



# Evidence Core 实证核心

The evidence and reporting layer for community organizations: capture the proof of your work with near-zero friction, let AI do the structuring, keep humans in charge of every decision — and generate funder-ready reports from one trusted store.

Audience: executive directors, boards, program directors, funders    Reading time: ~8 minutes

## 1. The problem every nonprofit already knows

Community organizations do enormous work and document almost none of it — not for lack of will, but because documentation competes with the work itself. The proof exists: photos on staff phones, WeChat and text messages, voice notes, sign-in sheets, spreadsheets, forwarded emails, last year's annual report. It is scattered, unlabelled, and mostly forgotten by the time a grant report or newsletter is due.

The result is a familiar cycle: a frantic quarterly scramble to reconstruct numbers, stories that go untold, and reports built on memory rather than evidence.

**Why existing tools have not fixed this.** CRMs, case-management systems and impact platforms all begin with the same demand: *fill in this form, configure this database, run this survey*. They ask busy staff to structure information at the moment they are least able to. Adoption fails, the database goes stale, and the organization pays for software that describes what it wished it had recorded.

**Evidence Core inverts the order.** The system accepts chaos as input; structuring is the system's job; humans only confirm.

## 2. What Evidence Core is — in one picture




Figure 1 — The four-stage loop. People spend their time only in stages 1 and 3; the system owns stages 2 and 4.

|                                                                    |                                                             |                                                                                |                                                 |
|--------------------------------------------------------------------|-------------------------------------------------------------|--------------------------------------------------------------------------------|-------------------------------------------------|
| <b>&lt; 2 min</b><br>to capture an event from a phone in the field | <b>&lt; 1 min</b><br>for a supervisor to confirm the drafts | <b>≈ \$0.04</b><br>AI cost per capture; hard daily budget cap per organization | <b>0</b><br>data-entry forms staff must fill in |
|--------------------------------------------------------------------|-------------------------------------------------------------|--------------------------------------------------------------------------------|-------------------------------------------------|

### 3. What is genuinely different

#### Capture-first, not form-first

Staff use the tools already in their hands — a photo, a 20-second voice memo, a forwarded email, a text message. No login is required to contribute by email or SMS; senders are verified once and admitted by an ingestion gate. Sensitive channels stay silent to strangers.

#### AI proposes, people decide

The extraction model has no ability to act. It reads content as untrusted data, returns structured proposals, and every proposal is a draft until a person confirms it. Instructions hidden inside a document are quoted as evidence, never followed. Correction rate is measured so the system's usefulness is visible week over week.

#### Consent is architecture, not policy

Every item carries a scope — **internal** → **funder reports** → **public**. Widening scope is an admin-only decision. Photos with faces default to internal; location metadata is erased on ingest. Reports can only draw on evidence cleared for their audience, and this is enforced by the system, not by an editor remembering.

#### Provenance end to end

Every number in a report traces to the item it came from, the person who confirmed it, and the original message, photo or page. Revoking consent flags every report that used the item.

#### Bilingual by design

Chinese and English are first-class across programs, people, stories and reports — no translation add-on, no duplicate records when a person appears under two names (matches attach; only a human merges).

#### Your data, fully

A one-click export delivers the entire organization — every record and every media file — in open formats. No lock-in.

**Positioning.** Evidence Core is not a CRM, not a case-manager, not a survey tool. It is the **evidence and reporting layer** that sits underneath all of them: low-friction capture + human-governed AI + provenance + outcome linkage + defensible reporting.

### 4. Consent and security — the two chokepoints



Figure 2 — Consent scope is enforced at exactly two places every piece of evidence must pass through. A public newsletter can physically only see public items; a funder report adds funder-cleared items; internal reports see everything.

Under the hood: each organization's data is isolated at the database level (row-level security), administrators use two-factor sign-in, all AI input is fenced as data, media is stored privately with expiring links, backups run nightly with a monthly restore drill, and every confirm, scope change, export and merge is written to an audit log.

## 5. A day in the life — workflow scenario




Figure 3 — Total human effort for one event: about three minutes, spread across two people, on the devices they already carry.

## 6. Report generation — from evidence to deliverable

Reporting is where the loop pays off. An administrator picks a format, a period and, optionally, a program or specific evidence items. The system retrieves only *confirmed* evidence the audience is entitled to see, drafts the document, and lists any gaps it could not fill with evidence — so the writer knows exactly what to add rather than what to invent.

| Format                                      | Audience | Draws on                |
|---------------------------------------------|----------|-------------------------|
| Newsletter · flyer · promo                  | Public   | public items only       |
| Monthly / annual report · grant application | Funder   | public + funder-cleared |
| Event report · internal review              | Internal | everything confirmed    |

- **Evidence-only rule.** Drafts cite what was confirmed; unsupported claims are not generated. Editors' notes flag gaps.
- **Traceability.** Every generated draft records which items it used; if consent is later revoked, affected drafts are flagged.
- **Outcome board.** Mission and goals feed an AI consultation that proposes an outcome ladder (output → quality → outcome → impact) with indicators bound to metrics you already collect. Live observations roll up automatically once confirmed.
- **Cold-start in a day.** Past annual reports, yearbooks and newsletter archives are mined into draft evidence, so year one starts with history, not an empty database.

## 7. How it compares

| Capability                                               | Evidence Core | Impact platforms (Sopact, Clear Impact, UpMetrics) | Nonprofit CRM / case mgmt (Salesforce NPC, Bloomerang, Apricot) | Story-coding tools (ImpactMapper) | DIY (Drive, Sheets, WeChat, Airtable) |
|----------------------------------------------------------|---------------|----------------------------------------------------|-----------------------------------------------------------------|-----------------------------------|---------------------------------------|
| Starts from capture (photo, voice, email, SMS), no forms | ●             | forms / surveys                                    | data entry                                                      | upload & code                     | 🕒 capture, no structure               |
| AI structures raw input; human confirms every item       | ●             | 🕒 analysis after collection                        | —                                                               | 🕒 assisted coding                 | —                                     |
| Consent scope enforced by the system, per item           | ●             | —                                                  | 🕒 record permissions                                            | —                                 | —                                     |
| Provenance from report line back to source               | ●             | 🕒                                                  | 🕒                                                               | ●                                 | —                                     |

| Capability                                          | Evidence Core  | Impact platforms<br>(Sopact, Clear Impact, UpMetrics) | Nonprofit CRM / case mgmt<br>(Salesforce NPC, Bloomerang, Apricot) | Story-coding tools<br>(ImpactMapper) | DIY<br>(Drive, Sheets, WeChat, Airtable) |
|-----------------------------------------------------|----------------|-------------------------------------------------------|--------------------------------------------------------------------|--------------------------------------|------------------------------------------|
| Bilingual (zh/en) records and reports               | ●              | —                                                     | —                                                                  | —                                    | 🕒 manual                                 |
| Generates newsletters, reports, grant drafts        | ●              | 🕒 dashboards                                          | 🕒 templates                                                        | —                                    | —                                        |
| Outcome ladder + indicators linked to live evidence | ●              | ●                                                     | 🕒                                                                  | 🕒                                    | —                                        |
| Historical documents mined into evidence            | ●              | —                                                     | —                                                                  | 🕒                                    | —                                        |
| Full data export, no lock-in                        | ●              | 🕒                                                     | 🕒                                                                  | 🕒                                    | ●                                        |
| Indicative annual cost, small org                   | ~\$1.2k – 2.4k | \$5k – 30k                                            | \$3k – 15k + setup                                                 | ~\$5.5k                              | “free” + staff hours                     |
| Time-to-first-value                                 | days           | weeks – months                                        | months                                                             | weeks                                | immediate, then decays                   |

● full · 🕒 partial · — absent. Competitor characterizations and prices are indicative, based on public materials and 2026 market research; they describe typical configurations, not every plan. The comparison is by design of approach, not feature count: none of the alternatives begin from low-friction capture, which is where nonprofit data actually fails.

## 8. Where it stands and what comes next

**Today.** Evidence Core is live at cyphia.app serving two organizations (Mental Health Association for Chinese Communities and Kindness Network), with all capture channels, the confirmation queue, evidence library, document and newsletter mining, staff assignments calendar, reporting drafts and the outcome/impact board in production. Nightly backups and monthly restore drills are in force. Data is isolated per organization; adding an organization is a single onboarding step.

**Success is measured behaviourally,** not technically: staff capture real activities without a developer present, and the AI correction rate falls week over week.

**Roadmap.** Evidence Core is the first of three views of one system — Evidence, People and Workflow — sharing the same store, consent model and identity. Near-term: field pilots with client staff, richer outcome analytics, and semantic search across all evidence. Deferred by design: connectors to external CRMs, white-label packaging, dedicated-instance hosting.

**Ask.** A pilot needs one organization, one program lead with a phone, and one supervisor with ten minutes a week. If capture takes longer than two minutes or confirmation longer than one, we stop and fix that first — that gate is written into the build plan.

Contact: [info@cyphiasoft.com](mailto:info@cyphiasoft.com) · Product: cyphia.app