


Reading a Grievance Register Against Itself

Stakeholder and grievance monitoring system · Sociable Systems

A low grievance count and a high closure rate look like good performance. They are equally consistent with a mechanism nobody can reach. This system is built so that the second reading is as visible as the first — the exception layer runs before the dashboard, not after it.

What the register holds



The four checks that do the work

Check	What it does	What a comfortable report would have shown instead
Timeliness / SLA	Flags every individual case breaching its acknowledgement, investigation or closure window.	A mean resolution time, inside which the long tail of stalled cases is invisible.
Pattern	Clusters by location, category and repeat complainant.	A flat category count, in which a systemic issue reads as unrelated one-off complaints.
Completeness	Surfaces records that are incomplete, internally inconsistent or stalled without a state change.	A closure rate computed over whatever records happened to be complete.
Standards cross-reference	Ties each finding to the specific IFC Performance Standard requirement it bears on.	A compliance assertion with no route from the claim to the clause.

Three design decisions worth stating

- 1. History is retained, not overwritten.** A case that moved from open to closed to reopened is a different object from a case that was closed once, and only one of the two carries a warning. Current-state-only registers destroy that distinction on write.
- 2. Exceptions are counted, not filtered.** Records that fail a check stay in the population and appear as a number on the report. A pipeline that quietly drops what it cannot parse produces a cleaner dashboard and a less true one.
- 3. Reports are generated, not assembled.** Reporting outputs are produced from the validated register, so a figure in a board pack and a figure in the operational data cannot diverge without the register itself being wrong — which the completeness check would already have flagged.

The question this system is built to answer

Not “how many grievances did we receive and close.” That number is available from any register and settles nothing. The question is: **which people who had a grievance never appear in this data at all, and what in the intake path explains their absence?** The pattern and completeness layers are how that question is made answerable from operational data rather than from assurance.

Built on twelve years of resettlement data operations under IFC Performance Standards. Related work: the Evidence Chain Read (sociable.systems/advisory/evidence-chain-read) and Grievoice, multilingual grievance intake (sociable.systems/the-watchdog/grievoice).

Liezl Coetzee · Sociable Systems · sociable.systems · liezlc@gmail.com