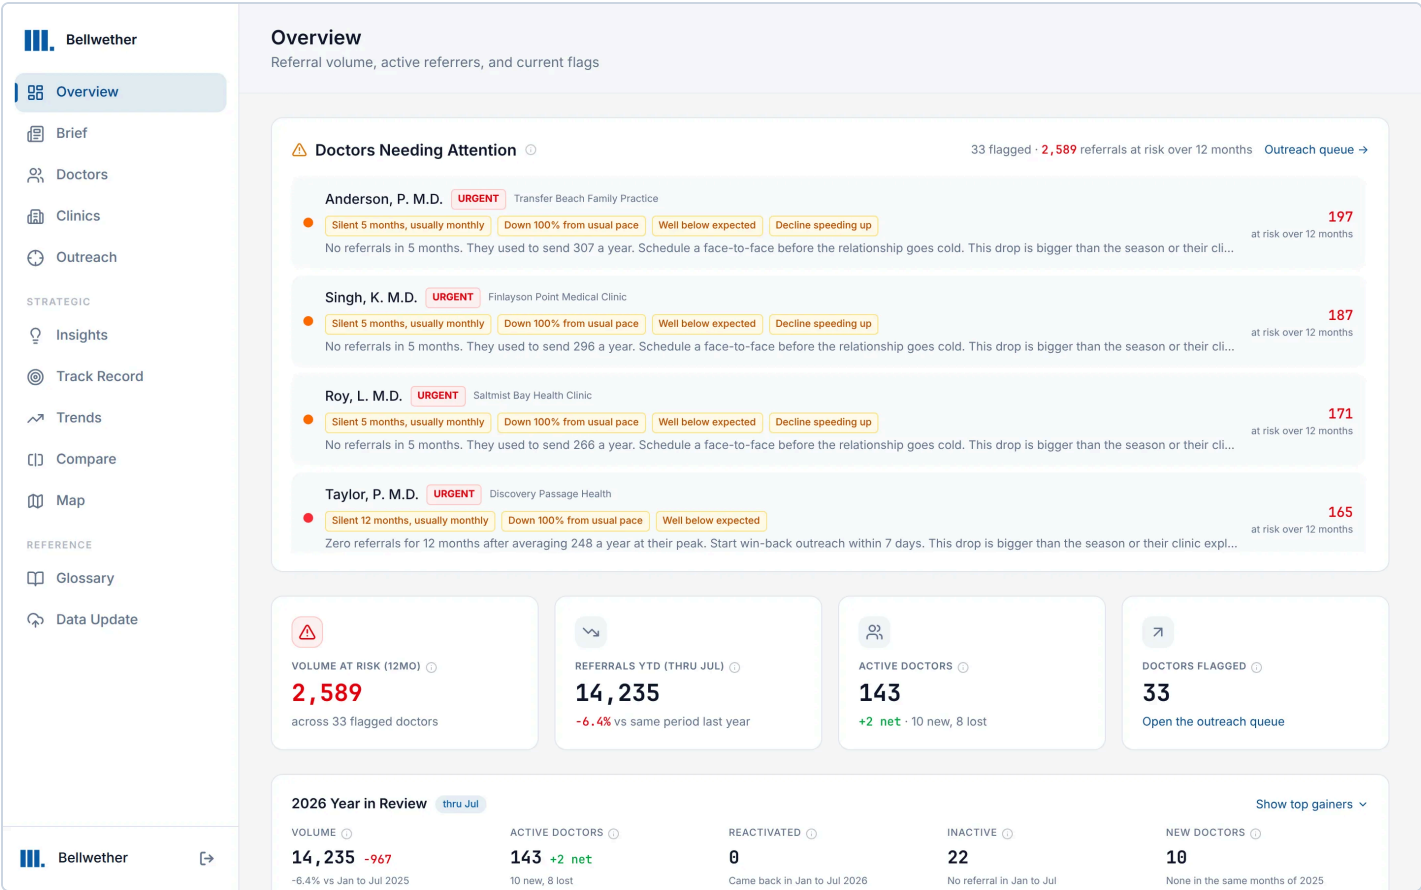


Referral early-warning analytics for community medical imaging

Bellwether reads a clinic's monthly referral export and identifies the referring clinicians whose volume is declining, while there is still time to act. It runs on one export a month from the billing or RIS system the clinic already uses, requires no integration project, and publishes how often its own scores have been right.

The name is the method: a bellwether wears the bell, and when the bell goes quiet the flock has already moved.



The overview screen. Every figure on it is computed from the clinic's own referral history. The lead figure prices the at-risk referral count in fee-schedule dollars where the fees display is configured. Screens in this document are captured from the demonstration deployment, which runs the full product on synthetic data.

WHAT IT NEEDS

One export a month, uploaded by drag and drop. Three fields are read: referring clinician, service date, and modality. A long-format billing export, for example Teleplan, is converted to monthly counts at upload, before anything is stored.

WHAT IT PRODUCES

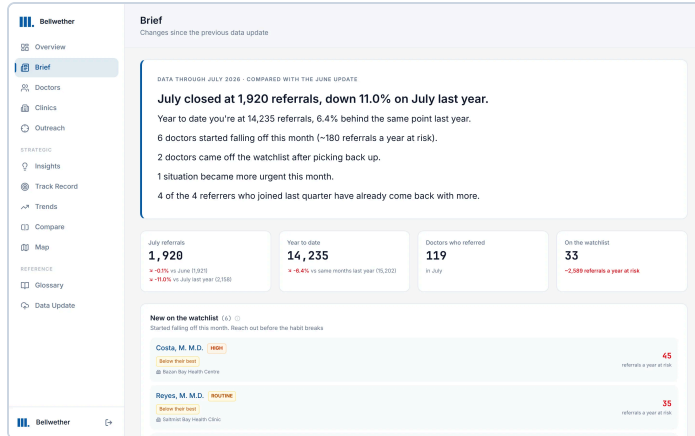
A ranked outreach queue with a stated reason for every name on it. A written brief after each data update. A six-month volume forecast with its error rate printed beside it. A track record of past calls, graded against what happened.

WHAT IT COSTS

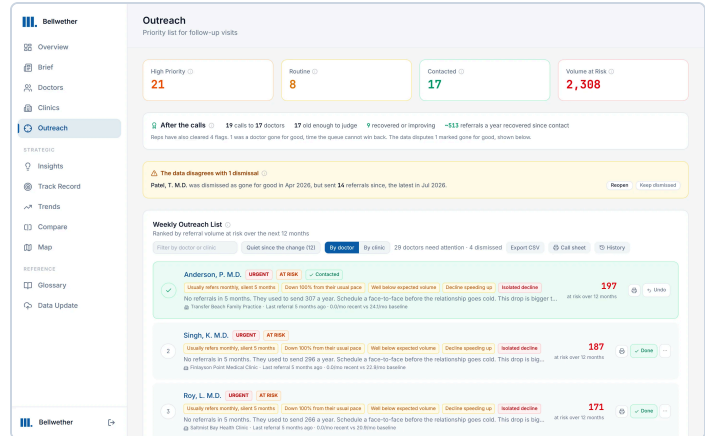
A flat monthly fee, from \$495 at founding rates, plus GST. It is not a share of billings and not a charge per referral, so the price does not change as referral volume changes. The full pricing table is on page 4, and setup is waived for founding clients through December 2026.

HOW A TEAM USES IT

After each data update, the brief states in plain sentences what changed: who entered the watchlist and why, who recovered, how the panel moved, and how earlier calls turned out. Between updates, the team works the outreach queue. Each row names the clinician, the signal that flagged them, the referral volume at risk over twelve months, and the office to visit. Actions are shared: when one person marks a clinician contacted, snoozed, or dismissed, everyone sees it.



The brief after a data update, in plain written sentences.



The outreach queue, ranked by referrals at risk and grouped by office.

THE SIX SIGNALS

- Cadence break.** Their own median gap sets the clock, and their own recurring quiet months raise the bar mid-break.
- Slippage.** Below their own baseline, two snapshots in a row.
- Deceleration.** Corroborating only, never flags alone.
- Below expected.** An exact Poisson or negative-binomial tail.
- Erosion.** Under half their best sustained pace, still falling.
- Sparse break.** Gap rhythm for the few-a-year sender.

Thresholds are absolute, judged against each clinician's own history and variability, never against whoever else is on the list that month. A clinic-wide dip is told apart from an isolated one before a name reaches the queue.

PAPER FOR THE VISIT

Every queue row prints a one-page visit briefing: why the clinician is flagged, the numbers, recent touches, and blank lines for notes. A separate per-office checklist records which offices hold the clinic's current fax number, requisition form, and EMR favourite.

AFTER THE CALLS

Each logged contact is graded on its own. The verdict waits for two full months of data after the call, and the comparison is against that clinician's own pace before the call, recorded at the moment of contact so it can never drift afterwards. Recovered, still declining, and went quiet are counted separately, and the results appear on the same screen the calls are made from.

Dismissals are kept and tallied by reason and by the signal that raised the flag, so the team can see which signals produce actionable calls. The product does not publish a comparison of contacted against uncontacted clinicians, because the team chooses whom to call and the two groups are not comparable.


THE MEASURES

MEASURE	RESULT	GATE
Ranking accuracy (AUC)	0.90	min 0.85
Watchlist precision	79%	min 70%
Calibration error	0.014	max 0.05
Volume at risk	$r = 0.88$	min 0.60
Forecast error	11.9%	max 20%

Measured out of sample on rolling cut-off dates in the production deployment (a community imaging clinic in Alberta, ~1,000 referring clinicians, five years of monthly history, running since 2025). Re-measured on your own history at setup, so you see the result before you commit.

CALIBRATION

The probability attached to each clinician is calibrated against outcomes: across clinicians given thirty percent, close to thirty in a hundred go on to stop referring. The measured calibration error is 0.014. The reliability curve below, shown inside the product as well, plots predicted probability against observed outcome rate; the figures printed on it are the demonstration panel's own.



The reliability curve: predicted probability against observed outcome rate, out of fold, here on the demonstration dataset.

THE TRACK RECORD

Each data update registers the system's live claims in a ledger before any outcome exists: who is flagged, at what recorded baseline, with what probability, and the six-month forecast with its band. An entry is never rewritten once a later month's update has arrived; two early entries were reconstructed from history with today's models and are labelled as such on screen. When enough months have passed, each entry is graded: recovered, still declining, or went quiet per clinician, forecast months against actuals, and predicted against realised volume at risk. Beside each precision figure the product prints the same test applied to every clinician it did not flag that was active and sending often enough to be eligible for a flag, so a reader can compare the flagged rate against a like-for-like base rate rather than judge a bare percentage.

HOW THE NUMBERS ARE PRODUCED

Every measure is computed out of sample. The system is rolled back to a past month, asked to score only what it could have seen then, and graded against what the following months actually held. In the production deployment, the same checks run automatically on every code change, each with the hard floor shown in the table, and a ratchet compares each run against the last accepted level so accuracy cannot drift down slowly between releases.

A month still in progress is detected at upload, held out of every model, and labelled on screen, so a mid-month export never reads as a cliff.

SETUP

- 1 Send one referral export from the billing or RIS system. Long-format claim files, for example Teleplan, are accepted directly and aggregated to monthly counts at upload.
- 2 Bellwether re-measures its accuracy on the clinic's own history. You see the result, against the same gates printed on page 3, before you commit.
- 3 The dashboard goes live under the clinic's own name, about two weeks after the first export. No on-site server, no access to patient records.
- 4 A monthly export by drag and drop keeps it current, with a full preview before anything is written. Weekly works too.

DATA HANDLING

Referring clinician, service date, and modality are the only fields used.

Long-format exports are aggregated to monthly counts at upload, before anything is stored.

Nothing patient-identifying is read, stored, or committed. There is no patient field in the system.

Spelling variants of one clinician are caught at upload and merged only when you confirm it.

A month still in progress is held out of every model and labelled on screen.

PRICING (FOUNDING RATES)

PLAN	FOUNDING	LIST
One location	\$495/mo	\$850
2 to 4 locations	\$995/mo	\$1,700
5 to 9 locations	\$1,495/mo	\$2,400
10 or more locations	quoted	.
Setup	waived	\$2,500

The fee is flat. It is not a share of billings and not a charge per referral, so the price does not change as referral volume changes. Founding rates for the first three clients, locked for as long as you stay. Setup is waived for founding clients through December 2026; from \$6,000 for a multi-site rollout. Ten or more locations are priced on locations and the size of the referring panel. Twelve-month term; a quarterly readout included. All prices are in Canadian dollars and do not include GST.

QUESTIONS

Does it need an integration? No. One export a month, uploaded by drag and drop, is the whole data path.

Who sees the data? The clinic's team, behind a password. The dashboard runs under the clinic's own name.

What if a flagged clinician has retired? The flag is dismissed with a reason, the reason is kept, and the tally is reported by reason, separately from the accuracy figures.

What does the demonstration show? The full product on synthetic data, so nothing confidential is on display.



The six-month volume forecast with its ninety-five percent confidence band and its average backtested error printed on the chart, here on the demonstration dataset.

ACCESS

Email samuel.putra101@gmail.com for demonstration access or to start with your own export. Requests are answered by a person within one business day, and nobody follows up with a pitch. The demonstration deployment is at bellwether-canada.vercel.app.