



When the Paperwork Steals the Care

Paperwork becomes a care issue when it displaces the conversation, observation and follow-up that people came to the service for. A careful redesign can remove some repeated work, but it must also stop the organisation from filling every recovered minute with another form.

David Saliba . 2026-07-25 . aimonger.com/whitepapers/ngo-admin-burden-care-time/

The operational problem

If the form takes longer than the session, the system is winning and the child is losing.

This brief applies the leadership approach in [A Practical AI Roadmap for NGO Leadership Teams](#) to one workflow. This is an operational technology brief, not clinical or safeguarding guidance. Local professional standards and escalation routes remain in force.

Nobody needs to accept a headline figure to see the shape of the problem locally, but a headline figure is a useful place to start arguing. The Federation of European Social Employers put the Dutch care-worker administrative share at roughly **36%** of working time in its 2023 research, and worker representatives in Austria and Belgium separately described paperwork getting heavier while the software meant to lighten it did not deliver. None of that transfers automatically to a specific team; a service running lean rotas and few funders may sit well under that share, one juggling several statutory returns may sit above it. The only way to know is to time the actual week rather than accept either the vendor's promise or the sector average.

That measurement matters more once staffing is already stretched. Eurofound's 2023 social-services report ties workforce shortages to both harder working conditions and patchy digitalisation across the EU care sector, and OECD's Beyond Applause? makes the related point that time pressure in long-term care is a working-conditions problem, not just an efficiency one. WHO Europe's MeND survey adds that workload of this kind carries a wellbeing cost, not only a productivity one. One consequence follows directly for a paperwork pilot: about a fifth of social-care workers across the EU never use a digital device at work, so a redesign built around the fastest typist on the team will make the rota more unequal rather than lighter.

What a real week looks like

A key worker finishes a productive session, writes a note while the details are fresh, then opens three systems asking for overlapping fields. A parent needs an update, a school needs a form, and a funder deadline has become urgent. None of those tasks is unreasonable alone. Together they turn the evening into the place where the organisation stores unfinished work.

The practical unit of improvement is not "admin" in the abstract. It is a specific repeated journey: session to signed note, referral to completed form, or meeting to assigned actions. Once that journey is visible, leaders can remove a handoff or field rather than asking staff to become faster at carrying the same burden.

Building the map is a job for the key worker who fills the form each week, not a steering group working from a flowchart. Only the person doing the task knows which field genuinely changes what a colleague decides next week and which field is copied forward purely because deleting it once caused an argument with a funder. A short-staffed rota cannot afford to remove the wrong one: cut a needless handoff and the week gets lighter, cut a field a reviewer actually reads and the team has just created a new risk.

What the evidence says in 2026

Paperwork is, on this evidence, the workflow where the sector has already placed most of its AI bets. The **Charity Digital Skills Report 2026** (launched 9 July 2026, n=807) found **63%** of charities already using AI for administration and project management, ahead of every other named use case bar drafting itself. That is the single most relevant number in the whole survey for this brief: administration is not a hypothetical future use of AI in the sector, it is the current one.

The same report puts that adoption in a governance context managers should not skip past. Overall AI use reached **79%** of charities (up from 76% in 2025 and 61% in 2024, and 92% among large charities), yet only **28%** had a documented digital strategy, itself down from 44% in 2025. A sector can be using a tool on 79% of its desks and still not have decided, in writing, who is accountable when it produces something wrong. For a paperwork pilot specifically, that gap is the difference between “staff are drafting with an AI tool” and “the organisation knows what that tool is allowed to see.”

The barriers point towards a specific pilot size. As of 2026-07, **63%** named squeezed finances as their biggest digital barrier, **56%** named lack of skills or technical expertise as the biggest AI barrier, and **35%** said they did not trust AI tools outright; staff training was the top funding need for **44%**, while only **17%** had received dedicated digital funding. Read against the administration figure above, the sensible reading is not “roll out AI drafting everywhere it might help” but “start with the one paperwork task the team already half-trusts a tool to touch, fund the training properly, and prove the case before asking for more budget.”

Where technology helped - and where it stalled

Shared templates, a reliable action register and approved drafting support can reduce the blank-page problem. They do not repair duplicate data entry by themselves. The technology stalls when the old form remains mandatory, a second platform becomes the unofficial record, or the minutes saved at one stage become extra inbox work at the next.

Put plainly, the honest limit on what drafting assistance can do for paperwork is this: it can locate a fact that already exists, arrange it under the right heading, propose routine wording and flag a missing required field. It cannot decide that a record is accurate, that a plan is the right one, or that a disclosure is proportionate. Those remain judgement calls, and a tool that quietly makes them anyway is not saving time, it is moving risk downstream to whoever signs the record without re-checking it. The trap specific to paperwork is subtler than a bad answer: a faster draft that gets filed straight into a second system nobody has decommissioned yet, so the organisation now maintains two records instead of one.

That trap is more likely where governance is thin. The 2026 Charity Digital Skills Report rated **33%** of charity boards poor on AI skills, which in practice means a third of boards cannot say what their own drafting tool is allowed to see. Without a named data boundary and a named reviewer, staff facing a Friday-afternoon deadline will reach for whatever finishes the form fastest, sanctioned or not. The fix is not a policy document; it is making the approved route genuinely quicker than pasting the case file into a personal account.

Not every field is missing an AI system

Not every step in this paperwork map needs an AI system sitting inside it. The week-mapping exercise earlier in this brief already found the honest shape of the problem: a field a key worker keeps filling

because deleting it once caused an argument with a funder is a missing ownership decision, not a drafting gap. The same goes for a countersignature queue with no agreed turnaround time, or a review date nobody has diarised. A pilot that reaches for drafting assistance before anyone removes a needless field or names an owner will make the wrong step faster, not lighter.

The pattern worth automating sits one level above any single form. Across a caseload and a funding cycle, the same fields drift, the same countersignature queue backs up on the same days each month, and the same funder deadline produces the same evening scramble every quarter. Connected agents reading that recurring pattern across many weeks, rather than one note at a time, can tell the named owner a return is trending late three days before the deadline instead of waiting to be asked. Gartner's June 2025 analysis of agentic AI project cancellations calls the alternative "agent washing": a vendor relabels an existing chatbot as an autonomous agent and sells it as the fix for a backlog that a missing ownership decision would have solved for free. SAS's July 2026 research frames the safer version as keeping automation in service of a person's judgement rather than replacing it. AIMonger's position follows that logic: an agentic layer over the routes and templates that already work here, so each pilot adds to what works instead of resetting the team to a generic process.

A safer AI-assisted path

1. **Spend five ordinary days mapping one role's non-direct work.** This step has no AI role: only the key worker filling the form each week knows which field a funder actually reads and which is copied forward out of habit. Record the task, system, waiting time and every repeated fact, then ask the people doing the work whether the map feels true. That honest baseline is what later drafting and pattern-watching can be measured against.
2. **Choose one recurring document with an owner and a stable quality standard.** Simplifying the template is human work with no AI role: delete needless fields and name an owner before any tool touches the document. Once that structure is settled, approved drafting may arrange known facts under those headings, but a named reviewer must still check every external output against the source record.
3. **Use approved systems for any sensitive content and retain a named human reviewer for external outputs.** Drafting support may arrange known facts under headings, but it must not invent observations or decide what should be shared. The reviewer should be able to see the source material and correct the draft.
4. **Run the path for four weeks and decide from the results whether it should continue.** If the watching layer is flagging the same duplicate fields, countersignature backlog or evening scramble each quarter while notes look tidier, a named owner should decide whether to remove a step, retire a duplicate system or halt the pilot. Measure whether evening completion, duplicate entry and direct-support time actually changed, not whether another drafting tool was added.

Where humans must intervene

Managers choose the workflow and protect the recovered time. Practitioners remain responsible for factual accuracy, safeguarding leads retain escalation decisions, and a named information-governance owner sets the data boundary. No output to a family, school, funder or partner should leave automatically.

NIST's Generative AI Profile gives a workable checklist for a paperwork pilot specifically: know what the drafting tool cannot verify, keep a record of where each fact in a draft actually came from, and treat a completed form as a consequential output that needs a human sign-off, not a routine one. That discipline does not lapse because a funder's deadline is close; GDPR and, where it applies, the EU AI Act still govern what may be typed into the tool in the first place.

Risks and failure conditions

The most likely failure here is not a dramatic error but a quiet volume increase: a drafting tool makes each form faster to produce, so the organisation asks for more forms rather than fewer, and the paperwork share of the week does not fall. A second, related failure is a two-speed team, where one confident colleague adopts the tool and effectively re-keys everyone else's notes because nobody redesigned the underlying process. Both point to the same root cause: a sponsor for the tool, but no named owner for the outcome once the pilot goes live.

A further risk is mistaking motion for progress. The 2026 report found **81%** of charities reporting some digital progress in the last year, a figure broad enough to include almost any change at all. It says nothing about whether a specific key worker got an evening back, whether the notes that came out were actually better, or whether anyone checked. A paperwork pilot only counts as progress if it can point to fewer minutes on the chosen form and no drop in what the record captures.

Measures that matter

Track median time to close the chosen document, same-day completion with a quality sample, and the number of fields re-entered from an existing record. Pair those with protected direct-support hours and evening work, otherwise a local time saving can simply move the pressure elsewhere.

Put these numbers in front of the people actually filling the form at the same point every week, and ask them where the time went rather than telling them. A key worker who says the form still takes as long, just with better wording, has just told the manager the redesign targeted the wrong step. Treat that as the finding it is before deciding whether the pilot deserves a wider rollout.

Decision questions

- 1. Which repeated document is costing the most skilled time each week?** Name it, name who currently owns fixing it, and state what happens to that person's week if the answer turns out to be wrong.
- 2. What will actually be removed, not merely drafted faster?** If the answer is only "the blank page", the pilot has not identified a real saving yet.
- 3. Who can approve the output and halt the pilot if quality falls?** That person needs the authority to stop the rollout, not just the job title to sign it off.
- 4. What data is prohibited from unapproved tools?** Write the boundary down before the first draft is typed, not after an incident forces the question.
- 5. Where will recovered minutes go: direct support, supervision, or another demand?** If nobody can answer this in advance, the organisation has already decided the minutes will be absorbed by whatever is loudest that week.

FAQ

Does this mean fewer records? Not necessarily. It means records should be purposeful, entered once where possible, and owned by someone who can remove stale requirements.

What can start without a platform budget? A five-day week map, a cleaner approved template and a clear rule about source records can expose waste before software is purchased.

References

1. David Saliba, A Practical AI Roadmap for NGO Leadership Teams, AIMonger (2026). <https://aimonger.com/whitepapers/ngo-practical-ai-roadmap-leadership/>
2. Federation of European Social Employers / research partners, Evolving jobs, skills and training needs in the social services sector and the role of social partners in managing changes (2023). https://www.socialemployers.eu/wp-content/uploads/2023/09/report_Evolving_jobs_skills_and_training_needs_in_the_social_services_sector.pdf
3. Eurofound, Social services in Europe: Adapting to a new reality (2023). <https://www.eurofound.europa.eu/en/publications/all/social-services-europe-adapting-new-reality>
4. Charity Digital Skills Report 2026 (Zoe Amar Digital / Nissa Ramsay; launched 9 July 2026, n=807). <https://charitydigitalskills.co.uk/report/>
5. NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, AI 600-1 (2024). <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.600-1.pdf>
6. Regulation (EU) 2016/679 (GDPR). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
7. Regulation (EU) 2024/1689 (EU AI Act). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
8. OECD, Beyond Applause? Improving Working Conditions in Long-Term Care (2023). <https://doi.org/10.1787/27d33ab3-en>
9. WHO Regional Office for Europe, Mental Health of Nurses and Doctors survey in the European Union, Iceland and Norway (MeND, 2025). <https://www.who.int/europe/publications/i/item/WHO-EURO-2025-12709-52483-81031>
10. Gartner, “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled By End of 2027” (25 June 2025). <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
11. SAS, “For Agentic AI ROI, Invest in Human Judgment” (22 July 2026). <https://www.prnewswire.com/news-releases/sas-for-agentic-ai-roi-invest-in-human-judgment-302831330.html>

Published by AIMonger . <https://aimonger.com> . Malta . Europe