sopact

SOPACT INTELLIGENCE LIBRARY
BOOK 03 · USE-CASE GUIDE

MANAGE → REVIEW → PROVE

sopact

3 CHAPTERS · 22 SECTIONS

FROM GRANT MANAGEMENT TO GRANT INTELLIGENCE.

BOOK 03

Unmesh Sheth · Founder & CEO, Sopact · 2026

sopact.com

CONTENTS

What's inside.

Three chapters, one continuous record. The arc runs the length of a grant—from the software that holds it, to the review that awards it, to the years of outcomes that prove it mattered.

PURPOSE · WHO THIS IS FOR

A guide for anyone who funds, reviews, or proves.

If you run a grant program—a foundation, a fund, a corporate giving team, an accelerator—and your year is spent collecting applications, scoring them fairly, and reporting outcomes a board will trust, this guide is for you. It explains why grant **management** software is giving way to grant **intelligence**, and what changes for each stage of the work.

The argument is simple. Collecting the data was never the hard part—any tool can take an application or a report. The cost was always what came after: the cleanup before analysis, the reviewer drift no one caught in time, the three-week scramble to assemble a portfolio for the board. Grant intelligence moves the work the software should carry off the team's plate, and keeps one record per stakeholder from first contact to year-three outcome.

1

Chapter 1

From grant management software to **grant intelligence** — what it means, and the persistent record it stands on.

2

Chapter 2

How **application review** changed, and how to design reviewer bias out of the default path.

3

Chapter 3

Post-award & compliance — the years after the check clears, where outcomes are actually proven.

i

One record runs the whole arc

The applicant in Chapter 2 becomes the grantee in Chapter 3 on the same persistent Contact ID. Read the chapters in order and you watch one record gather context instead of resetting—application, to award, to outcome.

CHAPTER ONE

What grant intelligence means, and how it changed.

Every other capability in this book depends on one idea: a single record per stakeholder that never resets between cycles. This chapter defines it.

THE DEFINITION

What is grant intelligence?

Grant intelligence is the practice of holding one persistent record for every stakeholder a funder works with—applicant, grantee, investee, partner—so that data from every form, document, and report accumulates on that record instead of resetting each cycle. It turns scattered, one-off feedback into a continuous, connected view of each relationship.

> **?** **The test is simple**
> Can your team answer what changed for a specific grantee over three years—without rebuilding the data first? If the answer is “not without a week of cleanup,” you have grant management. If the record already knows, you have grant intelligence.

IT DESCRIBES A CAPABILITY

Not a dashboard, not a single feature. It is the discipline of carrying the relationship forward—so the next decision builds on everything already known about that stakeholder.

USED ACROSS THE FUNDER’S WORK

Grant review · scholarship & fellowship programs · impact-fund portfolios · accelerator cohorts · corporate giving · multi-year outcome tracking.

THE RESET PROBLEM

Most tools start from zero, every cycle.

The data exists. The problem is that it lives in tools that forget the stakeholder between one cycle and the next—so every cycle starts with cleanup, not analysis.

TOOLS THAT RESET · CONTEXT IS LOST AT EVERY HANDOFF

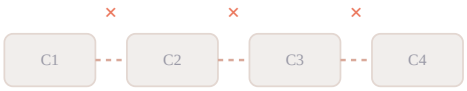

Every cycle begins at zero. The record of **who this stakeholder is** does not survive the handoff between tools.

GRANT INTELLIGENCE · CONTEXT COMPOUNDS

Every cycle builds on the last. The record carries forward on one **persistent Contact ID**—each interaction makes the next one sharper.

What changed in 2026: the cost stopped being hidden

It used to be absorbed quietly in staff time. Now AI made it loud—point an AI tool at clean, connected records and it produces traceable answers; point it at scattered exports and it produces confident guesses. At the same time, funders and LPs stopped accepting an annual snapshot as evidence.

THE ARCHITECTURE

The persistent Contact ID is the moat.

A persistent Contact ID is one identifier attached to a stakeholder that stays with them across every form, survey, interview, document, and check-in—regardless of an email change, a re-spelled name, or a new address. It is the difference between a pile of responses and a record.

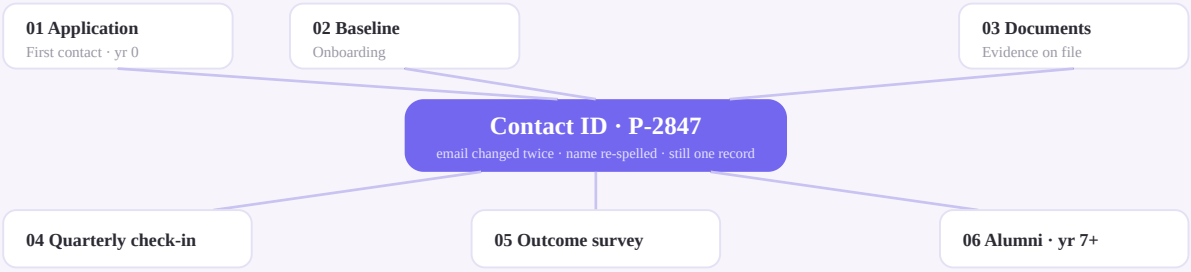

Six touchpoints across seven years — one identifier holds them together.

This is the line that separates a feature from a system. Survey tools were built around the survey instrument. Sopact was built around the persistent Contact ID—and has been since 2014. Continuous feedback bolted onto a survey tool is a feature; continuous feedback architected around the Contact ID is the system.

HOW THE ANALYSIS WORKS

Four layers of analysis, on one record.

Two layers operate the moment data is collected; two operate when reports run. All four work because every record carries the same Contact ID—so analysis never has to ask which row belongs to whom.

At collection

Intelligent Cell

SCOPE: ONE FIELD

The moment a document or essay arrives, it is read against the rubric your team defined—with its reasoning attached to the record.

At collection

Intelligent Row

SCOPE: ONE RECORD

Resume, recommendation, intake survey, and prior history combined into one coherent brief—no holding five tabs open at once.

At reporting

Intelligent Column

SCOPE: ONE QUESTION, EVERY RECORD

Themes extracted across hundreds of open-ended answers; sentiment trends across a multi-year cohort, ranked and tagged.

At reporting

Intelligent Grid

SCOPE: THE FULL DATASET

Cohort-versus-cohort comparison and funder-ready outputs—without weeks of spreadsheet reconciliation.

→ **Read the grid left to right**

A single cell, then a whole record, then one question across every record, then everything at once. The same data, four scopes—one for every question a team actually asks.

THE AI-NATIVE SHIFT

AI is only as reliable as the record beneath it.

A general-purpose AI tool can summarize anything you hand it. That is exactly the problem. Grant intelligence fixes the input, not the model.

AI OVER SCATTERED EXPORTS

A fluent answer over five exports of one cohort. It changes every run, nothing traces back, and the codebook is whatever the model inferred this time. Confident, and impossible to defend.

- answer drifts each run
- no source to open
- codebook guessed

AI OVER ONE CONNECTED RECORD

The same question returns the same answer. Every score traces to the source text, and the analysis runs against the codebook the team defined—not a fresh interpretation each time.

- reproducible
- citations attached
- your codebook

The point of the architecture is not a cleverer model—it is an answer a program officer can defend in a meeting.

REPRODUCIBLE BY DESIGN · THE SAME QUESTION, ASKED TWICE, RETURNS THE SAME ANSWER

WHAT CHANGED · AND WHO IT CHANGED

Software collects. Intelligence carries the record forward.

Survey tools end at collection. Application platforms end at the award decision. Bundled CRMs swap depth for breadth. Each is good at one slice; none was built to hold the whole relationship.

THE ERA THAT IS ENDING **Grant Management**	WHAT REPLACES IT **Grant Intelligence**
A persistent ID is a custom field or an export-time merge.	**Native to every record,** from first contact onward.
Analysis happens later, in a separate dashboard.	**Cell and Row analysis at intake,** traceable to the rubric.
Text in one tool, numbers in another, reconciled by hand.	**Open text, documents, and numbers on one record.**
After collection you have a spreadsheet to clean.	**A continuous, connected view,** ready to report.

WHO IT CHANGED FOR

Foundations

Intake, mid-cycle reporting, and outcome evidence on one record per grantee—annual reports without the reconciliation week.

Funds & LPs

Due diligence through year-seven exit on one record—quarterly signals, coded outcomes, a risk baseline that updates.

Accelerators

Cohort-versus-cohort comparison, because every participant sits on the same architecture and the same ID.

CHAPTER TWO

How application review changed—and how to reduce reviewer bias.

Collection is solved. The new bottleneck is the review itself: reading every application fairly, scoring it against the same rubric, and proving why each score is what it is.

BEYOND READING EVERY APPLICATION BY HAND

Review intelligence has begun.

Manual review doesn’t scale and it doesn’t stay fair. Reviewers open long PDFs cold, scores drift between people and across the afternoon, and reputation, halo, and anchoring quietly bend the outcome. The fixes everyone knows—blind review, a tight rubric, calibration—are hard to hold to by hand when the pile is hundreds deep.

The application is the unit of work

Sopact reads each application on arrival, scores it against your rubric with a citation trail back to the text, and makes blind, calibrated review the default rather than the exception—one application, one record, one defensible decision.

READ ON ARRIVAL

The committee opens to a shortlist, not a queue. The reviewer still decides—from evidence, against one rubric.

THE SHIFT

The era of reading every application by hand is over.

Not because reviewers stopped mattering—because asking a human to read hundreds of applications cold, hold one rubric in their head, and stay unbiased all afternoon was never going to be consistent. The judgment is the point; the reading and first-pass scoring is what the software should carry.

MANUAL REVIEW **The era that’s ending**	REVIEW INTELLIGENCE **Sopact**
A 30-page PDF, read cold, late in the pile.	**A structured summary + draft score,** each point cited to the text.
Rubric in a separate doc, applied from memory.	**Rubric encoded once,** applied identically to every application.
Consistency drifts; discovered at the decision meeting.	**Outliers flagged during review;** calibrate before deciding.
Bias—reputation, halo, anchoring—quietly bends the outcome.	**Blind review by default;** identity masked, evidence shown.
Defending a decision: “the committee felt…”	**Every score traces** to the text it came from.

That’s review intelligence: human judgment, machine consistency—and the decision is auditable afterward.

DESIGNING BIAS OUT OF THE DEFAULT

You can’t legislate bias away. Make the fair path the easy one.

Funding and reputation bias creep in when reviewers can see who applied and lean on prior funding, brand, or relationships rather than the application in front of them. Four mechanisms move the fair process from exception to default.

1

Blind review by default

Applicant identity and any fields you choose are masked—so anchoring, halo, and reputation effects lose their hook.

2

One explicit rubric

The same criteria and weights apply to every applicant, scored from cited evidence rather than impression.

3

Calibration during review

Reviewers see the same evidence-linked summary; divergent scores surface while there’s still time to re-calibrate.

4

Outlier flags

A scorer consistently far from the committee average is flagged—before the score decides who gets funded.

×

The boundary: AI drafts, humans decide

Sopact does not decide who gets funded. The AI reads and drafts scores from evidence; the rubric is human-set and the decision is human-made. Anyone selling an “AI picks the winners” black box is selling the thing you’ll have to defend later.

FROM THE FIELD · OPEN PLAY FOUNDATION

What happens when the review can finally read the application.

Open Play Foundation evaluated the way most committees do: stacks of narrative, a rubric in a separate document, and reviewers doing their honest best across far too many submissions. The judgments that mattered were buried in text no one had time to read closely, and the patterns across the pool stayed invisible.

When Open Play moved that work onto Sopact, the AI read what the committee couldn't read by hand—every response, on arrival, against the same rubric. What had been impossible to see across hundreds of submissions surfaced as soon as the review could read itself.

> "Those statistics that 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

That is the difference between review and review intelligence. Manual review tells you the committee got through the pile. An intelligent review tells you which scores don't hold up, where bias crept in, and why each decision was made—in time to fix it before the funding goes out.

THE LIFECYCLE

Five stages, one application record.

Every review program runs the same cycle. Most tools cover intake and a scoring form; the spine below keeps the rubric, the AI read, the calibration, and the audit trail on one record—intake to defensible decision.

STAGE 1 · INTAKE

White-label forms, one persistent applicant ID, and missing-document flags on arrival—review starts from complete applications, not a chase.

STAGE 2 · DEFINE THE RUBRIC

Criteria, weights, and what good looks like—set in plain English, versioned, applied identically to every application.

STAGE 3 · AI READS & SCORES

Each application read on arrival; a draft score per criterion, every number linked by citation trail to the text.

STAGE 4 · REVIEW & CALIBRATE

Blind review by default, conflict-of-interest rules, and outlier flags—so you calibrate before the decision, not after.

STAGE 5 · DECIDE & AUDIT

The decision attaches to the same record, every score traceable to its evidence—defensible to a board or a compliance review.

LIVE IN DAYS

The whole spine is configured in plain English, not by a consultant—a first review cycle is live in days, not a quarter.

REPORT SHAPES

Four reports a review process actually needs.

The final ranked list gets the attention. But the reports that keep review fair are simpler—and rarely built, because the evidence is stuck in scoring spreadsheets and reviewer inboxes. Sopact ships all four off one record.

01 · MISSING

Applications & scores not yet complete

Applications missing a document before the deadline; reviewers who haven't scored their assignments. Surfaces the gap while there's still time to close it.

02 · UNUSUAL

Scores that don't look like the rest

A reviewer consistently above or below the committee; a wide score spread; a possible bias signal—seen before it decides anything.

03 · COMPREHENSIVE

The defensible decision record

For each funded and declined application: rubric scores, citation trail, and reviewer notes—the full audit record, as one query.

04 · AGGREGATE

The fairness & throughput view

Score distributions by reviewer, time-to-decision, and outcomes by applicant segment—is the process fast **and** fair, cycle over cycle.

THE PLATFORMS YOU'RE COMPARING

Where the review platforms stop, and what carries on.

Submittable, SurveyMonkey Apply, OpenWater, Good Grants, and SmarterSelect run intake and rubric scoring well, and several do blind review. The two rows most stop short of: reading each application on arrival, and an auditable citation trail behind every score.

CAPABILITY	SOPACT	SUBMITTABLE	SM APPLY	OPENWATER	GOOD GRANTS
Time to first cycle live	Days	Weeks	Weeks	Wk–mo	Weeks
AI reads each application on arrival	Yes	Add-on	No	No	No
Custom rubric scoring	Yes	Yes	Yes	Yes	Yes
Citation trail behind each AI score	Yes	No	No	No	No
Blind review	Yes	Yes	Yes	Yes	Partial
Reviewer-outlier / bias flags	Yes	No	Limited	No	No
Carries result onto an outcome record	Yes	No	No	No	No

Honest reading: these are real, capable platforms. Where none was designed to compete is reading each application on arrival with AI, attaching a citation trail to every score, and carrying the result onto the record that proves the outcome later. Vendor capabilities change; confirm current details before deciding.

CHAPTER THREE

Post-award & compliance—tracking outcomes, not just disbursements.

Post-award grant management is what happens after the check clears: the planning phase, recurring metrics, mid-cycle reports, and the multi-year portfolio rollup. It doesn’t end at the report. It begins there.

AFTER THE AWARD DECISION

The award letter is the handoff, not the end.

Take a regional community foundation deploying $1.8M across three multi-year program grants. Three-year term, six payment milestones tied to deliverables. The application was the easy part. Everything between May 2026 and September 2029 is the work.

$1.8M
deployed across 3 multi-year grants

6
payment milestones tied to deliverables

3 yr
term · planning phase to portfolio rollup

Most grants management software treats the award letter as the finish line. Sopact treats it as the handoff. The grantee's Contact ID from application carries forward. The proposed metrics from their RFP response become the indicators tracked in post-award. The external evaluator joins the same record. **The grantee never re-enters what they already wrote.**

→ **This is the same record as Chapter 2**
The applicant you scored on a rubric is now the grantee you track for three years—one persistent ID, no re-keying at the boundary. The hardest part of switching platforms is preserving exactly this.

AWARD → REPORT → ANALYZE

Three stages, on one persistent grantee record.

The application Contact ID becomes the grantee record; proposed metrics become indicators; reviewer commentary stays attached for context when next year’s officer picks it up.

STAGE 04 · AWARD

The award letter is the handoff, not the end of the record

Application Contact ID becomes the grantee record. Proposed metrics become indicators. Reviewer commentary stays attached. `Planning phase begins`

STAGE 05 · REPORT

Quarterly metrics, semi-annual narratives, milestone-tied payments

Structured KPIs, unstructured narrative, and uploaded artifacts on one form per cycle. Schema evolves across cycles without breaking historical comparability. `Through Sep 2029`

STAGE 07 · ANALYZE

Portfolio rollup the board can read in one sitting

Spend vs. budget by program, outcome distribution, geo concentration, common barriers. Role-based views for board, program officer, and external evaluator. `Portfolio dashboard`

WHERE MOST FOUNDATIONS LOSE THE THREAD

The planning phase, as a structured workflow.

Many RFPs require a three-month planning phase before implementation: a logic model, an evaluation and learning plan, metric selection, and baseline data. Most systems don’t model this phase—the grant just “starts.” Sopact treats it as a workflow on the same record, so the logic model inherits the RFP narrative and the grantee only fills the gaps.

1

Logic model creation

Inputs → activities → outputs → outcomes, on the grantee record. External evaluator collaborates in place.

2

Evaluation & learning plan

What data, how often, who collects it, how it’s analyzed, and what decisions it informs.

3

Metric selection from library

Process metrics plus short-, medium-, and long-term outcome metrics for the program type.

4

Baseline data collection

Pre-implementation baseline on the same indicators that will be tracked quarterly.

i

By the time implementation starts

Every cycle’s metric submissions already know which indicators to ask for, what the baseline is, and which outcomes belong to which logic-model branch. Grants without learning plans tend to produce reports without insight.

REPORTING CADENCE

Quarterly metrics, semi-annual reports—without the cleanup.

Most foundations spend 80% of post-award effort reconciling data submitted across email, spreadsheets, and PDFs—a different format, definition of “served,” and baseline year every cycle. Sopact collects structured metrics, narrative, and documents on one form per cycle, on the grantee’s persistent record.

QUARTERLY SUBMISSION

Process metrics on the same form as the narrative behind them

AI reads the narrative against the metric values and flags discrepancies for the program officer. Same form, same definitions, every cycle.

SEMI-ANNUAL ASSEMBLY

Mid-year and annual reports auto-pull from the quarterly stream

Sopact auto-pulls metrics, narrative, and artifacts; financial actuals come from the compliance handoff. The grantee reviews and adds context.

✓ **Schema evolution, not schema breakage**

Year 1 grantees keep Year 1 questions; Year 3 grantees get Year 3 questions; cross-cycle rollups still work because indicator definitions are versioned, not overwritten.

PORTFOLIO-LEVEL ROLLUP

From 28 grants to one dashboard the board can read.

Mid-cycle qualitative depth and multi-year tracking land in the same place: the portfolio dashboard. Role-based views serve the board, the program officer, and the external evaluator—without forcing anyone into a tool they don't use.

Portfolio rollup · national health nonprofit

28 grants · multi-year · FY24

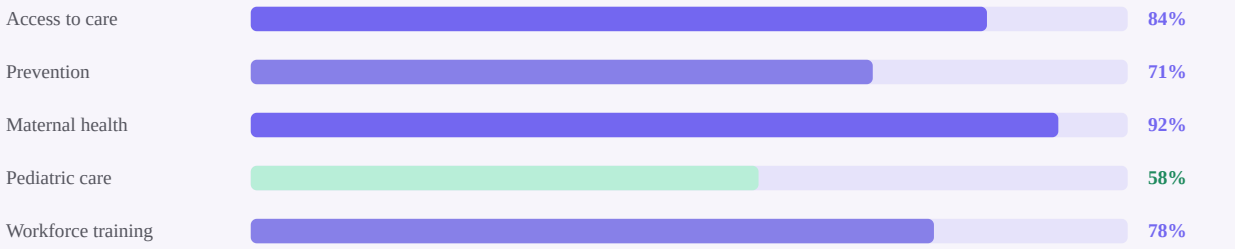

28	61	Staffing
active grantees	mid-cycle interviews	most-flagged barrier

"Before Sopact, 28 grantees meant 28 spreadsheets and a 3-week assembly window. Now I open one dashboard, filter by program, and the narrative pulls automatically from each record."

— PROGRAM OFFICER, NATIONAL HEALTH NONPROFIT

WHERE POST-AWARD HANDS OFF

Once the data is clean, compliance is mostly assembly.

Finance and compliance live on a dedicated layer because the buyer is usually different—a finance lead, not a program officer—and the integration partners are different. But the architectural advantage carries through: the same grantee record that holds quarterly metrics holds the finance disbursement record.

STAGE 06 · FINANCE

ERP integration for budget and AP (Workday, NetSuite, Sage Intacct). Disbursement reconciliation per grantee, tied to milestone completion.

STAGE 08 · COMPLY

990 reporting and state AG charitable-trust filings. Board compliance packets auto-assembled. An audit trail across the full grant lifecycle.

→ **No re-keying, no second product, no broken thread**

The same persistent ID that surfaces in the portfolio dashboard surfaces in the board compliance packet. For public agencies with audit requirements, the persistent-record model has the strongest audit trail—every cycle of data lives on the same record as the original RFP response and the award decision.

POST-AWARD STEP	SPREADSHEETS + EMAIL	FLUXX / FOUNDANT MODULE	SOPACT PERSISTENT RECORD
Grantee onboarding	Re-enter proposal data	New record, manual link	Application ID becomes the grantee record
Metric submission	Format drifts each quarter	Fixed form, hard to evolve	Same form, schema evolves additively
Portfolio rollup	3-week assembly per meeting	Manual export to board format	One dashboard, role-based views
Multi-year tracking	Yr 1 can’t speak to Yr 3	Careful schema management	Additive, versioned indicators

THE THROUGH-LINE

Three chapters, one record.

The arc of this book is the arc of a single Contact ID. The applicant scored in Chapter 2 is the grantee tracked in Chapter 3, on the record defined in Chapter 1. Nothing resets at the boundaries—and that is the whole point.

DEFINE IT ONCE

One persistent record per stakeholder, qualitative and quantitative together, analysis at collection.

AWARD IT FAIRLY

Blind, rubric-driven, calibrated review with a citation trail behind every score.

PROVE IT FOR YEARS

Quarterly metrics, mid-cycle depth, and a portfolio rollup that holds up to an audit.

SEE IT ON YOUR DATA

Grant intelligence starts with one record.

Bring one cohort, one grantee portfolio, or one active grant—its planning artifacts and the last cycle’s metric submission. The walkthrough uses your records, not a demo account.

SOPACT · GRANT INTELLIGENCE

Read every application. Score it fairly. Prove what it changed.

No demo theater, no discovery phase. Tell us your rubric, your reviewers, and your reporting cadence—we’ll show you the full lifecycle on Sopact, against your own data, live this cycle.

Book a walkthrough →

sopact.com

Stakeholder Intelligence | Grant Application Review | Post-Award Grant Management | Sopact Sense