

SOPACT INTELLIGENCE LIBRARY · BOOK 01



BEYOND CASE MANAGEMENT

A field guide to **case intelligence** — where a case is a stakeholder or a portfolio, and the requirements document becomes a working workflow you grow as you go.

By **Unmesh Sheth** · Founder & CEO, Sopact · 2026 · sopact.com

COMPLETE GUIDE · 21 SECTIONS

What's inside.

01	The shift	11	Data dictionary
02	Seven words	12	Configuration
03	The spine	13	Four reports
04	Lifecycles	14	Interoperability
05	Data design	15	Four properties
06	Multi-language	16	Worked example
07	Four channels	17	Sense + Skills
08	Transformation	18	Planning
09	Intelligent Suite	19	The long view
10	Prompt-craft	20	The library

● PURPOSE

A guide for choosing a system in the AI era.

If you're evaluating a survey tool, a CRM, or a case-management system right now, this guide gives you a clear, defensible case for **case intelligence** — and explains why the survey, CRM, and case-management eras have ended at the same time.

THE OLD WAY · WATERFALL

A requirements document, pages long, written before anyone touches real data. Signed off — then obsolete by launch.

THE NEW WAY · ITERATE

Start from an end-to-end workflow, collect real data, read it, refine. The requirement emerges from use.

i

A case takes one of two shapes

A **stakeholder** is a person you serve across cycles — applicant, scholar, trainee. A **portfolio** is an organization — investee, grantee, supplier. The relationship differs; the discipline is identical.

● THE DIFFERENTIATOR

The era of case management is over.

The bottleneck used to be **collection**. That moved. The new bottleneck is the workflow that **reads every document on arrival** — every essay, attestation, note, transcript, KPI report — and surfaces the failure before it becomes one. The case is the unit of work, and the case has to be intelligent.

THE ERA THAT ENDED · CASE MANAGEMENT

Months or years to configure even a moderate workflow.

Documents don't read themselves — they wait for a human who never has time.

Harder to maintain every year as changes pile up.

Analytics never produce what each stakeholder wants → export to Excel.

Cost of ownership an order of magnitude higher, with high odds of failure.

WHAT REPLACES IT · CASE INTELLIGENCE

Days to start. Configuration described in natural language, not built by integrators.

Every document is read on arrival; risks surface as data lands.

Focus on what matters day one; add scope without re-platforming.

Reports answer the question each stakeholder actually asks.

Start small and grow — ROI is immediate and risk stays small.

● FROM THE FIELD

An impossible reading, caught in **minutes**.

After heavy rain, Open Play's water system reported it had run out of rainwater — which Marco knew was impossible. Because the reserve, extraction and harvesting figures were now **live**, he cross-checked in minutes and surfaced the real cause: a probable reservoir leak or a failed float switch — caught while still fixable.

"Those statistics we're now running on Sopact immediately showed me there's something significantly wrong ... things like that, we would never have been able to do in the past."

— MARCO BOTHA · CEO, OPEN PLAY FOUNDATION

● THE SHARED LANGUAGE

Seven words this whole library is built on.

This is the parent guide. Every use-case guide that follows inherits these definitions without restating them. Get them precise here, once.

01**Data**

Raw inputs you collect on purpose — PDF, transcript, survey, financial, photo. Not yet meaning.

02**Workflow**

The 3–5 phases a case moves through. The workflow is the system; everything hangs off it.

03**Framework**

The theory of why your program works — ToC, Logic Model, IRIS+. It shapes what you measure.

04**Data Dictionary**

The signed vocabulary of measurement: definition, source, calc, owner. Drafted from your data.

05**Transformation**

The AI work in each phase — scoring, tagging, extracting, joining. Reading replaces filing.

06**Reports**

Intelligence outputs generated as data arrives — missing, unusual, comprehensive, aggregate.

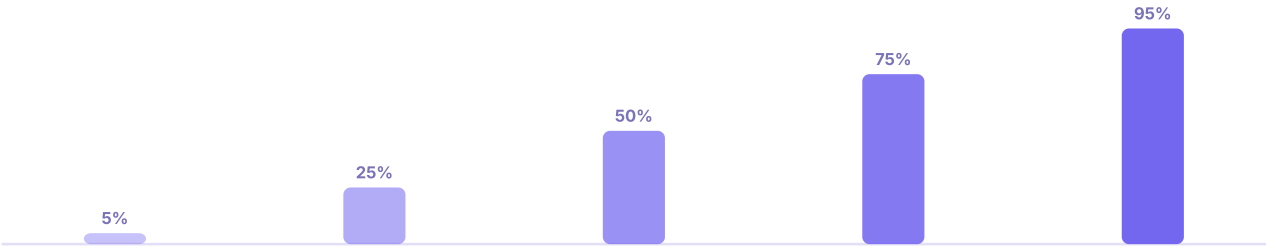
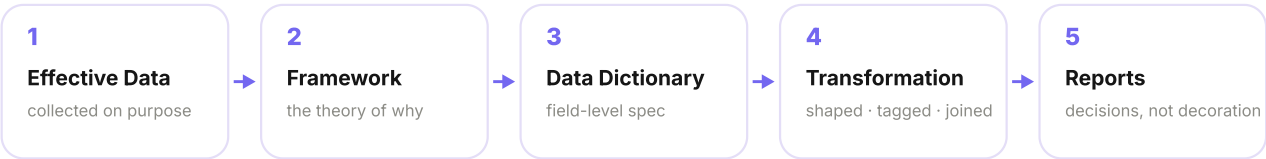
07**Context**

Accumulating intelligence that survives across phases — ~5% at intake to ~95% at exit.

● ONE WORKFLOW, REPEATED

The workflow spine.

Every effective case workflow — whatever the program — moves through the same five stages. Strategy is simply this one spine, repeated for every case.



CONTEXT ACCUMULATES · 5% → 95% FROM INTAKE TO EXIT

AI without a workflow is a clever intern with no desk. The teams winning with AI are the ones whose case data has a place to land.

● THE RECORD CARRIES THE ARC

Shape the lifecycle the record carries.

Before you configure anything, name the lifecycle. Two families — stakeholders and portfolios — and twelve common shapes between them.

STAKEHOLDER · SAME PERSON ACROSS CYCLES

- Application Review apply → admit

- Scholarship Mgmt award → alumni

- Awards & Competition submit → judge

- Training Programs pre → mid → post

- Surveys & Feedback collect → analyze

- Multi-Year Cohorts year 1 → year 5

PORTFOLIO · ORGANIZATIONS ACROSS YEARS

- Impact Portfolio onboard → year 5

- Investment Monitoring DD → exit

- Grant Tracking apply → report

- ESG Diligence screen → monitor

- Supplier Audits onboard → audit

- Cohort Programs cohort 1 → N



The rule

Pick the single arc your record must carry, then assign a persistent ID at the first step so the case stays whole from the left edge to the right.

● DECIDE BEFORE YOU COLLECT

Three design choices, made on day one.

Three decisions determine whether a case can ever become intelligent — plus a fourth layer that decides who ends up in your data. Retrofitting any later is effectively impossible.

01 · METHOD**Mixed method**

The number says **what**; the words say **why**. AI reads both at the same scale.

02 · TIME**Longitudinal**

The same subject across cycles, tied by a persistent ID. No ID, no story.

03 · BASELINE**Pre / post**

A baseline at intake makes every later reading mean something.

4**The fourth layer · designing for the field**

Field conditions decide who's in your data. Plan for offline capture, skip logic, and language before launch — and design for attrition. Assign the ID at first contact, capture a baseline before the program starts, collect words alongside numbers from day one.

● THREE INDEPENDENT LAYERS

Collect in Swahili. Analyze in English. Report in Portuguese.

Designed correctly, language is three independent layers that compose any way — nobody translates by hand, and no quote is read without source-language verification.

1 · COLLECTION

The respondent's language

Skip logic, validation and error states all translate.

- EN "What barriers?"
- SW "Changamoto gani?"

2 · PROMPTS

Read in the original

AI reads responses in source language — no lossy translation first.

- SW response → tag: childcare · transport

3 · REPORTS

Each audience, their tongue

Same evidence, generated from the same dataset.

- EN "confidence rose 38%"
- PT "a confiança subiu 38%"

The end-to-end flow



Swahili response → analyzed in source → English structured tags → Portuguese funder report. 40+ languages on collection, any combination on output. The participant never sees a translated question they didn't ask for.

● COLLECTION IS SOLVED — IF IT JOINS

Four channels. One stakeholder ID.

Most case data isn't in the survey. Online forms, offline field capture, uploaded documents and recorded transcripts are four sides of one case — collection is solved only if every channel lands in the same record.

CHANNEL 01

Online forms

Unique link tied to the case ID — re-entry updates, never duplicates.

CHANNEL 02

Offline field capture

Works without signal; syncs to the same record on reconnect.

CHANNEL 03

Uploaded documents

A 12-page PDF is a primary source — read and scored, not filed.

CHANNEL 04

Recorded transcripts

A 40-minute call becomes structured, timestamped evidence.



Unique-link, not anonymous

Every submission joins the right record automatically — no manual matching.



Documents are data

A primary source read and scored like any field — not an attachment to file away.

● WHAT TEACHES THE AI

Where raw input becomes shaped output.

An open-text answer becomes structured fields — and there are three composable ways to teach the AI what to extract.

```
// open-text · applicant #1042
"i had been laid off from my warehouse job last spring and the program helped me get a CDL. now i drive for a regional carrier, making about 1100 a week, my kids are okay."

prior_status laid_off credential CDL employed_t6 TRUE
weekly_wage $1,100 sentiment positive themes income_gain · family_stability
```

MODE A · DEFAULT

Auto-discovery

Sopact reads your records, infers structure, and writes the Data Dictionary itself — you confirm.

MODE B · SHIPPED

Built-in Skills

Prepackaged playbooks that know common impact patterns.

- sroi-calculator
- iris-plus-mapper

MODE C · CUSTOM

External Skills

Your own markdown — rubrics and domain language in plain English, versioned in your repo.

● SCOPES × OPERATIONS

Four scopes. Four operations. One dataset.

Every AI call is a **scope** — what it operates on — crossed with an **operation** — what it does. Naming both lets you ask a precise question and trust the answer.

	SCOPE ↓		OPERATION →	
	Rubric	Join	Compare	Summarize
Cell one value	1			
Row one case				
Column one field				
Grid whole table		2	3	

A precise call is **one scope × one operation**. The worked sequence: ① **score every essay** (cell × rubric) → ② **join to the Logic Model** (grid × join) → ③ **compare to last cohort** (grid × compare). One dataset, no ambiguity about what was asked.

✕ Scope × operation, in one line
 Score every grantee essay (cell × rubric) → join the scores to the Logic Model (grid × join) → compare to last cohort (grid × compare). Three precise calls, one dataset, no ambiguity.

- A GOOD PROMPT HAS FOUR PARTS

Context. Emphasis. Task. Constraints.

C

Context

What this is and who said it — program type, wave, persona.

E

Emphasis

Where to look hardest. "The inflection moment, not the average tone."

T

Task

The action, in one verb — extract, score, cluster, summarize.

C

Constraints

What the output must not do — "Score 1–5 only. No explanation. No null."

// assembled · one prompt

"In a 12-week coding bootcamp learner exit reflection at T1, **score the response on technical confidence**. **Focus on the inflection moment**, not the average tone. **Score 1–5 only. No explanation. No null.**"

● DRAFTED FROM YOUR DATA

One definition per field.

A framework says **what** matters; a dictionary says **how** it is measured, from what source, with what calculation, owned by whom. Case intelligence drafts it from your data, then asks you to confirm.

Field	Definition	Source	Calculation	Owner
case_id	Persistent identifier — joins every source.	first contact	UUID	Ops
readiness_score	0–100 program-fit at intake.	app + transcript	weighted rubric	PO
welder_wage_gain	Wage change vs. baseline, USD/wk.	survey	IMP "How Much"	M&E
placements_t6	Placements 6 months out.	follow-up	IRIS+ OI4188	M&E
narrative_theme	AI-extracted theme from open text.	survey/transcript	AI extract	M&E

Single owner per field — no collective ownership.

Every calculated field shows **its formula** — no hidden math.


One **case_id** ties every row across every source.

● NATURAL LANGUAGE, NOT INTEGRATORS

Configured in language, not over quarters.

You describe the change in plain language and it takes effect. Three things you configure in a sentence:

<p>ROUTING</p> <p>Send to the next step</p> <p>"Route any application scoring 80+ to the interview stage and assign the regional lead."</p>	<p>FLAG INCOMPLETE</p> <p>Surface unfinished work</p> <p>"Flag any grantee whose quarterly report or financials are still missing this week."</p>	<p>IMMEDIATE NOTIFY</p> <p>Alert the right person</p> <p>"Notify the case manager if a report mentions a safeguarding risk or staffing gap."</p>
---	---	--

 **Why this changes the economics**
When configuration is conversation, the cost of being wrong collapses. You no longer pay an integrator for every adjustment — you start with the few things that matter and adapt as the work teaches you.

● REPORTS ARE QUESTIONS, NOT FORMATS

Four report shapes, by what they answer.

ONE STAKEHOLDER → COHORT

UNUSUAL

Off vs. last cycle or the cohort

A score that swung, a metric far from its peers, a tone that changed.

AGGREGATE

The cohort, longitudinally

Patterns across every case over time — what predicts success, where to put the next dollar.

MISSING

What didn't arrive that should have

The overdue follow-up, the document never uploaded, the member who went quiet.

COMPREHENSIVE

One stakeholder, longitudinally

Everything known about a single case, across every source and phase.

AS DATA ARRIVES · POINT-IN-TIME

LONGITUDINAL · OVER TIME

Why this matters for evaluation



Missing and unusual run continuously, so problems surface as data arrives, not at quarter-end. Comprehensive and aggregate give operator and strategist the same truth at two scales — without anyone exporting to a spreadsheet.

● OUT TO YOUR STACK

One source. The full picture needs several.

Case intelligence is the engine; your stack — BI, warehouse, MCP, GenAI — is the actionable layer. The superpower isn't faster dashboards, it's **joining sources**: unified, they become evidence; siloed, they stay anecdote.

BI & WAREHOUSE

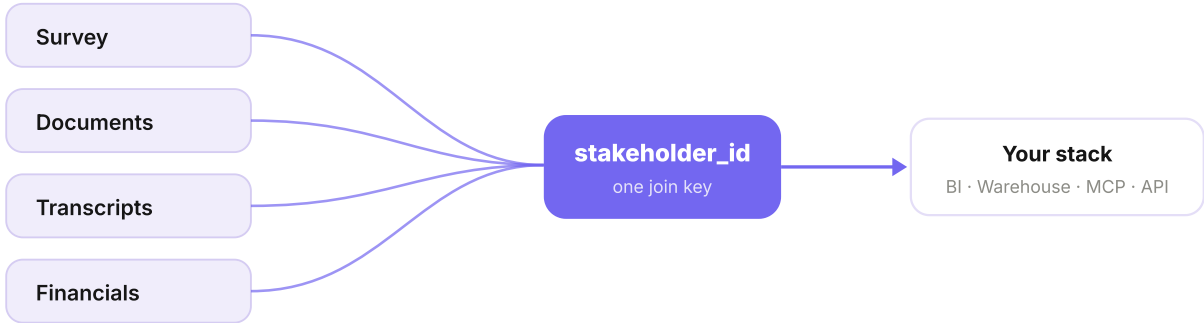
Dictionary-defined data flows out

Tableau, Power BI, Looker, Snowflake — your BI team keeps the dashboards they have, fed by data that finally reconciles.

MCP & API

Readable by AI agents

Exposed MCP-native and via API — a funder can ask a question in natural language and get an answer grounded in your records.



Every source attaches to one persistent **stakeholder_id** — **match once, unify forever**. Unified, the sources become evidence; siloed, they stay anecdote.

● WHAT'S GENUINELY NEW

Four properties a survey or CRM cannot offer.

1

Definitive AI

The same input gives the same answer. Generative tools aren't deterministic — even at temperature 0, accuracy degrades as context grows (Ouyang 2023; Liu 2023). Sopact runs the model to read, then **locks the answer** — reproducible and auditable every time.

2

Two-layer qual + quant

One layer drafts the dictionary from your data; another analyzes words and numbers together — joining a coded theme to a wage number to a rubric score on the same case.

3

Secure, enterprise-grade AI

Runs under GenAI SLAs, stored securely and encrypted at rest — reading every document doesn't mean surrendering control of it.

4

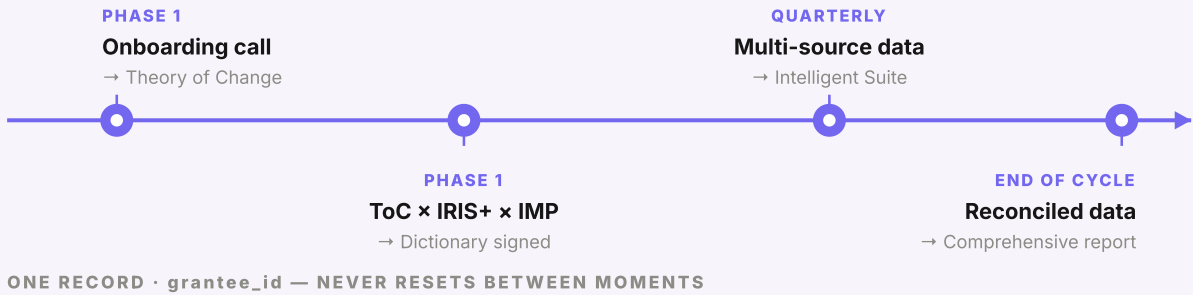
Configured on the fly

New rule, new field, new report — described in language and live immediately. The system bends to the work, not the reverse.

● ONE CASE, END TO END

One grantee. Four moments. Same record.

Watch a 30-minute onboarding call become a board-ready impact report — every metric tracing back to a commitment the grantee made themselves. The record never resets.



Σ **The board-ready result**
18 placements (IRIS+) · **+\$1,100/wk** wage gain · 12-mo recidivism **14% vs. 44%** baseline — every number sourced to the record.

● THE ACCELERANT

The spine is the idea. Sense is where it runs.

Sopact Sense is where the five-stage spine actually runs — collection, AI, warehouse, BI, in one place. **Skills** are how you teach it, and you can compose all three on one case.

A · SHIPS WITH SENSE**Out-of-box Skills**

Playbooks that turn on when a task matches — no setup.

`sroi-calculator``funder-report-composer`**B · YOURS****Organization Skills**

Your markdown — rubrics, weights, domain language. Versioned; override the built-ins.

`disability-employment-rubric`**C · THE DIFFERENTIATOR****Intelligent Reporting Skills**

Give it data plus a question — it builds indicators, attribution and narrative from your prompt.

**Why this compounds**

Each skill encodes a battle-tested playbook. As more programs run through Sense, the skills sharpen — your team starts from the best version of last quarter's work, not a blank page.

● START SMALL, GROW DELIBERATELY

A planning checklist for any use case.

You don't need a forty-page requirements document — you need clear answers to a short list, then a working workflow you grow.


- 01 **Name the lifecycle the record carries** one arc, left → right

- 02 **Decide stakeholder or portfolio, assign a persistent ID** set at first contact

- 03 **Capture a baseline and plan for the field** pre/post · offline · language

- 04 **List the documents that arrive, and what each reveals** each a primary source

- 05 **Configure one lifecycle, one framework, one report — then grow** add as work reveals



Start small, grow deliberately
Configure the lifecycle, one framework, and one report first. Add fields, rules and report shapes as the work reveals what it needs — because configuration is a sentence, not a project.

● HOW THE SYSTEM AGES

It only gets better with age.

Old systems fail the same way: years of rules and hard-won knowledge set up by hand, growing more complicated, understood by almost no one. Case intelligence works the other way.

1

Described, not built

You write what should happen in plain words. Changing it is another sentence, not another project.

2

Read, not retyped

Point it at the reports and transcripts you have; it drafts your definitions for you to confirm.

3

Days, not years

Nothing is rebuilt from scratch, so moving over is fast — and you're not held by the cost of starting again.

4

Sharpens over time

Every cycle improves the rubrics and adds to the record. It grows more useful, not more fragile.



And it keeps up on its own

Not tied to one AI model, the system keeps improving as the technology improves — without you rebuilding anything. You're no longer maintaining a system. You're growing one.

• THE LIBRARY · USE-CASE GUIDES

Guides that inherit this base.

This is the parent guide. Every guide below moves a field from **case management** to **intelligence** — same vocabulary, same five-stage spine, no term redefined.

- 03 From application → application intelligence apply → admit → alumni
- 04 From grant → grant intelligence apply → report
- 05 From ESG case mgmt → risk intelligence screen → monitor
- 06 From investment → portfolio intelligence DD → exit
- 07 From learner → learning intelligence pre → mid → post
- 08 From nonprofit → program intelligence the whole spine

See one of your own cases, made intelligent.

Bring a real application, grant report, or transcript. We'll run it through the spine live — start small, prove it in a day.

Book a working session
sopact.com

The bottleneck moved. Start with one workflow — and let the case carry the rest.