

Stack Teardown — sample (synthetic)

Context

- **Industry:** B2B SaaS (fictional sample)
- **Headcount:** 50-150
- **Stage:** Series B
- **System of record:** Ledgerly

Stack map

1. Ledgerly → MailMotion

- **Sync:** one-way, nightly batch
- **Flow:** billing accounts → marketing contacts
- **Notes:** Billing is the only system with reliable identity; the nightly batch means marketing works with day-old truth.

2. StreamPipe → Glacier

- **Sync:** one-way, streaming
- **Flow:** product events → warehouse
- **Notes:** Glacier is the analytics system of record — but TrackPanel ingests the same client-side events in parallel, so there are two versions of every funnel number.

3. MailMotion → PushPilot

- **Sync:** none (manual CSV, ad hoc)
- **Flow:** email audiences ↔ push audiences
- **Notes:** Two messaging tools, two contact stores, no sync. Duplicate and conflicting contact records are structural, not accidental.

Failure modes

No single customer identity across billing and engagement

Impacted: Ledgerly, MailMotion, PushPilot

A paying customer in Ledgerly can be two different contacts in MailMotion and PushPilot. Churn-save campaigns and dunning emails routinely target the wrong record.

Double-counted product events

Impacted: TrackPanel, StreamPipe, Glacier

TrackPanel and StreamPipe both fire client-side. Activation and funnel metrics disagree between the product team's dashboard and the warehouse — and both are defended in meetings.

Marketing operates on day-old billing truth

Impacted: Ledgerly, MailMotion

The nightly batch means upgrade/downgrade signals reach messaging up to 24h late; win-back sends fire at customers who already renewed.

Cost leakage

Overlapping messaging (MailMotion + PushPilot)

- **Range:** \$18,000 – \$36,000 / year
- **Assumptions:** both tools priced on contact volume; ~40% audience overlap held in both
- **Redundancies:** duplicate contact storage, two audience-builder teams-of-one

Redundant event ingestion (TrackPanel + StreamPipe)

- **Range:** \$24,000 – \$48,000 / year
- **Assumptions:** event-volume pricing on both; same client-side events shipped twice
- **Redundancies:** parallel pipelines for identical events

Manual reconciliation labor

- **Range:** \$30,000 – \$60,000 / year
- **Assumptions:** ~0.25–0.5 FTE across RevOps and finance; monthly board-metric fire drills
- **Redundancies:** spreadsheet exports bridging Ledgerly ↔ Glacier every close

Metric impact

The metric everyone argues about — activation rate — differs by up to 14 points between TrackPanel and Glacier because of double-fired client events. Until ingestion is unified, neither number is wrong; both are unaccountable.

Target architecture

Ledgerly stays the identity anchor. One event pipeline (StreamPipe → Glacier) with TrackPanel reading FROM the warehouse rather than ingesting independently. Messaging consolidates to a single contact store with push and email as channels, not silos.

Migration sketch

Phase 1: reconcile identities read-only and quarantine every ambiguous merge (no writes). Phase 2: dual-run messaging on the unified contact store with a completeness ledger proving every record's fate. Phase 3: cut over channel by channel behind human approval; decommission the duplicate pipeline last.

Hook

You have two versions of every customer and two versions of every number. The fix is not another tool — it is an accounting of the ones you have.