based, no genetics, no modalities.
Kitzu	All sources unified + modalities + local privacy	Only platform that cross-references 5+ sources locally, measures modalities objectively, and never stores client data externally.

13. Revenue Model & Go-To-Market

Stream	Type	Ballpark (CAD)
Onboarding + genetic analysis setup	One-time	$500 per client
Monthly platform subscription	Recurring	$149–749/mo (by tier)
Quarterly blood draw analysis + protocol	Included in Optimize+	—
Modality sessions (facility-based SKUs)	Per session or package	$75–150/session
Modality equipment referrals (home-based)	Commission / markup	15–25% margin
Ad-hoc consulting (Shelley)	Hourly	$200/hr
Corporate wellness programs	Contract	$5K–15K/yr per company

Annual revenue per client (blended estimate): $4,000–12,000 CAD depending on tier and modality usage.

Go-To-Market: Proof of Concept First

Bill and Shelley are the first pilot clients. Their personal health journeys — tracked end-to-end through the platform — serve as the live proof of concept. Conference presentations show a live dashboard with real data, walk through genetic insights, demonstrate a modality stack, and show that morning's wearable signals. That's the kind of demonstration that converts audience members into clients — and Shelley converts interest into long-term relationships.

14. Success Metrics

Platform Milestones

Milestone	Target
All data sources ingested and cross-referenced	30 days
Daily automated brief running	2 weeks
First fully adaptive protocol generated	1 month
Modality tracking + response profiling operational	6 weeks
First client onboarding (pilot)	3 months
Conference-ready live demo	4 months
Repeatable onboarding flow (under 1 hour setup)	4 months
3 pilot clients through full onboarding	6 months
At least 2 modality partnerships (facility access)	6 months
Corporate wellness pilot proposal ready	6 months

This document is a living specification. It will be updated as the platform evolves, new data sources become available, and client feedback shapes the offering.

KITZU | Wellness Intelligence Platform | March 2026

Shelley Syrros (Founder & Visionary) | Bill Syrros (Technology & AI) | Claude (AI Partner)

15. Platform Update — v0.2

April 2026 · Built by Bill Syrros with Claude (AI Technical Partner)

This section documents changes made to the Kitzu platform in April 2026. The core architecture described in sections 1–14 remains unchanged. Updates below reflect new dashboard capabilities, inference engine additions, and two bug fixes resolved during active development.

Version History

Version	Date	Author	Summary
v0.1	March 2026	Bill Syrros	Initial architecture specification
v0.2	April 2026	Bill Syrros + Claude	Exercise feed, supplement schedule card, bug fixes, README

Summary of Changes

Area	Change	Type
Inference Engine (Layer 3)	Daily brief now generates supplement_schedule — a structured, data-driven supplement list with timing, dose, ar	NEW
Presentation Layer (Layer 5)	Two new full-width dashboard cards: Exercise Feed (7-day workout history from Oura) and Supplement Schedul	NEW
Ingestion Pipeline (Layer 2)	Exercise/workout data now correctly stored in data lake. HRV trend now reads from full data lake (14 days) rathe	FIX
Technical Docs	README.md created at kitzu/ root covering architecture, data schema, known issues, bug history, and all CLI co	NEW

Updates to Section 4: Architecture

Layer 3: Inference Engine — AI Cross-Referencing

The daily brief engine now produces a sixth output field: **supplement_schedule**. This is a fully structured, priority-ordered list of supplements with timing slots, doses, and data-driven priority tiers computed each morning from the unified profile.

Priority	Meaning	Example Trigger
1 — Must-have	Directly targets a critical or at-risk marker	LDL > 130 → Omega-3 = Priority 1
2 — High value	Strong supporting role regardless of marker status	CoQ10, B-Complex, Copper
3 — Good to have	Maintenance / longevity; drop first when travelling	Resveratrol, Probiotic

Layer 5: Presentation — Dashboard Cards (Updated)

The Kitzu dashboard at v0.2 renders 12 cards in the following order:

#	Card	Data Source
1	Header (date, readiness ring)	brief.date, readiness section
2	Today's Priorities	brief.priority_actions
3	Sleep	brief.sections → Sleep
4	Readiness & Activity	brief.sections → Readiness + Activity
5	HRV Trend (14-day bar chart)	profile.oura.hrv_trend
6	Exercise Feed ★ NEW	profile.oura.exercise_feed (last 14 days)
7	Supplement Schedule ★ NEW	brief.supplement_schedule
8	Biomarker Watch	brief.sections → Biomarker Watch
9	Body Composition	brief.sections → Body Composition
10	Blood Pressure	profile.vitals.blood_pressure
11	Cross-Reference Insights	brief.cross_references
12	Data Freshness	brief.data_freshness

Updates to Section 5: Implementation Plan

Phase 1 and Phase 2 are complete as of April 2026. Phase 3 protocol adaptation rules are partially operational through the priority actions and supplement schedule engines.

Phase	Name	Target	Status
1	"Wire the Data" Oura, blood, iHealth, Viome collectors + unified schema + daily scheduled pull	Week 1–2	✓ Complete
2	"Connect the Dots" Daily brief engine, HRV trend, exercise feed, supplement schedule, cross-reference insights	Week 2–3	✓ Complete
3	"Make It Adaptive" Full adaptation rules engine, compliance detection, dashboard auto-updates	Week 3–4	In Progress
4	"Close the Glucose Loop" Freestyle Libre 3 CGM integration	Month 2	Planned

Bug Fixes Resolved

Bug	Root Cause	Fix
HRV trend chart showed only 2 days instead of 14	hrv_trend was built from records collected [illegible]	Changed to load_[illegible]
Exercise feed always empty despite [illegible]	The sessions endpoint 404 was inside a single try/except block [illegible]	Split into independent [illegible]

New Data Fields — Unified Profile & Brief Schema

Field	Location	Description
exercise_feed	profile.oura.exercise_feed	List of up to 30 workout sessions from the last 14 days. Each entry: date, type, t
supplement_schedule	brief.supplement_schedule	Ordered list of supplement entries. Each entry: time_block, time, name, dose, pr

This document is a living specification. It will be updated as the platform evolves, new data sources become available, and client feedback shapes the offering.

KITZU | Wellness Intelligence Platform | April 2026
Shelley Syrros (Founder & Visionary) | Bill Syrros (Technology & AI) | Claude (AI Technical Partner)