sopact

SOPACT INTELLIGENCE LIBRARY
BOOK 02 · USE-CASE GUIDE

APPLY → ADMIT → ALUMNI

sopact | 21 SECTIONS

FROM APPLICATION MANAGEMENT TO APPLICATION INTELLIGENCE.

BOOK 02

Unmesh Sheth · Founder & CEO, Sopact · 2026

sopact.com

PURPOSE · WHO THIS IS FOR

A guide for teams running applications under deadline.

If you are about to launch a grant cycle, a scholarship round, a fellowship call, an accelerator cohort, an awards process, or a contest — and the rubric isn't finished, the reviewers aren't booked, and the deadline is three weeks away — this guide is for you.

Two things are different about application work. The **deadline is external** — applications close when they close, the panel meets when it meets, the funder needs the cohort report by the date on the letter. And the **decision is contested** — every yes is a no to someone else, and the rationale has to be defensible to a board chair, an auditor, and the applicants themselves. Application Intelligence is the discipline of making both easier, not harder, as volume grows.

THE OLD PROMISE · WORKFLOW + REVIEWER MATRIX

Intake form, conditional logic, portal, status emails, rubric inside the app, scoring matrix. A configuration project that ships before anyone reads an essay.

THE NEW PROMISE · READ ON ARRIVAL

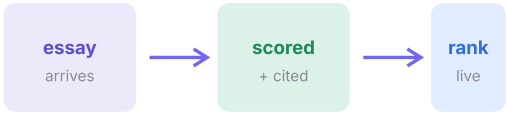

Every application read against the rubric the moment it lands. Citations attached. Ranking updates live. The committee opens to a shortlist, not a queue.

HOW TO READ THIS GUIDE

The first sections give the new vocabulary and method. Then the case study — Carnegie Mellon's Project Olympus, an accelerator that needed to launch in one week and decide in three. The middle applies the five-stage spine stage by stage. The last sections take the lifecycle past the decision — the year-one, year-two, and year-three follow-up where the program actually pays off.

— THE DIFFERENTIATOR

The era of application management is over.

For ten years the category sold two things every team needed in 2014: a built-in workflow and a built-in reviewer matrix. Both were genuine wins. Both are now the reason a cycle takes ten weeks instead of three.

The workflow became the configuration project — months to launch, repeated for every program. The matrix records numbers but can't tell you whether the essay was read. **Reviewer drift surfaces in the export after the committee has met** — and the fix waits for next cycle. The new bottleneck reads every application on arrival and turns the queue into a ranked, citable shortlist.

THE ERA THAT ENDED **Application Management**	WHAT REPLACES IT **Application Intelligence**
Takes **two to three months** to configure before the first cycle launches. Each new program repeats the work.	**Live in days.** Rubric described in plain words, intake drafted in a sentence, scoring active the moment the first application arrives.
Reviewers **read what they had time for**. The shortlist is the first forty applications, not the best forty.	**Every essay is read** against the rubric overnight, citations attached. The committee opens to 40 admits, 97 borderline, 363 cleared.
Drift **surfaces in the export** after the committee has voted. The fix waits for next cycle.	Drift is flagged **mid-cycle**, broken out by reviewer, track, and dimension. The chair recalibrates before committee.
The record **resets at the award letter**. Year-one outcomes live in a different system with a different ID.	One applicant ID from intake through alumni. Re-applicants show up with prior essays, scores, and outcomes attached.
The cohort report is a **two-to-four-week CSV merge** across four sources.	The cohort report is **one query**. Scores, decisions, demographics, and follow-up sit on the same row.

★★★★★

"Three days from setup to live. Applications opened Monday. By Tuesday morning we had scored results more detailed than anything our review committee produced in previous cycles."

— Program Director · Fellowship Selection · top-10 U.S. research university

— THE SHARED LANGUAGE

Seven words specific to application work.

The parent guide named seven words the whole library is built on — data, workflow, framework, dictionary, transformation, reports, context. This guide adds seven more, specific to running an application cycle.

01 **Applicant Record**	**One row per applicant, kept across every stage.** Intake, clarification thread, rubric scores, panel comments, decision rationale, year-one follow-up — all on the same persistent ID. Not five spreadsheets joined by name.
02 **Rubric**	**The framework, applied to one applicant.** The dimensions you score on, the weights, the scale — described in plain words and applied to every essay automatically. A rubric is a framework with weights and a scale.
03 **Threshold**	**The minimum total score required for consideration.** Below it, documented as not qualified — with the per-dimension breakdown that says why. Above it, the applicant enters the borderline pool the panel reviews by hand. What makes a long queue tractable on Tuesday.
04 **Citation Trail**	**The evidence behind every proposed score.** Each score points back to the essay paragraph, recommendation sentence, or budget line that produced it. When the board chair asks why this 40 and not those 40, the answer is on the record.
05 **Reviewer Drift**	**When one reviewer scores systematically above or below the panel mean** over a multi-week cycle. Detected during the cycle, broken out by reviewer, track, and dimension. Reviewer B at 18% high on the climate track is fixed on Tuesday, not next cycle.
06 **Tiebreaker**	**The named dimension that breaks ties at the threshold** — geographic distribution, demographic targets, first-generation status, regional commitment. Defined before the cycle opens. A policy made in advance, not a vote made in the room.
07 **Follow-up Thread**	**The post-decision continuation of the record** — year-one milestone survey, year-two outcome check-in, year-three alumni reflection. Same applicant ID, same row, same evidence base. The record does not reset at the award letter.

THE METHOD SPINE

The five-stage spine, applied to one applicant.

The parent guide named the spine. This is what it does when the case is an application and the deadline is Friday. Five stages, one applicant, three weeks of elapsed time.

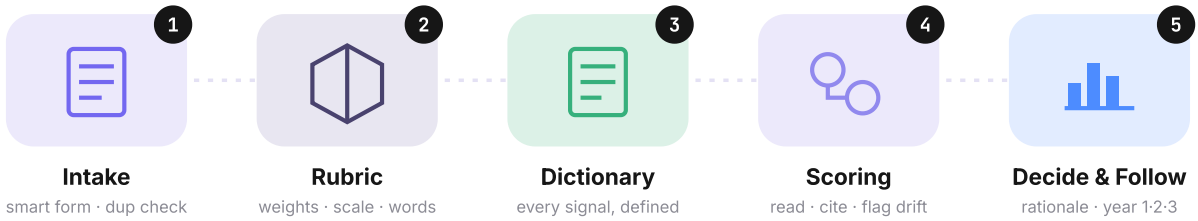

> **The shortlist isn't the best forty applicants — it's the first forty your team had time to read. Fixing that is a system problem, not a discipline problem.**

— THE SIX SHAPES

Six program shapes.
One applicant record.

The lifecycle changes, the dominant work changes, the reports the funder wants change. The record underneath — one applicant carried across every stage — does not.

01
Scholarship
essay → award → alumni

02
Grant
LOI → proposal → report

03
Fellowship
apply → interview → cohort

04
Accelerator
pitch → cohort → demo day

05
Award
nominate → judge → laureate

06
Contest
submit → round → winner

— WHERE THE WORK WEIGHT SITS — BEFORE & AFTER THE DECISION

BEFORE · HOW THE APPLICATION READS **Pre-decision weight**	AFTER · HOW THE FOLLOW-UP UNFOLDS **Post-decision weight**
Scholarship. Essay-heavy review, recommenders, transcripts, financial-aid context.	**Light.** Annual academic progress, alumni outcome surveys, milestone evidence to donors.
Grant. LOI, full proposal, budget PDF, theory of change. Multi-round review.	**Heavy.** Multi-year reporting, drawdowns, amendments, compliance. The record extends for years.
Fellowship & accelerator. Pitch deck, founder questionnaire, reference calls, panel selection.	**Medium.** Curriculum, check-ins, demo-day metrics, alumni engagement. Earlier cohorts inform the next.
Award & contest. Nomination, blind review, conflict-of-interest routing, tiebreakers.	**Light.** Announcement, outcome tracking, juror agreement statistics for the next round.

THE RULE

Pre-side workflow weight and post-side intensity vary independently. A scholarship has heavy intake and light follow-up; a microgrant layered with case management has light intake and heavy ongoing service. One record per applicant handles every combination — the applicant's identity doesn't break when the decision is made.

— THE ECONOMICS

Three flaws baked into the legacy category.

Every team running on Submittable, SurveyMonkey Apply, WizeHive, or Award Force feels the same three pressures. Naming them is the first step to choosing differently.

1

Long configuration and workflow cycles

Two to three months from contract to first cycle live. Intake forms, conditional logic, the rubric, the scoring matrix, conflict-of-interest, reviewer roles, panel routing — each wired by hand, repeated for every new program. **The flexibility is the cost.** A bolted-on AI summary doesn't change the configuration cycle; it just adds a tab.

2

Analysis that doesn't match what reviewers need

The platform records numbers. It can't tell you whether the essay was read. The shortlist is determined by who opened which application first. Reviewer fatigue sets in at application #30. **Drift surfaces in the export after the committee has met** — and the fix waits for next cycle. The platform does what it was built to do in 2014; the analysis the committee needs in 2026 is somewhere else.

3

Total cost of review is an order of magnitude higher than it looks

The license fee is the smallest line. The real number is **configuration time + license + reviewer hours × reviewers + the two-to-four weeks** reconstructing the cohort report every cycle. Multiply by every program, by the subjectivity of judgment, by the cycles where the shortlist was the first forty instead of the best. **The economics collapse** the moment you ask what each defensible yes actually cost.

The reframe

Application Intelligence inverts all three. Configuration is a sentence, not a quarter. Analysis is the work the system was built for, not a feature bolted on. The economics improve as volume grows — because every essay is read on arrival, not by a reviewer with a deadline.

— THE NEW PARADIGM

Read on arrival.
Decide on rationale.

Application Intelligence is built on three convictions the older platforms cannot adopt without rebuilding their foundations. Naming them tells you what to look for in a demo.

CONVICTION 1

Every essay is read

Not the ones reviewers had time for. Not the first forty. **All of them** — against the same rubric, supporting paragraphs attached. The cost of reading the 481st essay equals the first.

CONVICTION 2

The rubric can change mid-cycle

Criteria evolve as the panel learns from the first hundred. **The rubric is a sentence, not a configuration project.** Adjust a weight, add a dimension — the whole cohort re-scores in minutes.

CONVICTION 3

Consistency, not subjectivity

The same input produces the same proposed score every time. **Reviewer judgment stays on the calls that require it.** Drift detection runs in the background; the bias audit is part of the report.

WHAT THIS LOOKS LIKE FRIDAY MORNING

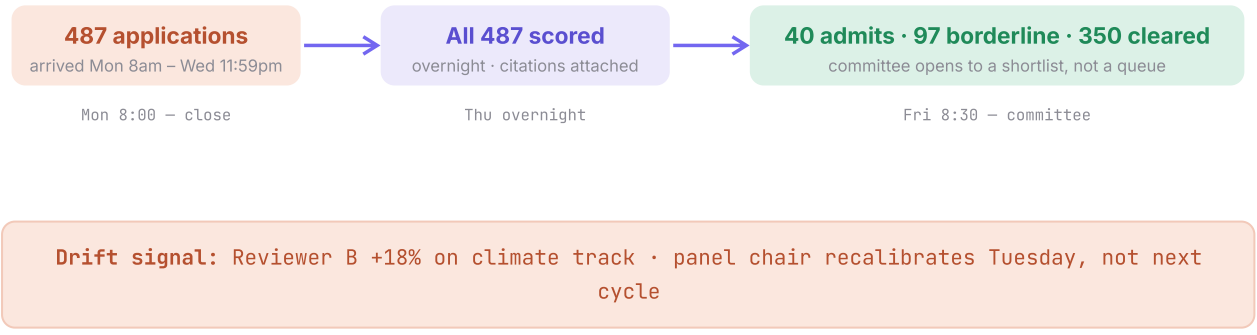

Drift signal: Reviewer B +18% on climate track · panel chair recalibrates Tuesday, not next cycle

— FROM THE FIELD · CARNEGIE MELLON UNIVERSITY

One week to launch.
Three weeks to decide.

Project Olympus is CMU's startup incubator for faculty, students, recent alumni, and staff turning research into ventures — open to entrepreneurs at any stage, from concept to business creation.

For their **Physical AI in Sports** track, the team needed to collect hundreds of applications, score them against a custom rubric, and make funding decisions on a hard external deadline. They spoke to Submittable and Sopact. Both claimed AI. Only one was AI-**native**; the other was AI-**bolted** — a workflow product with an AI summary tab on top.

SUBMITTABLE — AI BOLTED ON

Configuration cycle measured in **months** — past the one-week launch window before anything could be live. Reviewer matrix as in 2014. AI summary attached to each application as a tab; every review still done by hand. The deadline would have been missed before the program opened.

SOPACT — AI NATIVE

Configuration described in plain words. **Application live within days.** Every submission scored against the 80-point rubric on arrival. Citation trail on every dimension. Threshold filter applied automatically. Borderline pool surfaced for the panel. Tiebreakers documented before the cycle opened.

ABOUT THE PROGRAM · IN THEIR OWN WORDS

"Project Olympus offers support and resources for CMU faculty, students, recent alumni, and staff who want to turn their cutting-edge research and great ideas into startups. Applicants learn how to find product-market fit and develop a feasible business model to attract the funding that takes a venture to the next level."

— Project Olympus · Carnegie Mellon University · institutional mission

WINDOW TO LAUNCH	WINDOW TO DECIDE	RUBRIC
1 week	3 weeks	80 pts · 6 dims

FROM THE FIELD · SCORED ON ARRIVAL

A single applicant, scored the moment it arrived.

Every submission to the Physical AI in Sports track was read against the 80-point rubric on arrival, each dimension cited back to the source. Here is one record the panel opened to on Tuesday.

Application #PA-1138 · Physical AI in Sports — scored 2026-04-22 · 11:08 EDT

Deployability Path to a working pilot deployment		2 / 16
HW–SW Integration Hardware-software systems thinking		1 / 12
Pilot Traction Evidence of working pilots or trials		6 / 15
Defensibility Moat, IP, technical advantage		7 / 15
Business Model Revenue model, unit economics		9 / 16
PA Ecosystem Commitment Pennsylvania presence & engagement		6 / 6

Total · threshold 50 — **31** / 80

NOT QUALIFIED · below threshold
tiebreaker: PA ecosystem commitment

REVIEWER NOTE · AI-FLAGGED ON THE RECORD

"Applicant is outside Physical AI in Sports scope; scores interpreted conservatively per rubric. Items lacking evidence are marked as gaps." The reviewer sees the score, the per-dimension breakdown, and the option to override any line. **The panel decides what to do with #PA-1138 — but on Tuesday, not next week, and with the rationale already attached.**

STAGE 01 · EFFECTIVE DATA

Five kinds of input. One applicant ID.

An application is a form, plus a stack of documents, plus a recommendation letter, plus — increasingly — an interview transcript. And a quiet fifth input: the mixed-method fields — closed-ended ratings scored together with the open text beside them.

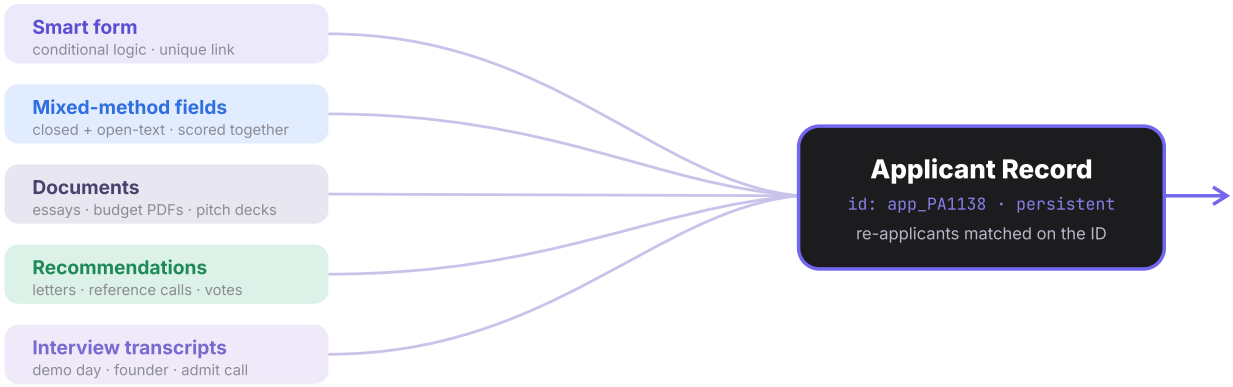

THE FIFTH INPUT · WHY MIXED-METHOD SCORING MATTERS

Most rubrics mix structured and unstructured fields on purpose — a Likert rating on team strength next to a two-paragraph explanation of **why**. Older platforms score these independently: the number lands in the matrix, the text waits for a reviewer who may not read it. Application Intelligence scores the **field cluster as one unit** — the closed-ended anchor and the open-text justification together, citations on both.

```
// closed-ended
team_size_q1: 3
revenue_stage: pre-revenue
pilots_running: 2
founder_full_time: yes
```

+

```
// open-text justification
“Two pilots are with a regional college
sports network; the second goes live in
May. We're pre-revenue but the second pilot
is paid…”
```

→ Pilot Traction · **scored 6/15 from both fields together**, citation pointing to “May” and “pilots_running: 2”

— STAGE 02 · FRAMEWORK / RUBRIC

A rubric, described the way you'd brief a panel.

In Application Management, the rubric was a configuration screen — dimensions, weights, ranges, picked from a dropdown. In Application Intelligence the rubric is a sentence the system reads and applies. Every dimension carries its own description, evidence requirements, and scale.

```
// the Physical AI in Sports rubric, described

"Score each application on six dimensions, summing to 80.
Deployability (0-16) — path to a working pilot, evidence of a prototype.
HW-SW Integration (0-12) — hardware and software thinking together.
Pilot Traction (0-15) — evidence of pilots, trials, or paying customers.
Defensibility (0-15) — moat, IP, technical advantage, regulatory barrier.
Business Model (0-16) — revenue model, unit economics, growth path.
PA Ecosystem Commitment (0-6) — Pennsylvania presence, hires, partnerships.

Threshold for consideration: 50. Tiebreaker: PA ecosystem commitment.
Mark gaps explicitly. Cite the paragraph behind every score."
```

Dimensions, in language

Each dimension is a description the AI reads, not a field name. Add a dimension, change a weight — a sentence, not a project.

Threshold & tiebreaker, in advance

Both decided before the cycle opens, not negotiated in the boardroom. Documented on every record within five points of the threshold.

Citation as a requirement

The rubric specifies that every score points to its supporting paragraph. The reviewer reads the evidence, not just the number.

WHY THIS MATTERS FOR THE DEADLINE

A traditional rubric configuration costs days of integrator time; changing it after the cycle opens costs another round. With the rubric described in language, the panel chair adjusts on Tuesday — add a tiebreaker, raise the threshold, sharpen a dimension — and the whole pool re-scores by Wednesday morning.

STAGE 03 · DATA DICTIONARY

Every field the rubric points to, defined once.

The rubric says what to score. The dictionary says how each signal is measured, where it comes from, who owns it — drafted from your data automatically, then confirmed before the cycle opens.

FIELD	DEFINITION	SOURCE	CALCULATION	OWNER
applicant_id	Persistent identifier — joins every source.	intake	UUID	Ops
deployability	Path to a working pilot, 0–16.	essay · deck	AI rubric · cited	Panel
pilot_traction	Evidence of pilots or trials, 0–15.	essay · ref	AI rubric · cited	Panel
rubric_total	Sum of six dimensions, 0–80.	computed	SUM(weights)	Ops
qualified	Total ≥ threshold (50).	computed	total ≥ 50	Ops
milestone_y1	Year-one milestone hit (admits only).	follow-up	survey + verify	PO

One owner per field — panel, program officer, ops.

Every score cited — the paragraph behind the number is on the record.

One **applicant_id** ties intake, scoring, and year-three follow-up together.

Defined once, cited forever

The dictionary is the contract between the rubric and the report. Because every field has a definition, a source, and an owner before the cycle opens, the cohort report at year-end is a query — not a forensic reconstruction.

STAGE 04 · TRANSFORMATION

Four transformations every application runs.

Every application lands as a stack of raw inputs — an essay, a recommendation, a budget PDF, a demo-day pitch. The Intelligent Suite runs four transformations on every one, turning unread material into scored, cited, comparable evidence.

```
// essay · paragraph 3
"We doubled enrollment among first-
generation students and shifted from
tutoring to project-based learning. Two of
our alumni are now teachers at the same
school."
```

→

```
theory_of_change 4.4 / 5
evidence: ¶3 s4-5
quant_anchor: "doubled"
confidence: high
status: pending confirm
```

Essay → rubric score. The reviewer sees the score, the supporting paragraphs, and the option to override. Any override is logged with reviewer ID and timestamp.

```
// recommendation letter
"I have worked with Maria since 2019. She
rebuilt our intake process, managed two
junior staff, and handled three difficult
board transitions with equanimity. I do not
write this kind of letter often…"
```

→

```
signal: strong recommender
tenure: 6 yrs
verifiable_claims: 4
calibration_flag: "do not write often"
status: flagged for read
```

Recommendation → signal. The AI doesn't score the letter — it pulls out the verifiable claims and flags the unusual language. The rubric score stays a human call.

```
// budget PDF · 4 pages
Personnel $84,500 · Fringe (22%) $18,590 ·
Travel $4,200 · Sub-grants $25,000 ·
Indirect (15%) $12,675 · Evaluation $8,000
— Total $164,965
```

→

```
eligibility: PASS
under_ceiling: ✓
indirect_in_policy: ✓
flag: sub-grants need Schedule R
flag: eval line absent from narrative
```

Budget PDF → eligibility check. Issues surface on Tuesday, not at panel. The program officer asks the applicant to clarify before the panel reads — not after the score is locked.

THE ENGINE ROOM

Four scopes. Four operations. One cohort.

The Intelligent Suite is Sopact's name for transformation in practice. Every AI call is a **scope** — what it operates on — crossed with an **operation** — what it does. For application work, the scopes are cell, row, column, grid; the operations are rubric, join, compare, summarize.

SCOPES · WHAT THE AI OPERATES ON

Cell	Row	Column	Grid
one score · one essay	one applicant · all dimensions	one dimension · all applicants	the cohort · across cycles

OPERATIONS · WHAT THE AI DOES

Rubric	Join	Compare	Summarize
score against criteria	essay to letter to budget	vs. mean · vs. last cycle	roll-up · drift report

Scope × operation, on a normal Tuesday

Score every pitch deck on Deployability (**cell × rubric**) → join the score to recommender signals on the same applicant (**row × join**) → compare Reviewer B's column to the panel mean (**column × compare**) → summarize the borderline pool by track for the chair (**grid × summarize**). Four calls, one dataset, defensible answers.

CALIBRATION

Reviewer drift, surfaced mid-cycle — not at year-end.

Every panel has drift. One reviewer scores high on the climate track; another marks down on writing quality at the end of a long week. In Application Management, drift becomes visible in the export — too late to fix this cycle. In Application Intelligence, drift is a Tuesday signal.

APPLICATION MANAGEMENT

Drift shows up in the cohort export at the end of the cycle, after the committee has already met. The fix is for the next cycle. The fix for this cycle is to defend the shortlist in the boardroom.

APPLICATION INTELLIGENCE

The drift signal surfaces on day 9 of 14. The panel chair sees it Tuesday and recalibrates. The committee does not vote on a cohort that has already drifted.

HOW THE WORK GETS DONE

A new tiebreaker Tuesday. Re-scored cohort by Wednesday.

Application Management charges for every adjustment — every new track, weight change, status email, new role. Application Intelligence treats configuration as conversation. Describe the change in language and it takes effect.

THREE CONFIGURATIONS THAT USED TO BE PROJECTS, NOW A SENTENCE

Route

Send the application to the next stage, panel, or person when a condition is met.

> "Route any application scoring 65+ on the Physical AI in Sports track to the interview round, assign the regional lead."

Re-weight mid-cycle

Adjust the rubric after the first hundred reads — the whole pool re-scores on the new weights.

> "Raise Deployability weight from 16 to 20 and reduce Defensibility from 15 to 11. Re-score the open pool."

Tiebreaker change

Change the tiebreaker before the panel meets; the new rationale appears on every record at the threshold.

> "Change tiebreaker from PA Ecosystem to First-Generation Founder. Apply within 4 points of threshold."

Why this changes the deadline math

Project Olympus needed one week to launch and three to decide. The old category needed three months to configure and two more weeks to reconfigure once the team saw the first hundred applications. Configuration-as-conversation closes the gap between what the panel learns and how the system responds — without an integrator, without a re-platform, and without missing the deadline.

STAGE 05 · REPORTS

Four report shapes, tied to application work.

The parent guide named the four report shapes — missing, unusual, comprehensive, aggregate. In application work, each has a specific job and a specific audience. Naming the shape tells you what the audience is really asking for.

MISSING

What didn't arrive that should have

The recommendation letter that hasn't shown up, the budget PDF still missing, the year-one survey nobody returned. Run continuously — surface gaps before the panel reads, not after.

UNUSUAL

What looks off vs. the cohort or last cycle

Reviewer B at +18% on the climate track. A score that swung after re-weight. A demographic distribution drifting from target. Anomaly against history or the group.

COMPREHENSIVE

The longitudinal view of one applicant

Everything known about #PA-1138 — first cycle to today. Intake essays, every comment, decision rationale, year-one to year-three. “Why this 40 and not those 40” in two clicks.

AGGREGATE

The longitudinal view of the cohort

What predicts success across cycles, where the portfolio is trending, which intake essays scored above 4.0 and then missed year-one milestones. The funder's annual report as a query, not a CSV merge.

WHY THIS MATTERS FOR THE BOARD CHAIR

Missing and unusual reports run continuously, so problems surface before committee. Comprehensive answers the question about any individual applicant — including the borderline ones the board chair will ask about. Aggregate answers the question the funder asks at the annual meeting. Same data, four shapes, four audiences, one query each.

— THE SIGNATURE COMPARISON

Five questions, asked on a normal Tuesday.

Eighty to eighty-five percent of an application-review team's work sits in these five questions — not in the year-end cohort report. The platform that answers them on Tuesday is the one that wins the program.

SUBMITTABLE · WIZEHIVE · AWARD FORCE **Application Management**	SOPACT **Application Intelligence**
"Did anyone read #447, and what did they think?" — **Probably not.** Three reviewers read the first 60. The other 440 are scored by the team that ran out of Thursday.	**Yes. Here's where the score came from.** Every essay read overnight, scored, paragraphs pulled. The borderline tab shows the 97 that need judgment — including #447. None of the 500 unread.
"Is Reviewer B drifting on climate?" — **Pull the export, run a pivot.** Drift shows up after the committee has met. The fix is for next cycle.	**Yes — by 18%. Calibration recommended before Friday.** The drift signal surfaces mid-cycle, broken out by reviewer, track, and dimension.
"Why this 40, not those 40?" — **The score is on file; the reasoning is in someone's head** or a margin comment. Reconstructing one rationale takes an afternoon.	**Pull up the rationale on any applicant on the bubble.** Per-dimension scores, the paragraphs behind each, panel comments, COI exclusions. #40 vs #41 is a click.
"Did this person apply in 2023? On another program now?" — **Maybe. Check the other system.** Re-applicants detected by name match — fails on married names, transliteration, email changes.	**Yes — applied to Cohort 2, declined, now Year 1 of the fellowship.** The applicant's history holds across programs on one record.
"What happened to Cohort 1? The funder wants outcomes next week." — **Two to four weeks of staff time.** Selection in one system, follow-up in another, comms in a third.	**Pulled. Demographics, milestones, year-one survey on the same row.** The cohort report is one query: who was admitted, what they said at intake, what year one said.

5

The shape of the answer

The five questions are where the program lives. The platform that answers them **on Tuesday with the citation trail attached** earns the next program inside the foundation. The one that answers them **after committee** earns the renewal conversation but loses the next program.

AFTER THE DECISION

After “congratulations.” The record does not reset.

Every application platform built before AI quits at the same place — the award letter. The applicant becomes a CSV export; the post-decision work moves to a different tool with a different ID. The record should continue past the decision, not start over.

01 APPLY

Intake & rubric

- Smart form
- Essay · pitch deck
- Recommendations
- Citation trail attached

02 DECIDE

Panel & rationale

- Borderline pool
- Drift recalibrated
- Tiebreakers applied
- Rationale on record

03 YEAR 1

Onboard & baseline

- Training kicks off
- Pre-program baseline
- Interview transcripts
- Commitment captured

04 YEAR 2

Milestone check

- Between-session check-ins
- Progress vs. promise
- Pattern emerging

05 YEAR 3

Outcome & loop-back

- Alumni reflection
- Predictive selection
- Next cycle smarter

5 yrs

One persistent **applicant_id** carries Application #PA-1138 from first essay through the year-three alumni reflection. Re-applicants in 2029 arrive with the full history attached. **Cohort 3 benefits from what Cohort 1 and 2 actually did.**

— THE ECONOMICS

The real cost per yes, compared.

Comparing license fees is misleading. The honest number is the cost per defensible yes — license plus configuration plus reviewer hours plus cohort-report reconstruction, divided by the count of decisions actually defensible to the board chair.

APPLICATION MANAGEMENT · PER CYCLE The real cost	APPLICATION INTELLIGENCE · PER CYCLE The real cost
Configuration time. 2–3 months at first launch, 4–6 weeks for each new program. Often outsourced.	**Configuration time.** Days. The rubric is described, not built; new programs inherit the pattern.
Reviewer hours. 12 apps/hour, fatigue at #30, the first forty read instead of the best forty.	**Reviewer hours.** Spent on the 97 borderline cases, not the 363 cleared. Two-to-three-times productivity.
Cohort-report reconstruction. 2–4 weeks of staff time every cycle. Nobody fully trusts the joined record.	**Cohort-report reconstruction.** Zero. The cohort report is one query against the records.
Cost of a wrong shortlist. Hidden, recurring. The applicant nobody read is the founder you missed.	**Cost of a wrong shortlist.** Quantifiable and small. Every essay is read; the shortlist is citation-backed.
License. Visible. The smallest line in the real number.	**License.** Comparable — and now the only line that resembles the legacy total.

THE REFRAME IN ONE LINE

Application Management priced for the workflow and gave you the reviewer matrix. Application Intelligence prices for the reading itself — and the reading does not get tired at application #30.

— THE LONG VIEW

Why selection sharpens with every cycle.

Application Management starts every cycle from zero. The same selection mistakes repeat. The applicant who dropped out of Cohort 1 reapplies in Cohort 3 — and nobody knows. Application Intelligence works the other way.

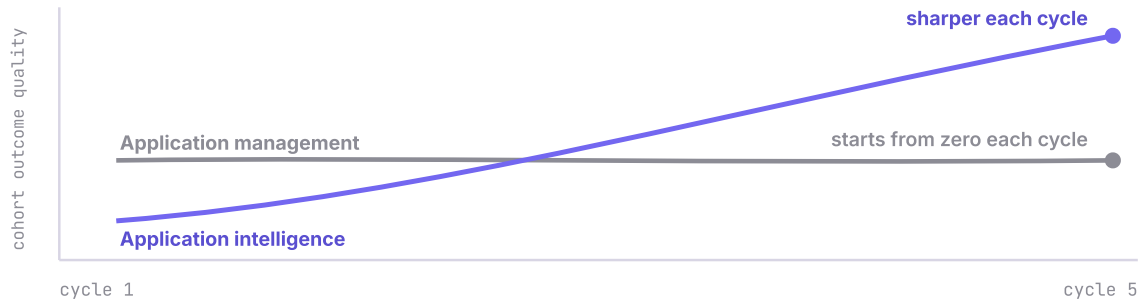

1 **Past outcomes improve future scoring**

Cycle 3 learns which intake-essay signals correlated with cycle-1 admits who hit year-one milestones. The rubric inherits the evidence, not just the language.

2 **Re-applicants arrive with their history**

The applicant who applied in 2023 and now applies in 2026 shows up on the same record. The panel sees prior essays, the prior decision, the prior reasons.

3 **The dictionary grows with the program**

New dimensions added in cycle 4 are still defined — and still cited — when the year-three alumni survey closes the loop in cycle 7.

4 **Selection bias surfaces, instead of compounding**

The bias audit is part of the cohort report. Trends across cycles — geographic clustering, institutional over-representation, demographic drift — are visible before they become uncorrectable.

WHERE TO GO NEXT

Where this guide sits in the library.

This guide is one of six use-case books in the Sopact Intelligence Library. Each moves a field from case management to intelligence, inheriting the same vocabulary and five-stage spine from the parent.

The parent guide

Beyond Case Management

The foundational guide every title below inherits from.

BOOK 02 · YOU ARE HERE

From Application Management to Application Intelligence

Stakeholder. apply → admit → alumni — scholarships, grants, fellowships, accelerators, awards, contests.

BOOK 03

From Grant Management to Grant Intelligence

Stakeholder · Partner. apply → report — a grant portfolio that reads its own quarterly reports.

BOOK 04

From Supplier & ESG Case Management to Risk Intelligence

Risk. screen → monitor — supplier audits and ESG diligence end to end.

BOOK 05

From Investment Case Management to Portfolio Intelligence

Portfolio · Stakeholder. due diligence → exit — IRIS+ and the Five Dimensions.

BOOK 06

From Learner Case Management to Learning Intelligence

Stakeholder · Portfolio. pre → mid → post — proving change across a learning cohort.

BOOK 07

From Nonprofit Case Management to Program Intelligence

Stakeholder. the whole spine across a multi-program nonprofit.

SUB-USE CASES THIS BOOK COVERS

Scholarship
essay → award → alumni

Grant
LOI → proposal → report

Fellowship
apply → interview → cohort

Accelerator
pitch → cohort → demo day

Award
nominate → judge → laureate

Contest
submit → round → winner

WHEN YOU'RE READY

Bring your last cycle. We'll score it in sixty minutes.

One cohort intake. One rubric. One round of reviewer scoring you've already done. We'll walk through how it would live as one record per applicant, what the AI would pull out of the long-form fields, and what the cohort report would look like coming out of one place.

A 60-minute working session

From a queue of four hundred unreviewed to a defensible, citation-backed forty.

Every essay read on arrival. Drift caught on Tuesday. The rationale already attached when the board chair asks why this 40 and not those 40. The shortlist isn't the first forty your team had time to read — it's the best forty, and that is a system you can run on a deadline.

Book a working session →

sopact.com

> **The shortlist isn't the best forty applicants. It's the first forty your team had time to read — and that is a system problem, not a discipline one.**