



Behaviour Support Plans Without the Copy-Paste Marathon

A Behaviour Support Plan is useful only when people can see the person, the evidence and the agreed response in it. Faster document production is valuable only if it protects that work rather than turning a past plan into authoritative-looking boilerplate.

David Saliba . 2026-07-25 . aimonger.com/whitepapers/ngo-behaviour-plans-without-copy-paste/

The operational problem

A living plan helps a child. A copy-pasted plan helps nobody and risks everyone.

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 educational advice. Apply Positive Behaviour Support practice, clinical governance, safeguarding policy and professional standards.

Behaviour Support Plans sit inside a wider documentation load that is worth naming before singling this one document out. The Federation of European Social Employers put the Dutch care-worker administrative share at roughly 36% of working time in 2023, and worker representatives in Austria and Belgium reported the same pattern the sector keeps describing: paperwork getting heavier, and the software bought to lighten it not delivering. A plan is one document among many competing for that time, which is exactly why it is vulnerable to shortcuts when the week is already full.

The shortcut risk is sharper here than in most paperwork, because Eurofound's 2023 social-services report links the same staff shortages and uneven digitalisation to a workforce that is already stretched thin, and a specialist under that pressure is precisely the person most tempted to start from an old plan rather than a blank page. It is also worth remembering that roughly a fifth of social-care workers across the EU never use a digital device at work at all, so any redesign of how a plan gets written has to work for the specialist without a laptop habit, not just the confident early adopter.

What a real week looks like

After a difficult incident, a specialist has staff observations, family insight, a school email, an older assessment and a meeting note with no clear action owner. The tempting route is to copy last year's plan and change a few nouns. That produces a complete-looking document before it produces a shared understanding of what happened or what should change.

A plan becomes unsafe when its authority is mistaken for its evidence. NICE NG11 expects a written plan to rest on a shared understanding of function and to include proactive, preventive and reactive strategies, risk management and continuing review. Those requirements ask for current professional work and participation, not generic language assembled from an archive. NICE NG93 and BILD's Positive Behaviour Support resources reinforce the need for service design and practice that keep the plan connected to lived support.

A specialist, not a steering group, is the only person who can map how a plan actually gets written, because only they know which sections come from a fresh observation and which get quietly carried forward from the last review. That distinction decides whether the map is honest. A steering group reconstructing the process from a template will not spot the copy-paste habit; the specialist doing the work will, if asked directly and given permission to name it without being blamed for it.

What the evidence says in 2026

The relevant caution from the **Charity Digital Skills Report 2026** (launched 9 July 2026, n=807) is not that AI is spreading through the sector, which it clearly is at **79%** of charities (up from 76% in 2025 and 61% in 2024, 92% among large charities), but that governance has not kept pace: only **28%** had a documented digital strategy, down from 44% the year before, and **33%** of boards were rated poor on AI skills. For a document that is meant to carry clinical and safeguarding authority, that gap between adoption and governance is the whole risk in one statistic. A tool being widely used tells a manager nothing about whether it is being used safely on a document this sensitive.

The barriers reported alongside that adoption explain why a specialist reaches for a shortcut in the first place: **63%** named squeezed finances as the biggest digital barrier, **56%** named a lack of skills or technical expertise, and **35%** said outright that they did not trust AI tools. Staff training was the top funding need for **44%**, yet only **17%** had received dedicated digital funding. Nobody trained a behaviour specialist to redesign a documentation workflow, so absent that training the fastest route to a complete-looking plan is the one already sitting in the folder from last year.

The same survey found **63%** of charities already using AI for administration and project management, evidence that drafting assistance has a real place in the sector's toolkit. It is not evidence that assistance belongs inside the part of a plan that states why a strategy is being used with a specific person, and the 2026 figures do not distinguish between organisations using AI to organise a meeting agenda and organisations letting it draft clinical formulation. That distinction has to be made locally, by policy, before any pilot starts.

Where technology helped - and where it stalled

A controlled system can gather approved fragments, flag missing review dates and produce a checklist for the meeting. It can also make stale text easier to reproduce at scale. Technology stalls when copied prose is treated as evidence, when objective data are absent, or when no owner has time to turn observations into a reviewed decision.

For a Behaviour Support Plan specifically, the honest limit is drawn around the formulation, not the formatting. Assistance can arrange confirmed observations under the right headings, flag a missing review date or draft the agenda for a review meeting. It cannot decide why a person behaves as they do, or whether a strategy that worked eighteen months ago still fits their current circumstances. A tool that produces a fluent-sounding formulation from old notes is not saving the specialist time; it is manufacturing the exact boilerplate-with-authority problem this brief exists to prevent.

That risk grows where nobody has drawn a line around what the tool may see. A third of charity boards were rated poor on AI skills in the 2026 survey, which in a behaviour-support context means a third of organisations cannot say with confidence whether a specialist has ever pasted a case history into a consumer chatbot to speed up a draft. The answer is not a policy nobody reads under deadline pressure; it is an approved route that is faster than the risky one, with a named reviewer who can see the source material the draft was built from.

The review date is not an AI problem, but the drift behind it is

Not every part of this workflow is waiting for an AI system to arrive. This brief's own account of how a plan gets written already located the fault line: a specialist reaches for last year's plan not because no tool exists to help them draft, but because nobody has set a working review date, or the plan has no

named owner willing to defend it at the next review meeting. That is an accountability gap, not a technology gap. A pilot that hands a specialist drafting assistance before anyone fixes the review calendar will only make a stale plan easier to reproduce.

The pattern that rewards an agentic layer sits above any single plan. Across a whole caseload, watching which plans are drifting past their review date, and which strategies keep being carried forward without a fresh observation, is a recurring signal a chatbot the specialist has to remember to open will never surface on its own. Connected agents reading that pattern across many plans can prompt the named specialist or clinical lead before a review lapses, not after a family notices the plan no longer fits. McKinsey's 2025 State of AI survey found most organisations using AI see little bottom-line effect because they layer a tool onto an unchanged process rather than redesign it, the same trap a vendor sets when selling drafting for the formulation section without learning that a PBS review cycle already exists. Gartner's 2026 analysis of abandoned generative AI projects points to the same root cause: unclear business value follows from skipping that discovery step. AIMonger's position is to place the agentic layer over the review cycle that already works, not instead of it.

A safer AI-assisted path

1. **Map how a plan changes after a material observation or incident.** No AI belongs in this mapping step: only the specialist doing the work can say which sections come from fresh observation and which are quietly carried forward from the last review. Trace the sources, handoffs, version used in practice and where family input is recorded, so the pilot targets drafting delay, evidence gathering, review or distribution rather than guessing.
2. **Build a small approved evidence pack for one plan type.** Separating observations, reported views, formulation, agreed strategies and outstanding questions is human design work; a retrieval layer can later surface the current approved fragment for each slot, but it cannot merge them or substitute for a missing observation. The source record must stay visible to the practitioner who reviews any arranged draft.
3. **Use assistance only for bounded tasks such as checking for missing sections, creating a meeting agenda or arranging confirmed content under an approved structure.** A qualified practitioner must decide the formulation, strategies and risk response. The person and family or carers need a real route to challenge the plan.
4. **Review a sample of plans after four weeks for individual detail, named owners and evidence of current review.** If the watching layer is already showing plans drifting past review dates or strategies copied forward without a fresh observation, a clinical lead should fix ownership and the review calendar in response to that signal. If staff still train from old copies or the plan cannot explain why a strategy is present, halt the pilot until the record system reflects those findings.

Where humans must intervene

The person, family or carers and multidisciplinary team remain central. Qualified practitioners own formulation and strategy; safeguarding and restrictive-intervention decisions follow the established governance route. No system should generate a behavioural explanation, risk decision or intervention without accountable professional review.

NIST's Generative AI Profile is a reasonable checklist to hold a behaviour-support pilot against: know what the system cannot verify (whether a strategy is still clinically appropriate is high on that list), keep the provenance of every observation the draft used, and treat a plan update as a consequential output requiring named human sign-off, every time, not just on the ones that look contentious. GDPR and, where it applies, the EU AI Act, still govern what special-category data about a vulnerable person may be typed into any tool in the first place.

Risks and failure conditions

The failure mode this workflow is most exposed to is boilerplate wearing clinical authority: a plan that reads as thorough because it is fluent, while actually carrying forward last year's strategy into a situation that has since changed, or smoothing over a genuine disagreement between the specialist and the family about what the behaviour means. A generic-sounding plan can give a stressed member of staff false confidence in exactly the moment they most need to notice that the plan does not fit.

A second risk is treating overall sector movement as proof this particular workflow has improved. The 2026 survey found **81%** of charities reporting some digital progress, a number broad enough to mean almost anything happened somewhere in the organisation. It says nothing about whether this specialist's plans are more current, whether families are more involved, or whether a stale strategy has been caught and retired. Only a direct audit of the plans themselves answers that.

Measures that matter

Audit the share of active plans with a named owner, family or carer involvement and a dated next review. Measure time from material observation to reviewed update, but pair it with senior sampling for personal strategies and cited observations. Completion alone is not proof that a plan helped.

Sit down with the specialists reviewing these plans at the same point each week and ask them directly whether a plan they signed off actually reflected current evidence, or whether it was close enough to pass review. Their honest answer, not the completion count, is what tells a manager whether the pilot has fixed the copy-paste habit or just made it faster to produce.

Decision questions

1. **What current evidence supports each active strategy?** If nobody can point to an observation from the last review period, the plan is running on inherited authority rather than current fact.
2. **Which document is the authoritative version in practice?** A pilot that cannot answer this has not solved the original problem; it has just added a better-looking copy to the confusion.
3. **Who decides that a change is clinically and ethically appropriate?** That person must be named, qualified, and reachable within the review cycle, not a title on an organisation chart.
4. **How are the person and family or carers actually involved, not just informed?** A plan drafted around them rather than with them will lose their trust exactly when it is needed most.
5. **What stops a generated summary from being read as the professional formulation?** Without a clear answer, the fastest-produced document will eventually become the one people trust, whether or not it deserves to.

FAQ

Can a tool personalise a Behaviour Support Plan? It can organise confirmed information, but personalisation is the work of people who understand the person, their context and the evidence.

What is the smallest safe pilot? A review-date and missing-section check on plans already owned by qualified staff, with no automatic changes to strategies or risk wording.

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.socialemmployers.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. NICE, Challenging behaviour and learning disabilities (NG11). <https://www.nice.org.uk/guidance/ng11>
9. NICE, Learning disabilities and behaviour that challenges: service design and delivery (NG93). <https://www.nice.org.uk/guidance/ng93>
10. BILD, Positive Behaviour Support resources. <https://www.bild.org.uk/positive-behaviour-support/>
11. McKinsey & Company, “The State of AI: Global Survey” (2025 edition). <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
12. Gartner, “Why Half of GenAI Projects Fail: Avoid These 5 Common Mistakes” (2026). <https://www.gartner.com/en/articles/genai-project-failure>

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