NEWELLOI

OWNER'S GUIDE / 03

Your Business Doesn't Need Another Dashboard

How to get the right decisions in front of you at the right time, and keep everything else off your desk.

THE SHORT VERSION

Your business produces far more information than you can read. More reports, alerts and charts don't fix that. They just move the pile. What helps is a dashboard built as a decision interface, not another stack of reports.

This guide describes a better approach: routine work you've approved gets done and checked, and only the decisions that need your judgment reach you, with the facts attached. It also shows how to tell whether a system like this is actually working, and why sometimes the most valuable thing it can do is bring a person in.

For owners of established, staffed independent businesses

By Newell Starks, Founder, NewellOI | November 2026

Examples are illustrative unless identified as documented cases. Nothing here is legal advice.

1. You can't read everything, and you shouldn't have to

Every day your business throws off more signals than any one person can follow: calls, estimates, cancellations, service notes, reviews, web inquiries, inventory changes, customer anniversaries and whatever the competition is doing. You already have the judgment to act on almost all of it. What you don't have is the time to see it all before the moment passes.

The costly problems usually start small:

- An unanswered complaint turns into a lost renewal.
- A good customer quietly goes longer and longer between visits.
- A recurring dispatch problem eats into profit per job.
- A supplier change causes rework nobody connects to it.
- A competitor makes booking easier, and calls slowly drift their way.

Each one looks minor on its own. Its importance only shows up in context and timing. The answer isn't to send every one of them to you. A good system works like a disciplined filter: it uses the rules you've set, the evidence in your records and what your people know to decide what needs action, who is allowed to take it, and when to stay quiet.

2. From noise to a decision

A system you can trust handles each event in the same six steps, and keeps a record of each one so that you, or anyone you appoint, can check what happened, find mistakes and correct them. This follows the same principles as the federal NIST AI Risk Management Framework: clear human oversight, defined roles and ongoing checks. [1]

Step	The question it answers	What it keeps on record
Notice	What happened, and where did we hear it?	Time, source, which customer or job
Sort	Is this routine, unusual or serious?	The rule or threshold used, and anything uncertain
Prioritize	What decision is needed, by whom, and by when?	Likely impact, urgency, what else is competing
Act or bring someone in	What is it allowed to do on its own?	The authority it acted under, any approval, what it did
Check	Did the intended result actually happen?	Independent confirmation and error checks
Learn	Should a rule or step change?	The outcome, cost, feedback and each rule change

Different situations need different handling. A customer asking your hours can get an approved answer automatically, checked for completion. A request for an unusual discount goes to a manager. A safety complaint, a major contract or a privacy concern comes to a person right away. A system that treats these the same isn't saving you time; it's creating risk.

> **A DRAFT IS NOT A DECISION**
> A recommendation is not an approval. A prepared email is not a sent email. A drafted work order is not permission to commit labor or money. A good system keeps every one of those lines clear, and it never crosses one without the authority you've given it.

3. Judge the work, not the automation

"AI handled 92% of inquiries" sounds impressive and means almost nothing. A system can "handle" every inquiry while making promises you can't keep, missing problems that needed a person, and leaving your staff to clean up afterwards. What matters is how the work turned out. Judge every case on two separate questions:

Question	Possible answers
Was the work done right? (execution)	Completed correctly • Completed incorrectly • Left incomplete • Failed
Was the right person involved? (intervention)	Rightly handled alone • Rightly brought to a person • Brought to a person unnecessarily • Should have come to a person but didn't

Keeping the two questions apart matters. A task the system was allowed to handle on its own, and did handle on its own, can still be done wrong. That's an execution error, not a missed escalation, and a single list of outcomes lets it slip between the categories. The best result is work completed correctly *and* rightly handled alone. An unnecessary hand-off is something to fix, so count the staff time it took. A missed escalation or an incorrect result is a failure: find out how serious it was, and prevent it.

Three habits keep this scorecard honest:

- **Check the work independently, mostly automatically.** A system that grades itself can look efficient simply by never asking for help. Use automated checks that don't rely on the system doing the work wherever they're reliable, such as confirming an appointment really appears on the schedule or an invoice matches the work order. Save human review for serious exceptions and for cases where an automated check fails or can't decide. The goal isn't to build a second staff just to watch the first.

- **Review a risk-weighted sample of the work it handled alone**, with more attention on work where mistakes cost more, looking for errors and for cases that should have come to a person. Those misses don't show up on their own.
- **Decide in advance which work counts as eligible**, and log any change to that list. Otherwise the success rate can rise just because the hard cases were quietly dropped. Work that always needs your approval is reported separately, not counted as a success or a failure.

Pair those numbers with the staff minutes each completed job still takes, the cost of fixing mistakes, and what happened to the customer. More automation that weakens customer relationships isn't progress.

4. Customer loyalty is a judgment call

Customers don't all want the same relationship with you, and what they want changes over time. How often they buy, what each job earns you, whether they want a personal touch, and how tempted they are by a competitor's offer all vary by trade, by area and by person. [2] An automatic "anniversary email" doesn't understand any of that. The point is to know when a thank-you, or an effort to fix something, will actually make a difference.

Milestones depend on how your customers buy. A 50th visit to a neighborhood salon deserves real thanks. In pest control, the meaningful milestone may be five years of uninterrupted renewals. For a roofer, a household that buys ten times is so rare that a milestone rule would make no sense.

Signal	Possible response	Why a person might matter
Tenth completed visit	Personal thanks, or a quick check on preferences	The relationship is moving past trying you out
Fiftieth repeat purchase	A note or call from a manager or the owner	A loyal customer may value recognition more than a discount
Longer gap than usual between visits	Ask whether something went wrong or needs to change	The reason is unclear; don't assume
Serious repeat complaint	Direct, personal effort to make it right	Empathy and accountability often matter more than speed

These are possible triggers, not proven ones. A milestone message can feel awkward or manipulative if it misreads the relationship. One customer appreciates a note; another would rather not hear from you. Personal touches should follow what customers have actually told you they want, not just what the software can send. And when customers are anxious or upset, research suggests that being able to reach a real person makes a measurable difference. [3]

5. Know how strong your evidence is

A good system doesn't recommend something just because it worked somewhere else. It weighs four levels of evidence, and they aren't equally reliable:

- **Industry**: published research and what's typical for your trade. Useful for knowing where to look.
- **Regional**: seasons, weather, labor costs and what customers in your area expect.
- **Local**: what nearby competitors are visibly doing, and your own results in your market. You can see a competitor's discount; you usually can't see whether it's profitable.
- **Individual**: a customer's history with you, their stated preferences and their service experiences. This is the strongest evidence of how someone behaves, but on its own it doesn't prove what *caused* that behavior.

Here's how that plays out. A customer has had 49 scheduled appointments over four years, is clearly satisfied, and is due for the 50th next week. In the past you've called important customers yourself. The system should pull together that history and suggest a short thank-you, and you decide whether to call. It shouldn't send a gift or promise an exception on its own unless you've given it that authority.

> **DID THE THANK-YOU CAUSE THE LOYALTY?**
> Your most loyal customers will naturally be the ones who receive milestone thanks. That doesn't prove the thanks made them loyal. To find out whether recognition works, compare similar customers who did and didn't receive it, over enough time, and where that's lawful and practical. At minimum, track repeat visits, profit per customer and complaints against a sensible baseline, and remember that other things may have changed too.
>
> As in Guides 01 and 02, keep three kinds of numbers apart: **observed** (counted in your records), **estimated** (based on stated assumptions) and **verified** (measured against a comparison that tests and reduces plausible alternative explanations).

6. A dashboard for decisions, not a wall of numbers

Most dashboards are collections of reports: charts and figures that still leave you to work out what matters. A decision interface works the other way around. It starts with what needs your attention, why, by when, and what your options are. The NewellOI Owner's Intelligence Dashboard is designed as a decision interface first. Its front page isn't a chart; it's a short list of decisions waiting for you. Each item should say:

- **what the issue is** and **what decision you're being asked to make**;
- **how urgent it is**, and **what's likely to happen** either way;
- **how confident the system is**, and **the evidence** behind it;

- **your options**, ready to act on.

The detailed operating measurements are still there, one step behind that list, whenever you want them. A handful of them tell you whether the business, and the system, are doing their jobs:

Area	What to measure	The question it answers
Customer relationships	Time between visits, retention, profit per customer, referrals	Are my best relationships getting stronger or weaker?
Operations	Jobs completed, problems fixed, rework, profit per job	Is the work producing better results?
Work done without you	Eligible work completed correctly and checked, with no unplanned help, and errors in work done alone	What is truly running on its own, and how well?
When people step in	Right escalations, missed escalations, staff minutes	Is judgment being applied at the right moments?
Speed of decisions	Time from an event to the right action	Are important decisions reaching me too late?
Net result	Added operating earnings minus the full cost of running the system	Is this actually improving the business?

Every number needs a definition. "92% resolved automatically" means nothing until you know which work was eligible, how results were checked, and whether anything that needed a person slipped through. "Retention improved" means little without knowing which customers, over what period, and whether the mix of work changed.

7. A worked example: the tenth visit

Picture a locally owned auto-service shop. A customer has made nine visits over six years and is coming up on a tenth. The shop has noted that she wants a clear explanation before any major repair. On her last visit, she complained that nobody could tell her when her car would be ready.

A basic loyalty bot sends her a "10th visit discount." A better system sees that the unresolved complaint matters more than the visit count. It checks whether the shop already fixed the problem, pulls together the timeline of her last visit, and suggests that the service manager call her. The manager acknowledges the problem and explains the shop's new way of updating customers on pickup times. Once her next visit goes smoothly, that's recorded. Only then does a thoughtful thank-you make sense.

This is an illustration, not a NewellOI customer case. It shows the difference between automating transactions and supporting an owner's judgment. The right move depends on the order of events, the relationship, what the customer has agreed to, what's at stake, and real human judgment.

Afterwards, the system records whether the complaint was resolved, whether she came back, whether she said anything about her experience, and how much staff time it took. None of those on its own proves the call caused her to return. Together, they help make the next decision a better one.

8. Protecting trust

The information behind good decisions is often sensitive: prices, customer messages, purchase history, employee performance. Access should be limited to what's needed, for a stated purpose, with clear rules on how long information is kept. One client's confidential information should never reach another, even when both use the same industry approach. The FTC recommends keeping only the sensitive information you need and limiting who can reach it. [4]

Marketing rules still apply. Under the FTC's CAN-SPAM guidance, routine transactional or relationship messages are treated differently from commercial ones, and having done business with someone before doesn't make every promotional email exempt. [5] A milestone "thank you" that includes an offer may count as marketing. Check with your attorney if you're unsure.

Every action the system takes should be visible and open to review. But not every action can be undone. A sent email can't be reliably recalled, a promise to a customer may create an obligation, and confidential information shared outside the business can't be taken back.

So actions fall into two groups. **Reversible actions** can be taken within the authority you've given and corrected if something goes wrong. **Actions that can't be reversed** need controls before they happen: your approval, a second check, or a rule that the system never takes them on its own. Where you draw that line is one of the most important parts of the authority you give the system.

For every action, keep the information it used, which version of the rule applied, the approval it acted under, what it did, and whether the result was checked. A system that acts on its own without that record isn't one you can depend on.

9. Start with one decision

Your first step doesn't need to predict every customer who might leave or recommend every price change. Pick one limited process with reliable data, a clear goal and a defined point where a person steps in. Answering common questions with an easy hand-off to staff is a natural start. So are appointment confirmations, alerts on unresolved service problems, and thanks for selected repeat customers.

Then follow the same Define, Deploy and Operate steps from Guide 01. Record how things work today and what staff currently do. Decide which work the system may handle on its own. Run the new process with visible checks and a reliable fallback. Measure customer results and the net financial effect. Expand only when the evidence shows the added capability is worth it and can be supported. The test isn't how clever the technology is. It's what your business can depend on.

Where this fits

Guide 01 is about finding and fixing the money you're losing. Guide 02 is about holding your ground as larger competitors arrive. This guide describes what connects them: an owner-controlled capability that keeps watching, handles the routine, and brings you the decisions that need you, with the evidence attached. The Owner's Intelligence Dashboard is where you see that each day, and each improvement adds to it.

You belong at the center of your business's judgment, not at the center of its paperwork. A good system makes that difference real every day: fewer distractions, better-informed choices, timely personal attention where it counts, measurable results, and a clear record of why the business acted.

> **WHAT REACHES YOUR DESK?**
>
> Tell us about the decisions that land on you too often, or too late. We'll help you sort what could be handled without you, what should always come to you, and how we'd measure whether the change is working.
>
> **newelloi.com**

Sources

Sources accessed November 2026. Figures and findings keep the scope and limits set by their publishers. **About these examples:** the businesses and customers described are hypothetical. This guide describes how NewellOI designs systems; it is not a claim that every capability described is in use today. Third-party research is not a NewellOI customer result, and nothing here is a forecast or guarantee of performance.

[1] National Institute of Standards and Technology (2023). AI Risk Management Framework 1.0: Core. Voluntary framework covering human oversight, defined roles and ongoing evaluation of AI systems. NIST AI RMF Core (nist.gov)

[2] Avery, J., S. Fournier & J. Wittenbraker (July 2014). Unlock the Mysteries of Your Customer Relationships. Harvard Business Review. HBR article (hbr.org)

[3] Shell, M. A. & R. W. Buell (April 2019). Why Anxious Customers Prefer Human Customer Service. Harvard Business Review. HBR article (hbr.org)

[4] U.S. Federal Trade Commission. Protecting Personal Information: A Guide for Business. FTC business guide (ftc.gov)

[5] U.S. Federal Trade Commission. CAN-SPAM Act: A Compliance Guide for Business. FTC compliance guide (ftc.gov)