LOW EFFORT FOR HOUSING · FIELD GUIDE

WHAT MAKES A GOOD MEASURE?

The Low Effort Guide to Measuring Housing Repairs Performance

A practical field guide for housing leaders: see beyond the dashboard, understand current capability and improve the way repairs really work.

79%
of non-emergency repairs in England completed within landlords' own target timescales, 2024/25

97%
of routine repairs 'within target' on one organisation's dashboard

216 days
the average time those repairs actually took, from first report to fixed

Stuart Corrigan
Descartes Consulting · loweffort.com/housing



CONTENTS

Ten short chapters, a worked example from a repairs van, and a one-page worksheet to take into your own service.

1	The Green Dashboard	3
2	How We Learned to Manage by Numbers	5
3	Same Organisation. Two Different Stories.	7
4	Principle 1 — Establish Current Capability	8
5	The Eight Measures a Repairs Leader Needs	14
6	Principle 2 — By What Method?	16
7	How to Study the Method	18
8	A Worked Example — What Is in the Van?	22
9	The Final Piece — System Conditions	24
10	Your Monday-Morning Starting Point	26
·	A One-Page Field Worksheet	28

CHAPTER 01

THE GREEN DASHBOARD

1

On 4 November 2025, the Regulator of Social Housing published its second annual set of Tenant Satisfaction Measures, covering April 2024 to March 2025.

During that year, landlords reported completing more than 11 million non-emergency responsive repairs. Of those, **79% were completed within their target timescales**. [1]

At first glance, 79% sounds reasonably impressive. But is it?

After all, we're talking about repairs to people's homes. Sometimes the problems are minor. Sometimes they're serious. And many of the people affected are vulnerable.

So before deciding whether 79% represents success or failure, there's a rather important question to answer.

Who decides what counts as completing a repair on time?

Suppose I decided to run a marathon and decided eight hours constituted a world-class performance. Suddenly becoming a world-class marathon runner wouldn't seem quite so difficult.

So before we award that 79% a gold star, perhaps we should ask who decided where to set the finishing line.

It can't have been the tenant. Surely they'd have demanded more.

And we wouldn't let those doing the work set their own targets, would we?

Except that, in an important respect, is exactly what we've done. Landlords define the applicable repairs timescales used in the measure. Those timescales must also be reported; the regulator does not pretend that a high percentage is meaningful without that context. [1, 2]

Marking our own homework is not a new phenomenon.

We can go back to 2003. To something even more important than repairs. Life and death.

EIGHT MINUTES TO REACH WHOM?

The ambulance service was measured against an eight-minute response-time standard for life-threatening emergencies. The number sounded clear. But the evidence presented to Parliament described a rather less clear reality. [3]

An ambulance could reach a tower block and have its arrival recorded while the patient was still several floors above. In different trusts, the clock could also start at different stages of handling the call — sometimes differing by as much as three minutes. [3]

Imagine two identical emergency calls. In one trust the clock starts promptly. In another it starts three minutes later. Both vehicles record an eight-minute response, although the callers have waited different amounts of time.

That comparison is an illustration, not a documented pair of patients.

The target hadn't necessarily changed. The meaning of the measurement had.

The housing regulator faced a similar question. During consultation it considered alternatives, including repairs completed right first time and the average days taken to complete a repair. It chose the percentage completed within each provider's target timescales, citing the difficulty of consistent definitions for some alternatives and the value of transparency about published standards. [2]

The regulator also publishes tenants' perceptions of repairs and of the time taken. Those measures are valuable. But they don't replace studying what actually happens between a tenant's first report and the moment the problem is fixed.

Which brings us to a rather more important question.

What makes a good measure?

Not simply one that tells us whether we've met a target.

One that helps us understand what's actually happening, why it's happening, and what we might do to improve it.

And if we were designing a housing repairs service around the needs of tenants, what would we measure?

More importantly, **how would we use those measures to make the service better?**

And before we go any further, there's one more assumption worth examining.

Where did we get the idea that setting a target was a good way to improve performance in the first place?

CHAPTER 02

HOW WE LEARNED TO MANAGE BY NUMBERS

2

Imagine the boardroom of a housing association.

A large square table, packed with directors and team leaders.

In front of them sits last year's performance report. Alongside it, a letter from the regulator calling for improvement.

On the television screen, a single slide.

NEXT YEAR: 95% FIRST-TIME FIX.

Everyone nods.

The tenants' association is unhappy. Nobody wants more scrutiny. And 95% sounds like a respectable ambition.

But there are two questions nobody seems to be asking.

Why 95%?

Why not really push the boat out and go for 100%?

And perhaps more importantly:

By what method?

What, exactly, are we going to change to achieve it?

This boardroom is an illustrative scene, not a claim about any identifiable housing association or a regulatory instruction to achieve 95% first-time fix.

THE PSYCHOLOGIST AND THE GOAL

Let's go back to 1935.

A British psychologist named Cecil Mace was interested in a deceptively simple question.

How do we get people to perform better?

If we give them a goal, will they?

In experimental work published that year, Mace investigated the effects of goals on task performance. Later accounts describe a complex computation task and a comparison between instructions to do one's best and goals related to previous performance. His work became an early landmark in goal-setting research. [4]

And the idea was appealing: specific goals could influence how well people performed a defined task.

But there was something else about this work. Something that becomes enormously important when we move from an individual task to an entire repairs service.

> Participants could influence the results through their own actions.

A plumber can influence how carefully he works. But can he determine whether a caller diagnoses the problem correctly? Whether the right parts are on the van? Whether the schedule allows enough time? Whether he is authorised to do the necessary work?

Having influence over your own task is not the same as controlling the system in which you work.

Mace did not invent targets, nor did he claim that every organisational outcome was within a worker's control. The problem begins when we forget the difference between an individual goal and the performance of a whole service.

JAPAN, 1950

In post-war Japan, industrial leaders faced a rather different question.

With factories damaged, resources scarce and an economy rebuilding, how could Japanese manufacturers compete with industrial giants such as Ford?

An American statistician, W. Edwards Deming, offered a radically different way of thinking. Especially about quality and measurement.

His challenge to management was not to assume that poor results could be fixed by demanding more from workers. He urged leaders to understand and improve the system that produced those results. [5]

Toyota's own history describes engineers installing statistical control charts in its machining operations in 1950. The charts were initially greeted by indifference, resistance and anxiety. But when defects appeared, the question became what in the process was causing them, and what countermeasure might prevent their recurrence. [6]

This was not a company without goals. It was a company learning to study how its work actually behaved.

> Not what leadership wanted the system to achieve.

> What the system was capable of achieving today.

Stop there. Because this distinction is fundamental.

CHAPTER 03

SAME ORGANISATION. TWO DIFFERENT STORIES.

3

Let's return to housing repairs.

On the dashboard of one organisation, the reports were impressive: **97% of routine repairs recorded as completed within target**. [7]

By the organisation's own measure, things appeared to be going rather well.

But John Little and the team asked a different question:

From the moment a tenant first reported a problem, how long until it was actually fixed?

The reporting system couldn't readily answer it. So the team reconstructed the histories of **560 repairs**, one at a time. [7]

The average?

216 DAYS.

Not 216 hours.

216 days.

The book records an upper control limit of about **1,081 days**, with some repairs taking around **2,000 days**. One leaking bay window generated **24 visits**. Another kitchen problem generated 26. A bathroom problem generated 17. [7]

And every month the reports were saying 97%.

Neither number necessarily had to be wrong.

The dashboard was reporting performance against the organisation's definitions and targets. The investigation reconstructed the life of the tenant's problem.

The dashboard told management whether specified jobs had complied with their deadlines.

The repair histories revealed what tenants were actually experiencing.

And, crucially, they gave us somewhere to begin improving.

We hadn't started by deciding what the average completion time ought to be.

We'd started by discovering what it actually was.

> Not what management wanted. What the system was capable of achieving today.

CHAPTER 04

PRINCIPLE 1 — ESTABLISH CURRENT CAPABILITY

4

Before deciding what performance ought to be, discover what your system is actually achieving.

This is a practical guide. You don't need permission to begin asking better questions. You do need to define your data carefully.

STEP 1. START WITH WHAT MATTERS TO THE TENANT

A tenant reports a leaking tap. What do they want?

Not a job number. Not an appointment recorded within a target. Not a line in a contractor's completion report.

> They want the tap fixed.

So start with end-to-end completion time: the elapsed calendar time between the tenant's first report of that problem and its genuine resolution.

A repair reported on 1 March and resolved on 21 March took **20 calendar days**. Not two hours because that was the time on site. Not five days because that was the appointment target.

Twenty days. That's what the tenant experienced.

STEP 2. DEFINE WHAT COUNTS AS ONE REPAIR

This is more important than it sounds.

One problem is not necessarily one job. A single leaking tap may create ten work orders. Counting each order as a completed repair can produce an immaculate-looking report while the tap continues to drip.

Agree the unit of analysis before you touch the spreadsheet:

- **Start:** the first contact about the identifiable problem.
- **Finish:** the point at which the problem is genuinely resolved, not merely administratively closed.
- **Reopened or repeat work:** link it back to the original problem where it represents unresolved work, rather than treating it as a fresh successful repair.
- **New unrelated problem:** treat it as new demand; don't force distinct issues into the same case.
- **Emergency, routine and specialist work:** analyse separately where their demand and conditions differ.
- **Still open:** retain separately in an age-of-outstanding-repairs view; do not silently drop them.

Document the definitions. If your original records cannot connect jobs to one underlying problem, reconstruct a representative sample manually and be explicit about the limitation.

STEP 3. COLLECT A MEANINGFUL SAMPLE

For each completed case collect the original report date, the genuine completion date, a unique case identifier, repair type and, where possible, the sequence of visits and rework.

Start with at least 100 consecutive completed repair histories if you can. That's a useful exploratory starting point, not a universal statistical minimum. More observations generally give a richer picture. The investigation in Housing Repairs ER used 560.

Aim for **at least 12 months of history** to reveal seasonality and sustained changes. If you have less, begin with it rather than postponing the investigation.

Do not cherry-pick cases. Check that your selection represents the work being studied. Check the open backlog alongside closed cases so you don’t mistake clearing easy work for better overall performance.

STEP 4. PUT REPAIRS IN THE ORDER THEY WERE COMPLETED

Record the elapsed calendar days for each case. Sort by **actual completion date**, earliest to latest, preserving the original report date for each case.

Repair	First report	Genuine completion	Elapsed calendar days
A	1 March	8 March	7
B	3 March	17 March	14
C	4 March	25 March	21
D	10 March	26 March	16

Illustrative example — not extracted from the 560-repair records.

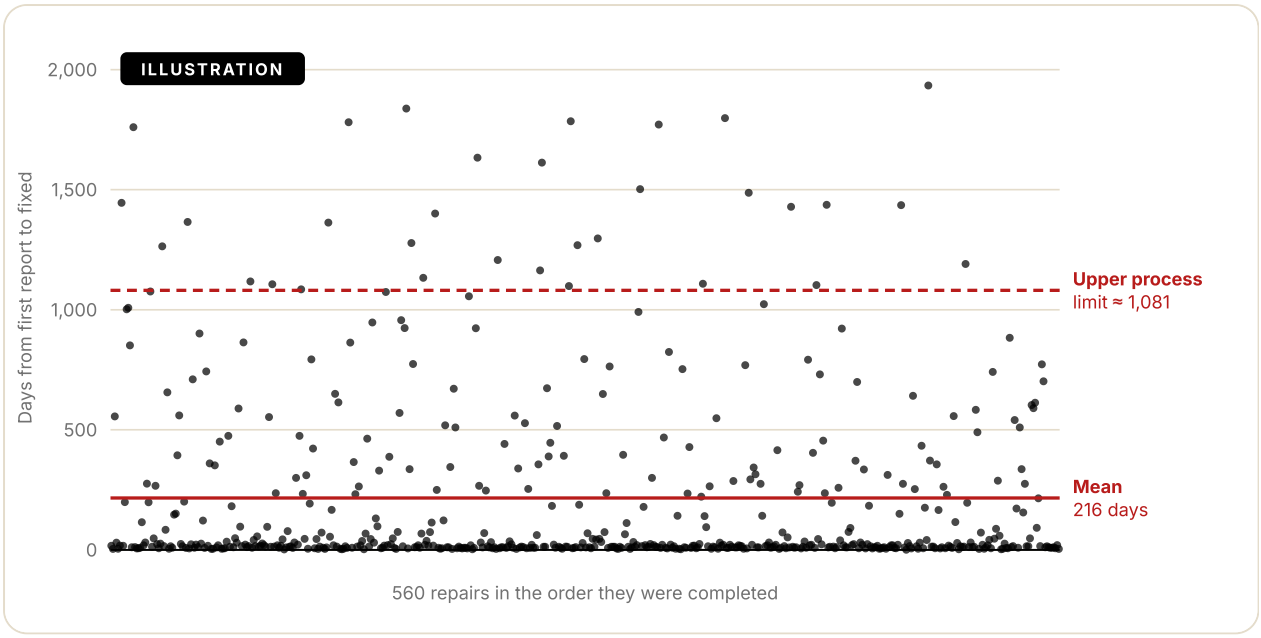
Don’t just produce one monthly average. **Keep the individual observations in time order.** Performance changes over time, and the sequence carries important information.

STEP 5. MAKE THE DATA VISIBLE

Put the repairs in completion-date order on the horizontal axis and elapsed days on the vertical axis. Each repair becomes a point.

Some will be quick. Some will be slow. A few may be extraordinarily slow. The dots let you see what a percentage hides.

Then draw a horizontal line at the mean. In the case in the book, that mean was 216 days.



Illustration, not the case data. Each dot is one simulated repair, set to match the figures published in Housing Repairs ER: 560 repairs, a mean of 216 days, an upper process limit of about 1,081 days and a few repairs near 2,000 days.

STEP 6. CALCULATE THE PROCESS LIMITS – NOT ANOTHER TARGET

A useful starting point for individual elapsed times is an **individuals-and-moving-range (XmR) chart**. Its purpose is to distinguish routine variation from signals that something changed in the process.

For a sequence of individual completion times:

- Calculate the mean of all the elapsed times.
- Calculate the absolute difference between each pair of consecutive observations. These are the moving ranges.
- Calculate the average moving range, written **MR-bar**.
- Estimate the three-sigma process limits as **mean + 2.66 × MR-bar** and **mean – 2.66 × MR-bar**. [8]

These limits are derived from the observed behaviour of the process. They are **not acceptable repair times** and they are **not performance targets**.

Repair times are often highly skewed. A negative calculated lower limit is not a physically possible repair time; don't imply otherwise in a chart. For heavily skewed or mixed repair categories, review the suitability and interpretation of the XmR chart with someone competent in process-behaviour analysis. The 1,081-day upper limit quoted earlier comes from the book; we cannot independently recalculate it without the original observations.

STEP 7. LEARN HOW TO READ THE CHART

This is where a measure begins to earn its keep.

If points fluctuate routinely inside the limits: the system may be exhibiting a reasonably consistent pattern. Don't explain every up-and-down movement as good or bad management. Look for recurring causes in the way work is organised.

If a point falls beyond a process limit: investigate what was unusual about that case. It may represent an exceptional interruption, a different kind of work or a genuine process change.

If there is a sustained run or shift: investigate what changed in the system — before attributing the result to effort or celebrating a single good month.

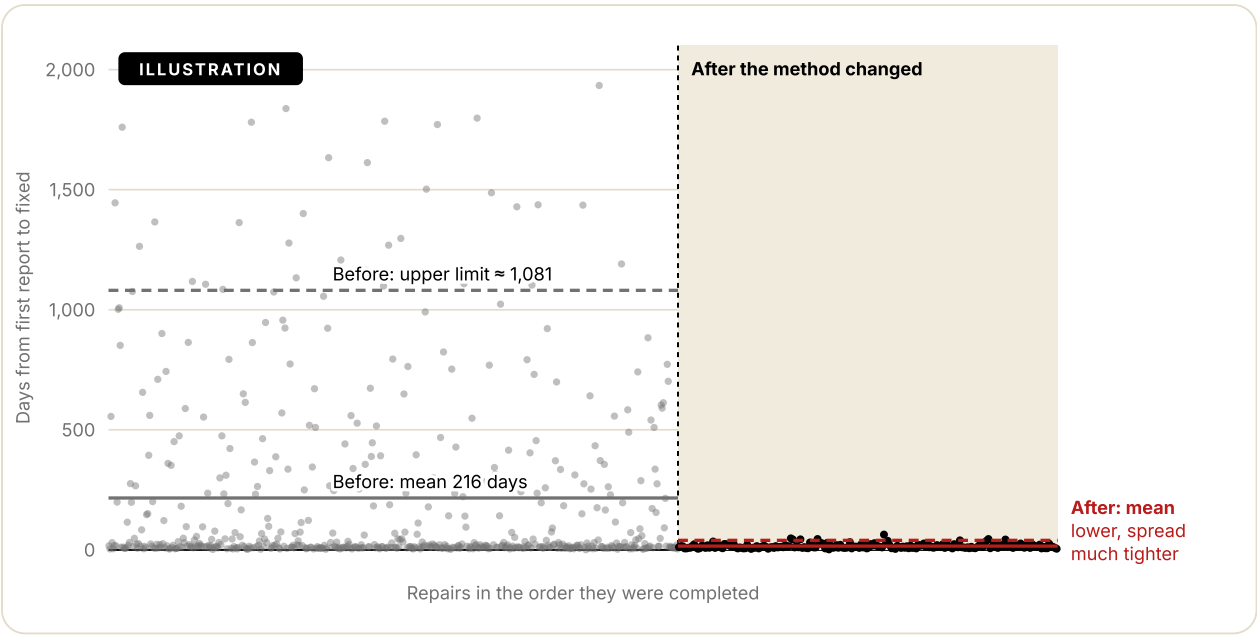
If the entire pattern remains poor: a process can be predictable and still give tenants an unacceptable service. A stable 216-day average is not something to congratulate ourselves about.

Don't confuse a control limit with a promise to the tenant. And don't continually redraw the baseline every time performance moves: retain a credible pre-change baseline so you can test whether an improvement is real.

STEP 8. COMPARE THE SYSTEM AFTER CHANGING THE METHOD

Once you change diagnosis, stock availability, scheduling or authority to complete work, continue measuring **the same customer outcome using the same definition**.

Compare the new time-ordered performance with your baseline. Has the mean shifted? Is variation lower? Are the slowest cases becoming less severe? Has the age of outstanding work improved too? Are tenants expending less effort?



Illustration, not results from the case. It shows the shape to look for after changing the method: the same measure, the old baseline kept for comparison, a lower mean and a much narrower spread.

One good chart can change a management conversation. Two comparable charts can show whether changing the method changed the result.

THE FIRST PRINCIPLE

Establish what your system is capable of achieving before deciding what you want it to achieve.

> Measure an outcome the tenant recognises. Establish its baseline. Study the variation. Investigate the work behind the numbers. Then change the method and measure again.

CHAPTER 05

5 THE EIGHT MEASURES A REPAIRS LEADER NEEDS

End-to-end time is the best place to begin, but it isn't the whole picture. If we were starting with a blank sheet of paper, I'd begin with two questions:

Are we delivering what tenants need?

What in the way we work is helping or preventing us?

Here are eight useful measures — **not eight new targets**.

WHAT IS THE SYSTEM ACHIEVING?

1. End-to-end completion time. From the tenant's first report of a problem to genuine resolution. Report individual days and their behaviour over time, by meaningful repair type.

2. One-stop repair rate. Of the repair problems with an initial operative visit, how many were genuinely resolved on that first visit? Don't count a job coded as complete if the underlying problem persists. Report numerator and denominator explicitly and investigate excluded repair types.

3. Age of outstanding repairs. For every still-unresolved repair, how many days have passed since the original report? Show the count and age distribution by sensible age bands, not merely the total backlog.

4. Failure demand. What share and volume of additional *customer contacts* arose because the organisation didn't do something, or didn't do it right? A tenant chasing an unconfirmed appointment is a candidate. So is a tenant reporting the same drip. A new, unrelated repair is value demand, even if the tenant has called before. Avoid calling internal repeat visits 'failure demand'; those are waste.

5. Tenant Effort Score. Ask after genuine resolution: **How much effort did you personally have to make to get your repair resolved?** Use a consistent five-point scale, from 1 (very little) to 5 (a great deal). Report the distribution, number of responses, collection method and possible non-response bias. Keep the question and scale stable so comparisons over time remain meaningful.

WHY IS THE SYSTEM PERFORMING THAT WAY?

6. Visits per completed repair. Count all operative attendances related to the underlying problem, including revisits. Examine the distribution, not just the mean, and distinguish necessary multi-stage work from avoidable return visits.

7. Time spent doing value work. Observe how much operative time goes into repairing homes rather than unnecessary travel, hunting for parts, waiting, handovers and paperwork. Define what counts as value work; do not simply infer it from electronic job statuses.

8. Total resource effort per completed repair. Look across the whole service, including tenant effort. Contact handling, diagnosis, scheduling, visits, administration, repeat activity and the tenant's own time all matter. Don't count only the cheapest departmental transaction and call that efficiency.

These measures are a starting set, not a compulsory universal scorecard. Keep operational measures close to the people improving the work. Retain mandatory regulatory reporting, but don't confuse it with the information needed to design a better service.

> Eight measures. Not eight targets.

The point is not to build a larger dashboard. The point is to make the reality of the work visible.

CHAPTER 06

6 PRINCIPLE 2 – BY WHAT METHOD?

We've established something rather surprising.

A repairs service reporting excellent performance was taking an average of 216 days to complete routine repairs.

Now we know **what** was happening.

But we still don't know **why**.

Which brings us back to W. Edwards Deming.

Deming had a question for managers who announced ambitious goals. Three words. [5]

By what method?

Our imagined housing association board has decided it wants 95% first-time fix.

Fine.

But what, exactly, is going to change?

To see why this matters, let me introduce you to John Little.

THE £700 TAP

One morning, John accompanied a plumber to a tenant's home.

The job appeared straightforward.

A dripping kitchen tap.

The plumber's handheld device instructed him to replace a washer.

But as they approached the front door, the plumber told John something interesting.

He'd been here before.

Inside, the tap was still dripping.

John asked what needed doing.

The plumber explained that the taps were worn out. Another washer might stop the dripping temporarily, but it wouldn't solve the problem.

Both taps needed replacing.

Did he have replacement taps on his van?

Yes.

So why not fit them?

Because the computer said **replace washer**.

And that was the job he was authorised to do.

John offered to authorise the replacement. The plumber declined.

Even when a senior housing manager agreed to authorise it, the plumber remained reluctant.

Think about this.

Here was a skilled tradesperson who knew what was wrong.

He knew how to fix it.

He had the materials.

He was standing in the tenant's kitchen.

And he was reluctant to do the repair.

Eventually, John persuaded him to replace the taps.

But the really interesting discovery came afterwards.

John examined the property's repair history.

> Ten visits in two years.

At approximately £70 per visit.

£700 SPENT ON A DRIPPING TAP.

The replacement taps cost around **£60**. [7]

And that £700 didn't capture all the inconvenience to the tenant, scheduling, administration and opportunity cost of sending a plumber back again and again.

Would setting that plumber a more demanding productivity target have solved the problem?

Of course not.

He wasn't short of effort. He wasn't short of skill. He wasn't even short of parts.

> He was working in a system that made following an instruction safer than solving the tenant's problem.

And that brings us back to Deming.

The question isn't how to get more effort from the plumber.

> It's what must change in the way the work is organised so the plumber can complete the repair.

For the full investigation — including the van-stock, diagnosis, scheduling and capacity discoveries — see Stuart Corrigan and Neil Alexander, Housing Repairs ER, especially Chapters 1–6. [7]

CHAPTER 07

HOW TO STUDY THE METHOD 7

You can start tomorrow. Not from a board report. Not from a meeting about KPIs.

> From the tenant's first call.

STEP 1. STUDY DEMAND

Sit alongside the people answering repairs calls. Listen. Write down what the tenant actually says.

Not the system category *plumbing enquiry*. Their own words.

> *"My kitchen tap is leaking. Could somebody come and fix it?"*

> *"My tap is still dripping."*

> *"When is the operative coming?"*

> *"I was told to call back and arrange another appointment."*

The first is **value demand**: a legitimate initial request for a service.

The others may be **failure demand**: extra customer contact arising because something wasn't done, wasn't done properly or wasn't explained. Check the circumstances before classifying; repeat contact is not automatically a failure.

Take consecutive calls across different days and times. Start with around 100 if that's manageable. Record the exact request, what caused it and whether it represents value or failure demand. Where you're unsure, follow the repair history rather than guessing.

Calculate **failure-demand calls ÷ all calls studied × 100**; also keep the absolute number and the reason for each failure call. Group the causes: appointment chasing, unresolved repair, missing information, rebooking, broken promises.

An unnecessary return visit is waste. A tenant calling because the repair remains unresolved is failure demand. Keep the distinction clear.

Every avoidable contact consumes capacity inside the organisation and effort from the tenant. **Improving the method can remove the need for some calls altogether.**

STEP 2. PICK ONE REPAIR PROCESS TO INVESTIGATE

Don't attempt to redesign everything. Find a type of value demand that is producing substantial failure demand and customer effort.

A tap repair. A broken window. Loss of electricity. A boiler fault.

For each substantial category, examine initial value-demand requests, related failure-demand contacts and the consequences for tenants. Compare both the volume of avoidable contact and its rate per original request.

Initial value demand	Original requests	Related failure calls	Failure calls per request
Leaking taps	200	150	0.75
Broken windows	100	180	1.80
Boiler problems	150	90	0.60
Electrical faults	80	40	0.50

Example numbers only — not historical findings.

Here, broken windows generate the highest additional calls per original request and the greatest total failure-call volume. That makes them a plausible starting point — but safety, tenant impact and resources consumed must also inform the choice.

Why spend all your time treating a mosquito bite when you have a broken leg?

> Find the largest source of avoidable effort you can safely and meaningfully tackle. Then follow it.

STEP 3. FOLLOW THE WORK – GO AND SEE

You cannot do this from behind your desk. You cannot do it by admiring a beautiful process-mapping booklet.

You have to go and see.

Listen to the first call. Watch the diagnosis. Watch the scheduling decision. Accompany the operative. Follow the case to a properly resolved outcome.

Ask the same four questions at every step:

1. What actually happens here? Not what the written procedure claims should happen.

2. What predictably goes wrong? Wrong diagnosis, unavailable parts, bad handovers, unnecessary permissions, waiting.

3. How often does it happen? Count it. Don't confuse one memorable anecdote with a recurring feature of the system.

4. What is the impact? Choose a relevant consequence: tenant days waiting, avoidable contacts, repeat visits, operative hours, administration or total cost.

What actually happens?	What predictably goes wrong?	How often?	Impact
Operative attends	Part unavailable	8 of 20 observed cases	Further visit; longer wait
Repair diagnosed	Wrong trade sent	5 of 20	Journey and delay
Job closed	Problem not resolved	3 of 20	Tenant calls again

Illustrative field notes. The aim is to replace them with what your people actually observe.

This is the difference between **work as imagined** — the procedure, software workflow and management report — and **work as done**, what happens when actual tenants, tradespeople, properties and constraints meet.

A process map tells you how work is supposed to happen. **The tenant experiences how it actually happens.**

STEP 4. PUT IT ON THE WALL. CHANGE SOMETHING.

Brown paper. Post-it notes. Photographs. Scribbled observations. Who cares?

It will look messy.

> That's because the process is messy.

Don't spend three weeks perfecting the diagram. Stand back and ask:

Where are we creating the most unnecessary effort for the tenant and the operative?

Pick that point. Get frontline people around you: call handlers, schedulers, operatives, supervisors and anyone else who understands the consequences.

Ask what would have to change to prevent the problem.

Then make a small, safe, deliberate change.

Not another target. Not a motivational poster. Something about the work.

Could an operative complete an appropriate repair without waiting for an unnecessary authorisation? Could the diagnosis be improved? Could the required part reach the home before the visit? Could a tenant get a reliable appointment without calling again?

Watch what happens and **measure the same outcomes again**. Keep safety, legal duties and tenant consent intact. Document what you changed and when; don't confuse a short-lived improvement with a proven new capability.

A NOTE ABOUT CONTROLLED EXPERIMENTS

In our consulting work, we go further. We use what we call a **clean room**: a deliberately controlled setting for testing a redesigned way of working, examining demand, capacity, roles, rules, materials and decisions together.

> It is a more rigorous approach than making one change on the spot. But you don't need that level of expertise to begin learning.

> Find something that predictably goes wrong. Change the conditions that cause it. Measure what happens.

CHAPTER 08

A WORKED EXAMPLE – WHAT IS IN THE VAN?

8

Let's return to the plumber's van.

One of the biggest obstacles to completing a repair first time is surprisingly simple: **the operative doesn't have the part required**.

Off they go to the merchant. Sometimes they return later. Sometimes a second appointment is needed. Meanwhile the tenant waits, the van travels and the organisation consumes capacity.

You might assume the answer is to load more stock. But that isn't necessarily true.

In *Housing Repairs ER*, John Little and the team opened a van and found materials accumulated over time — including **seventeen rolls of PTFE tape**. What the operative thought he used was not always what the actual records showed. The team examined consumption, variation and the needs of different operatives rather than imposing a single standard loadout. [7]

The problem was not simply how much stock was in the van.

> It was whether the stock matched the work.

Proper van-stock design is more complicated than a quick inspection. It needs consumption histories, replenishment times, variation, repair mix and an appropriate method of restocking. We have developed a system for addressing this through **Stock Right Now**.

But here's an easy place to start.

THE 40-PART VAN-STOCK TEST

Ask an operative if you can look through their van together. Pick a part.

"How often do you use this?"

Record the answer. Then check what the consumption records suggest.

Repeat across **30–40 different parts**. Include both commonly used items and things that appear to have been sitting unused for some time.

Part	Operative's estimated usage	Recorded consumption	Available when needed?
[Part]	[Estimate]	[Actual history]	[Yes / No / Unknown]
[Part]	[Estimate]	[Actual history]	[Yes / No / Unknown]
[Part]	[Estimate]	[Actual history]	[Yes / No / Unknown]

Use this as a conversation starter, not as a validated inventory design.

Don't use the findings to embarrass the operative. Human memory is a poor inventory system. People remember the occasion they ran out and forget the many ordinary days they didn't.

Ask what was missing when a repair could not be completed, how frequently it happened and what extra visits or travel followed.

> You can give an operative a tougher first-time fix target. Or you can change the method by which the right part becomes available.

One puts pressure on the person. The other improves the system.

For the fuller van-stock method and the original case narrative, see *Housing Repairs ER* and Stock Right Now. [7]

Stock Right Now: stockrightnow.com

CHAPTER 09

9 THE FINAL PIECE – SYSTEM CONDITIONS

There's a danger with everything we've just discussed.

You could follow the work, identify the delays, draw a process map and announce a business process reengineering programme.

Please don't mistake that for what we're suggesting.

Because the problem may not be the visible sequence of activities.

It may be the **rules, measures, contracts, permissions and managerial assumptions** that determine how those activities happen.

John Seddon's work calls attention to these **system conditions**: the circumstances created by the way organisations are designed and managed. His central connection is **Thinking → System → Performance**. [9]

Think of our £700 tap. The plumber had the skills and the parts. Improving the route to the property would not have fixed the permission rule that made changing the taps risky.

Or the van stock. Asking the operative to be more productive did not make missing materials appear.

Or the tenant who has to call repeatedly for information. Faster call handling does not remove the reason for calling.

These conditions didn't appear by magic.

They were created by management decisions: how we specify jobs, how we purchase work, what we count, who is authorised to decide, how work is divided and what managers believe control looks like.

> Management thinking shapes the system. The system shapes performance.

THE THINKING THAT REFUSES TO DIE

Remember our imaginary boardroom?

NEXT YEAR: 95% FIRST-TIME FIX.

It is worth asking where that confidence in target-setting came from.

Performance targets didn't begin with Tony Blair. They were part of British public-sector management long before him. But from 1998, New Labour's Public Service Agreements greatly expanded the role of measurable commitments across government. They were intended to improve delivery and accountability for public money. [10]

That intention was understandable. We should also be fair: targets can sometimes direct attention to neglected problems. The mistake is to assume that specifying a desired number automatically tells you how to change the work.

In housing, an operative can be working hard, a supervisor can be scrutinising every completion, and the board can be demanding 95% first-time fix — while the conditions responsible for repeat visits remain untouched.

Deming's question returns.

By what method?

Three words. And they send us back to the beginning.

The tenant wants their home fixed. The operative wants the means to fix it. The leader's job is not merely to demand a different number, but to understand why the current system produces the result it does — and change the conditions that shape that result.

> Performance is a feature of the system, not simply a measure of how hard individuals work.

CHAPTER 10

YOUR MONDAY-MORNING STARTING POINT

10

This guide is a beginning, not a qualification in statistical process control or systems redesign.

But you can begin now.

CHOOSE	one customer outcome that matters.
DEFINE	when the customer's problem starts and when it is genuinely resolved.
ESTABLISH	current capability, over time, from actual cases.
LISTEN	to the customer's demand, in their own words.
FOLLOW	the work, from first contact to finished repair.
ASK FOUR QUESTIONS	What happens? What predictably goes wrong? How often? What's the impact?
CHANGE	one meaningful system condition with the people doing the work.
MEASURE AGAIN	did tenants get a better outcome with less total resource effort?

That's enough to make a start.

And the next time someone walks into your boardroom with an ambitious new target, ask them a question.

It's only three words.

BY WHAT METHOD?

GO FURTHER

For the full investigations behind the £700 tap, van-stock and repairs-capacity examples, read **Housing Repairs ER**, by Stuart Corrigan and Neil Alexander.

Explore Housing Repairs ER: loweffort.com/housing

For help measuring capability, studying demand and testing a redesigned method of housing repairs, contact **Descartes Consulting**.

Start with the free presentation: loweffort.com/housing/presentation

A ONE-PAGE FIELD WORKSHEET

Repair type / customer problem: _____

What does the tenant need, in their own words? _____

How do we know the problem is genuinely finished? _____

Current capability: sample _____ period _____ mean _____ process limits _____

Age of outstanding repairs checked? Yes / No

What happens?	What predictably goes wrong?	How often?	Effect on tenant / total resource effort

The largest avoidable burden we found: _____

The system condition that appears to cause it: _____

The people who need to help us change it: _____

The safe change we will test: _____

Date of change: _____

How we'll know whether it worked (same outcome, same definition): _____

What changed in the data? _____

What did tenants and operatives tell us? _____

SOURCES

The housing repairs case: the historical figures and the £700 tap and van-stock stories come from Housing Repairs ER, by Stuart Corrigan and Neil Alexander, especially Chapters 1, 2 and 6. The 97% is that service's routine repairs; the 79% is the sector-wide figure from the regulator.

[1] Regulator of Social Housing, Tenant Satisfaction Measures 2024/25: Headline Report (4 November 2025): View source

[2] Regulator of Social Housing, Consultation on the introduction of tenant satisfaction measures, sections 8.13–8.15: View source

[3] House of Commons, Public Administration Select Committee, ambulance response-time evidence (2002–03): View source 1 ; View source 2

[4] C. A. Mace, Incentives: Some Experimental Studies (1935), Wellcome catalogue: View source 1 ; Paula Lynne Phillips, 1991 Academy of Management study of Mace: View source 2 .

[5] W. Edwards Deming Institute, The Futility of a Numerical Goal (including “By what method?”): View source 1 ; Deming’s 14 Points: View source 2

[6] Toyota Motor Corporation, 75 Years of Toyota — Establishment of quality control system: View source

[7] Corrigan & Alexander, Housing Repairs ER, Chapters 1, 2 and 6.

[8] JMP, Individual and Moving Range Chart (three-sigma moving range method): View source

[9] John Seddon, Vanguard Consulting, Thinking / System / Performance: View source

[10] Institute for Government, Public Service Agreements and the Prime Minister’s Delivery Unit: View source