Lifehouse Design Tech

A summary of the technology we are building | Working draft, system in testing | October 2026

What we are building

Lifehouse is building shared community tech so a neighbourhood can see itself: its people, households, skills, needs, offers, resources and relationships, and who is responsible for what. Much of this is already known locally, but it is scattered across conversations, spreadsheets, inboxes and people's memories. The system brings it together in one place, with consent, accountability and human review built in. The Workspace carries the line "Making the community legible to itself."

It supports the Lifehouse Design being piloted in Stroud and is intended to be reused elsewhere. As Buckminster Fuller put it, to change something, build a new model that makes the existing one obsolete.

How it came about

We expected to choose a few open-source tools, set them up and hand them over. It soon became clear this was not a software problem. Before the tools could be configured, we had to settle what counts as a skill and what counts as an offer, how a circle differs from a group, and who may see what. So the technology and the governance have been designed together, with over 200 hours of work estimated so far.

Abstract language also does not work in a database. Ideas like "interconnectedness" and "resource flow" had to become practical questions: who knows whom, who is isolated, who has tools, skills or space, and what permission exists to connect people. This gives the ideas handles rather than diluting them, and keeps our language open to the people we are mapping, who are already part of a commons.

How it works: one connected flow

- **Capture.** Lifehouse keepers and fieldworkers record what people tell them using structured forms. AI can propose structure from notes, but it never decides what is true.
- **Staging.** Every submission lands in a holding area first. Nothing becomes an official record yet.
- **Human review.** A reviewer sees the source, the proposed values and the consent information, then corrects, verifies or rejects. This is where a person takes responsibility.
- **The canonical record.** Verified information is published into a governed PostgreSQL data model covering people, households, places, needs, offers, skills, resources, relationships, circles, roles, tasks, decisions, consent and a full audit history. The value lies in how these connect. A skill is not automatically an offer, and a concern is not automatically a need.
- **The Workspace.** The everyday interface, built around Lifehouse concepts rather than a conventional CRM. Twenty, the CRM tool, sits underneath where useful but does not define what our information means.
- **Infoscape.** Our knowledge layer. It answers questions from an approved document library, shows its sources, prefers current documents over older ones, and says so when the evidence does not support an answer. It cannot create or change records.

The Workspace

The home screen shows the review pipeline (staged, under review, verified), 14 domains from People and Households to Circles, Governance, Capacity and Events, and panels for open tasks, saved searches and proposals. Proposals are suggested matches between needs and capabilities. They are always human-reviewed and never an assignment or obligation.

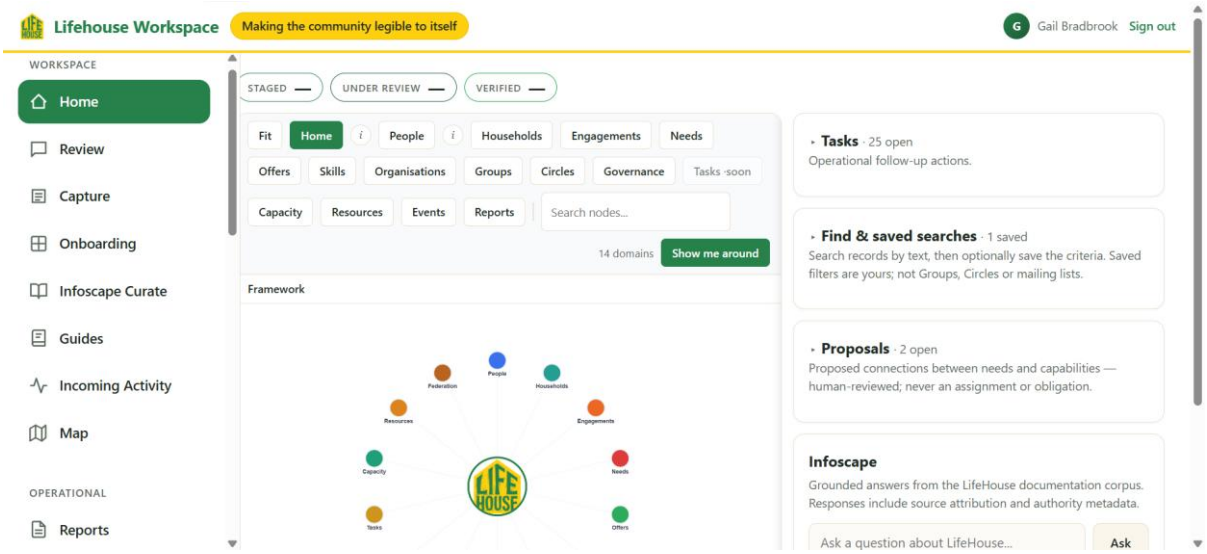

The Lifehouse Workspace home screen (test environment)

The Fieldwork Map lets you move from the whole area down to individual streets and households. Each property is marked as needing approach, under review or processed, so door-knocking can be planned and followed. Sensitive location information stays subject to access rules.

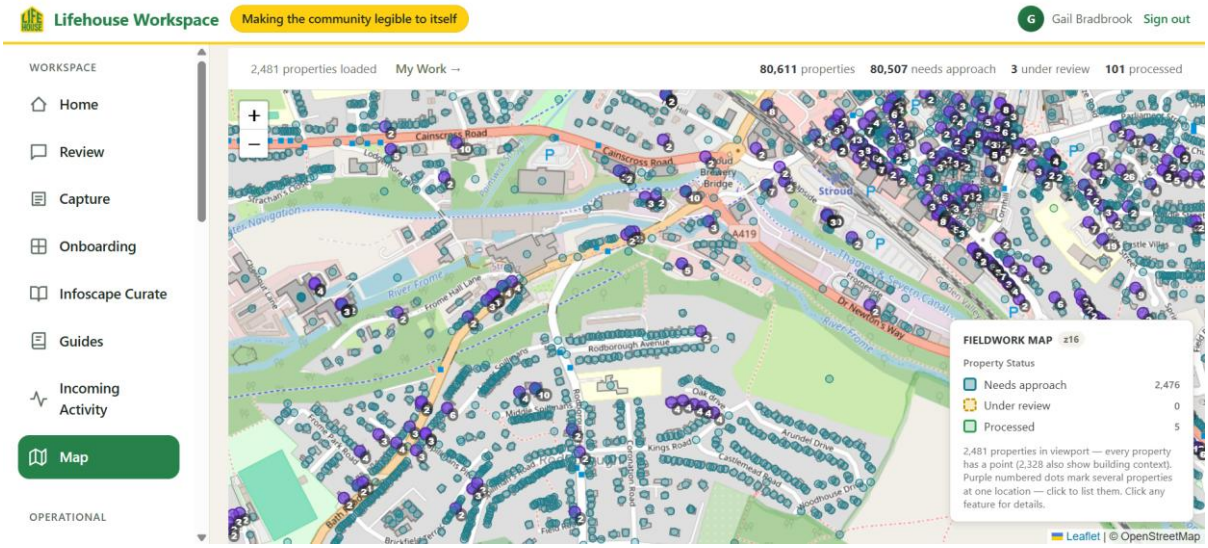

The Fieldwork Map: street and household level detail (test environment)

Principles that run through it

- **A data commons is not open access.** Information is shared for community purposes under clear rules. Shared value does not mean everyone sees everything.
- **Authority follows responsibility.** Access runs from circle to role to assignment to person, shaped by consent and sensitivity. People hold roles, roles sit in circles, and when someone leaves the role and its history remain.
- **Consent is recorded by purpose** (contact, photos, public profile, research and so on). History is kept and consent is never inferred from friendliness or participation.
- **UK GDPR is built into the structure** through provenance, audit trails, restricted access and support for access, portability and erasure requests. Judgements about lawful basis and retention stay with the organisation.
- **The map is not the territory.** The database is an aid to relationships, never a replacement for listening, care or action.

Built to be copied

The platform is designed as self-hosted, reusable community infrastructure. The core (data model, review process, consent and governance structures, knowledge layer, deployment procedures) stays stable. Each community sets up its own circles, mandates, people, branding and policies, and nothing is copied from Stroud's data. We will not call it replication-ready until an operator who has not worked on it can build a clean instance from the documentation alone. Until then it is designed for reuse and being made ready.

Where we are

The system is currently in testing and is not yet live. This stage is about showing the main journey from capture to governed use working end to end. Role-specific views, deeper refinement and further integrations are still to come, and the governance interpretation remains open to stakeholder review. People still decide the organisation's values, mandates and referral routes.