FIELD NOTE · INSTITUTIONS, GOVERNANCE & SYSTEMS

Every Team Was Green: A Field Note

A 30-day read-only count first, then an optional 90-day pilot, for giving one person the right to change a handover

> Operating detail for Every Team Was Green, for managers who want to count the repeat work their teams make for each other and, if it is big enough, take it to a peer team or test the executive option. The thresholds and clocks here are judgement, not evidence.

Lateral House · 22 November 2026 · lateralhouse.com/essays/the-last-mile-field-note/

How this was made. This field note was AI-drafted. The author directed it, shaping the argument, structure, sourcing and edits, rather than writing each sentence personally.

Operating detail for Every Team Was Green, which sets out the problem and the count. The executive arrangement near the end is an optional extension with no audited case, so every threshold and clock below is a design choice to change as needed.

Start Here: The 30-Day Count Pack

This is a read-only count: change nothing while you measure, and use a completed past month (a fresh month needs 30 days more). Pick one journey (the path a customer travels from first need to resolution across every unit) and give your analyst this query.

> **Query:** For [journey], take every case closed in [month]. Count the unique customers who made contact again on the same account, for any reason, within 30 days of that closure. Then read a random sample of 50 of those returning customers, one record per customer, and note whether each came back about the same problem and how many contacts they made about it. The headline number is the count times the share who did, as in the essay: the same customer, about the same problem, within 30 days of closure. The workload, the number of repeat contacts, is that figure times the average contacts per such customer. Report it as a rough figure, since a simple random sample of 50 has a worst-case margin of about 14 percentage points either way at roughly 95 per cent confidence, and report the any-reason count beside it. Where a customer had several cases closed in the month, start from the first and count them once. Use a past month whose 30 days of follow-up have finished; if you start from today, allow 30 more days after the month ends before

> the number is ready. Give the share of customers with a case closed in the month. For each repeat, record the last handover between units before it. If the systems are not joined, export two lists for one product and one month and join them on account ID in an approved analyst workspace.

The readout is one page: the count and share, with the definition used; a handover map showing how many repeats follow each point where work passes between units; a rough annual cost, arithmetic shown; and the limits, starting with customers who gave up or left.

If contacts cannot be linked to customers across teams, that is your first finding: read a sample of 50 second contacts by hand in the meantime. Matching is partial, because customers ring from other numbers and accounts, so check a sample of the linked contacts too.

Report the count beside previous months, since a small one swings by chance. Tag customers chasing an update separately: their remedy is being told what is happening.

Keep the count off every scorecard. Once a unit is judged on it, repeats get re-labelled as new problems and the number falls while nothing improves (Goodhart's law).

If Finance will not give you a cost per contact, build your own: average handling time times an agent's hourly pay plus overheads, labelled as your estimate. For most readers the count is the whole recommendation. Send the email below to the executive above all the teams, or to your own manager if you are further down.

> **Subject:** One number on [journey], 30 days, no changes
>
> [Name], I would like to measure one thing on [journey] from past records, changing nothing. From existing records, [analyst] will count customers who came back about the same problem within 30 days and map the handovers those repeats follow. It needs read access to [systems] and about a day of [analyst]'s time for the count, plus a few more to check the sample and map the handovers. You get a one-page readout: the number, the map and a rough annual cost. Could you approve the data access?

Price It, and What to Say

Put a price on the count. In an invented example, 10,000 repeat contacts a month (the workload, not the number of customers) at £6 a contact comes to £720,000 a year of staff time (10,000 × £6 × 12); substitute your own count and your own cost per contact. That is capacity released for other work, not cash saved, unless hours are cut or hiring is avoided.

If you need a sentence for your boss: "I'd like one number, how many customers came back about the same fault within 30 days last month. It takes an analyst about a day, and it changes nothing until we have seen it."

LATERAL HOUSE

The Second Number: Other Customers, Same Fault

Group by what customers wrote, not by the agent's category: a text-similarity query (for example trigram matching in Postgres) or an AI model that clusters contact text will propose the groups, and a person reads a sample of each cluster to confirm it. Use only approved systems, keep customer text inside them, and limit who can see it. Matching finds candidates; it does not replace linking by customer. The customer count misses a fault that fifty customers each report once. For the second number, group contacts by product, place and week, or by fault code if you trust it, and count how many arrived after a fix on the same fault was logged as done. It is rougher than the first number and the cause may sit anywhere in the chain, so treat it as a hint that sends someone to look, not as a count to score a team on.

Slow Fixes: The Problem Record

Some causes take months. A false date does more harm than a long one, and patching each customer one by one hides the cause behind closed tickets. Keep one open problem record per cause:

- **Owner.** One named person, who is not the person logging contacts.
- **Honest range and next update.** The team doing the work sets a range and a date for the next update, never the agent on the call. An unusually wide range is itself a flag. Keep the update date, so customers do not have to chase.
- **Workaround log.** Each patch is logged against the record, not closed as resolved.
- **Running count and cost.** Contacts so far and a cost priced once by Finance, so the owner cannot shrink or inflate it. It is the price of the delay.
- **Review date.** A date when someone decides whether the fix deserves more priority.

A contact chasing a promised date is a repeat contact, counted once against the journey whose team set the date. For the agent on that call, write a script for "no update yet" that says when the next update is due and what has been done meanwhile, and protect handling time for adding a contact to the record.

Taking It to a Peer Team

Show the numbers privately first, before they appear in any review. Say what is in it for them: the trial cuts their own second-call load, and a figure that worsens during it is read against the baseline you wrote down, not their target. Agree the stop condition in advance, for example stopping after two weeks if repeats on the journey have not fallen or if complaints have risen. Write down the baseline before the first day. If the count is low, say so and stop: the honest result is that this journey does not need the rule.

LATERAL HOUSE

Keeping the Count Honest

Any number that is watched gets gamed, and few people want to mark down a colleague. So the count is a flashlight, not a scorecard. To keep it that way:

- do not tie it to pay, bonus or ranking; use it to decide where to look;
- read it beside complaints and churn, since a falling count may only mean customers gave up;
- have someone outside the closing team read a sample of the cases each month, including new tickets raised soon after a closure, and share what it finds with the two peer managers before anyone above them;
- let reopening cost the agent nothing, so agents have no reason to open a new ticket to hide the link;
- expect the gaming to move, for example to transfers, “pending” states or other channels, and look there too.

Outside Telecoms

Elsewhere, define the same problem and a repeat locally.

SETTING	SAME PROBLEM	REPEAT
Bank	Same customer, same account and same complaint or payment category	Any further contact about it within 30 days, in any channel
Hospital	Same patient, same or related condition	Unplanned readmission or emergency attendance within 30 days of discharge, counted for the same or a related condition, not all-cause admissions
Council repairs	Same property, same repair type (leak, damp, heating)	A new request or complaint within 30 days of a job closed as complete; for damp, consider 90

An invented example. A council closes 1,200 repair jobs in a month; 216 properties, 18 per cent, report the same repair type again within 30 days. Most follow jobs a contractor closed after the first visit while a second trade, booked by another team, was still needed. The rule to change: a job stays open until the last trade has finished.

Hospitals offer a mixed analogy. In a 1999 trial of 363 patients aged 65 or over at two Philadelphia hospitals, nurse-led discharge planning and home follow-up cut readmissions within 24 weeks to 20.3 per cent, against 37.1 per cent for controls (Naylor et al.). A 2009 evaluation of 15 randomised Medicare care-coordination programmes found 13 with no significant change in admissions, one with fewer and one with more, and none with net savings; the authors concluded that coordination without a strong transitional-care component is unlikely to save money, while programmes with substantial in-person contact aimed at moderate to severe patients could be cost-neutral and improve some

LATERAL HOUSE

aspects of care (Peikes et al.). The evidence is mixed: coordination alone does not reliably reduce admissions or save money, and more intensive management of transitions has had better results. Neither study tested the handover approach in this note.

Running the Right and Its Two Protections

Peer agreement, as in the essay, is the first response. It lasts only while both managers stay in post, so write it down, or take it to the executive, before either moves. If the same dispute keeps returning, or nobody has the authority to settle it, and the number is big enough, the executive above all the teams can make the repeat work between teams one named person's job (the owner). The owner gets one written right: to change a handover rule, if the change is logged and can be undone. Two protections come with it: a baseline fixed before the trial, so a team whose figures worsen is judged against it, and a ten-working-day clock for objections. The executive can end any trial and carries any red figure for a few weeks. The owner works within existing delegated authority, change control and system-owner approvals, and replaces none of them. In operation:

- **The log.** Each change records the rule, the unit that owns it, the measure it is expected to worsen, the date and how to reverse it.
- **The baseline.** The data owner freezes each unit's figures for the 90 days before day 1. HR certifies before day 1, to day 90, that nobody is marked down for a logged change; Finance validates the cost figures.
- **The clock.** Ten working days run from a written objection. A disputed operational change runs as a trial unless the executive rules against it in that time.
- **The limits.** The right covers reversible operational handover rules only. A change that touches a legal, regulatory, safety, privacy or independent-control requirement needs the written approval of that control's owner first; silence never counts as approval, and the ten-day clock does not apply to it. Anyone may stop a change at once if it is harming customers, and it then goes to the executive.

The owner is usually a peer of the unit heads, so the rights are real only once the executive writes them into each head's objectives. Treat a change that quietly lapses as a dispute and put it on the clock; the data owner catches re-labelling by spot-checking "new" faults on accounts with one in the previous 30 days. If the executive will not rule against a head who ignores the log, the readout says the rights were not real.

If there is a project office, the owner does not replace it. Any project that changes a handover on the journey goes through the owner's log like any other change, so the two never issue conflicting rules.

LATERAL HOUSE

Choosing the Owner

Choose someone the unit heads treat as a peer, who has worked in at least two of the units and can read the ticket data unaided. Avoid the head of the unit with the most handovers, whose figure is likeliest to worsen. During the pilot a deputy takes the owner's routine work, and the executive says in writing that an honest readout, pass or fail, meets the objective.

What the Executive Gives Up

Rulings take time, and a ruling against a unit head is made in front of that head's peers. Rule on the written evidence and say why.

If the executive refuses outright, keep the read-only count running each month, method unchanged, and send it to the same reader. The refusal goes on record as a finding, and a number that keeps arriving is harder to set aside than a proposal.

Stage Two, Optional: A 90-Day Pilot

Run it only if all three gates are a yes; otherwise it fails for reasons unrelated to the idea.

1. **Exemption certified.** HR has signed that no unit is penalised for a logged change, against its own baseline at the start, and Finance has validated the cost figures.
2. **Executive pre-committed.** The executive has agreed in writing that a disputed operational change runs as a trial unless they rule against it within ten working days; control-sensitive changes are excluded, as above.
3. **Baseline readable.** The past 90 days of repeat contact can be read from existing records.

REPEAT CONTACT	Share of faults where the customer makes contact again about the same fault within 30 days
RESOLUTION PROXY	A passed line test and no repeat within 30 days. A customer who gave up also looks resolved, so a small monthly phone sample asks customers whether the service now works
DATA OWNER	A finance or quality analyst who is not judged on the result and holds the written definition of a repeat

Fix the definition of a repeat before day 1, and compare against a matched region that carries on as before. **Success** has three tests. Repeat contact is down by at least a fifth from baseline (a relative fall of 20 per cent or more). That relative fall is at least 20 percentage points larger than the relative fall in the comparison (for example 33 per cent against 5). And no customer segment named before day 1 is worse than its own baseline by more than a point. A pilot that passes the first two and fails the third is

LATERAL HOUSE

a conditional pass: it continues only with a fix for that segment and a date to re-test. **Stop** at day 90 if either of the first two tests is missed, if no logged change was made, or if disputes were left unruled past ten working days.

A Worked Example

Every number here is invented for illustration.

- **Baseline:** 6,000 faults in 90 days and 1,200 repeats (20 per cent); residential about 20 per cent, business lines 18 per cent; comparison region 20.
- **Week 1:** most repeats follow closures on a passed remote test. On day 5 the owner logs a change: no such closure after a repeat contact. Network operations objects; the executive rules for the change on day 13, and it is live by day 15.
- **By week 7, what showed up:** a cancelled visit now triggers a call-back, made by the agents who answer business lines, so business customers waited longer and rang again. Nobody had mapped that handover.
- **Day 90:** of 3,000 faults in the readout window, residential repeats fell to 12 per cent (300 of 2,500) and business-line repeats rose to 20 per cent (100 of 500). Overall, 400 repeats is about 13 per cent, a relative fall of a third, against 5 per cent (20 to 19) in the comparison. Network operations' closure time rose from about one day to two, not penalised and listed.
- **The decision:** the pilot passes the first two tests and fails the third, because business lines are two points worse than their baseline. It is a conditional pass, and the readout names the business queue as a harm. The executive continues, with call-backs moved to a separate slot, and the pilot does not count as a clean pass unless business lines are back to 18 per cent by day 150.

The Executive Ask

> **Subject:** A 90-day trial on one customer journey
>
> [Executive], our 30-day count on [journey] found [number] customers coming back about the same problem, worth about [annual figure] a year in staff time. I would like to run a 90-day trial in [region].
>
> Please give [person] the right and its two protections described in this note, in writing and in each unit head's objectives, with HR's certification signed before day 1.
>
> Cost: no new budget, about 45 working days over 12 weeks plus half a day a week from each unit. We compare against [comparison region] and stop at day 90 if repeat contact has not fallen by at least a fifth and by 20 points more than there, if a named segment is more than a point worse than its baseline and no approved fix and retest date exist, if no logged change was made, or if disputes were left unruled past ten working days.

LATERAL HOUSE

> Could we take 20 minutes this week to agree it?

If You Need Rigour

The join. Match contacts on account number (failing that, phone number); a repeat is a contact within 30 days of the earlier case's closure in the same fault category. It undercounts, and a shared phone can join unrelated contacts; say so and keep the method fixed. Matching joins personal data: use only access you already have, keep the matched list in the analyst's restricted store, and report counts, not names. If the systems share no key, sample 200 faults closed in one week and look each account up 30 days later, at baseline and day 90; only a large fall will show, and it is an exploratory check, not evidence of a 20 per cent effect.

Volume. At N repeats a week, ordinary weekly swing is about the square root of N, and a one-fifth fall should be at least twice that, so treat 100 repeats a week as a rough screening floor, not a power calculation: it ignores the denominator, the size of the comparison group and the customer mix, and it does not show that a one-fifth fall can be detected. A one-point segment rule can also raise false alarms in a small segment, so read it with the segment's size. Repeats also cluster around shared events such as storms or cabinet outages, so real week-to-week variance is larger than the square root suggests. Set decision-grade sample sizes and stopping rules with a statistician before the pilot starts; do so for any stop decision or bonus.

Comparison. Match the comparison region on baseline repeat rate, product and customer mix, fault types and volumes, and the trend over the 90-day baseline; a similar headline rate alone is not enough. Without a match, stagger the start across two areas a month apart. Before-and-after alone is weakest, since a bad baseline may drift back to normal unaided.

Unit. The 30-day count counts customers; the pilot counts faults, because a fault can be followed for its own 30 days and compared between regions. A customer with two faults counts twice, so say which unit each figure uses.

Timing. Exclude a declared network incident from both regions. The change goes live by day 15, so every fault counted was reported after it began. The day-90 readout counts faults first reported on days 15 to 60, each followed for its full 30 days; later faults go into a day-120 readout. The baseline spans 90 days so that it covers an ordinary run of weeks, while the readout window is shorter because each fault needs its 30 days of follow-up; check that volumes and seasons are not very different between them. Whether gains last is not yet known.

The clock at a glance.

- Before day 1: baseline fixed, segments named, HR certification signed.
- By day 15: the change is live.
- Days 15 to 60: faults are counted, each followed for 30 days.
- Day 90: readout. Day 120: readout for the later faults.

LATERAL HOUSE

- Day 150: re-test of any segment that was a conditional pass.

Sources

Naylor, Mary D., Dorothy Brooten, Roberta Campbell, et al. “Comprehensive Discharge Planning and Home Follow-Up of Hospitalized Elders: A Randomized Clinical Trial.” JAMA, 281(7) (1999), pp. 613–620. doi:10.1001/jama.281.7.613.

Peikes, Deborah, Arnold Chen, Jennifer Schore, and Randall Brown. “Effects of Care Coordination on Hospitalization, Quality of Care, and Health Care Expenditures Among Medicare Beneficiaries: 15 Randomized Trials.” JAMA, 301(6) (2009), pp. 603–618. doi:10.1001/jama.2009.126.

Goodhart, Charles A. E. “Problems of Monetary Management: The U.K. Experience.” In Papers in Monetary Economics, Vol. I. Reserve Bank of Australia (1975).

Cite this field note. Lateral House (2026). “Every Team Was Green: A Field Note.” Lateral House, 22 November 2026. https://lateralhouse.com/essays/the-last-mile-field-note/

© 2026 Lateral House. Published at lateralhouse.com. This PDF is for personal and educational reading; please link to the web version when sharing.

LATERAL HOUSE