



From Research to Practice: The Pilots and the Four Questions They're Answering

Every pilot we run exists to answer one specific question we genuinely don't know the answer to yet — not to prove something we've already decided is true.

Gemma Gardner — Founder, AURIS Institute™ · Public Research Briefing · July 2026

A pilot programme is not a proof point. It's a live test of a specific question, run carefully enough that the answer — whatever it turns out to be — is worth trusting. Three partner organisations are currently running, or about to run, the AURIS pathway with their communities: Trafford Veterans CIC, delivering since 12 July 2026; Out in the Bay, working with LGBTQIA+ adults in Morecambe and the wider Bay community; and Red Rose Recovery, preparing a veteran-specific pathway for people in recovery from drugs and alcohol. Each sits inside a different population, and each is quietly testing a different edge of the same underlying framework.

1. Why We Build a Bespoke Programme for Every Partner

Every AURIS pilot uses exactly the same fixed core: the same five-pillar framework, the same six-week pathway shape, the same coaching-not-therapy stance, the same safeguarding architecture. What changes, deliberately, is the wrapper — the language, the case studies, the referral routes, and the specific risks each pilot is built to hold safely. This is not customisation for its own sake. It is how the same underlying model gets tested honestly against genuinely different populations, which is the only way to answer whether the framework is universal or military-specific in the first place.

Trafford Veterans CIC works with veterans navigating the standard identity-defining transition this whole programme was originally built around. Out in the Bay works with LGBTQIA+ adults in Morecambe and the wider Bay community — a population whose identity-defining transitions may have nothing to do with a uniform at all, which makes it one of the most important tests of whether our framework generalises. Red Rose Recovery is preparing a pathway specifically for veterans in recovery from drugs and alcohol — a higher-acuity cohort where the safeguarding net has to be heavier by design, and where the relationship between identity loss and substance use is itself part of what we're trying to understand.

2. The Emerging Sequence We're Testing

Early coding of the Identity After the Uniform interviews points to a pattern: where people describe genuinely rebuilding a stable sense of self after a role that used to define them, it tends to follow a broader sequence — rebuilding identity, reconnecting with a community, finding a purpose that gives the new identity somewhere to go, and only then arriving at something that feels like stability. We call this the Identity → Community → Purpose → Stability sequence.

We want to be precise about what that sequence currently is: an emerging pattern from 22 of 108 planned interviews, not a proven causal model. Several interviews depart from it in instructive ways — a small number of interviewees described identity as portable and largely intact through their transition, because their sense of self was never primarily built on the uniformed role in the first place; at least one interviewee had prepared for years in advance of leaving and reported no significant identity disruption at all, suggesting pre-emptive identity work may function as a protective factor sitting outside the sequence as currently framed; and for some interviewees, stability appeared to be sustained primarily through continued activity or role-replication rather than through the fuller sequence, raising an open question about whether that represents a genuinely different, equally valid pathway, or a less durable version of the same one. These disconfirming cases are treated as part of the evidence base, not noise to be excluded — good evidence practice means reporting them alongside the pattern, not just the interviews that fit it.

It's a strong enough signal to design a pathway around and test formally — which is exactly what the AURIS Framework and its six-week pilot pathways do — but it is not yet a validated finding, and we won't describe it as one until the fuller interview set says otherwise.

3. The Four Questions — and Why Each One Matters

Behind the pilot programmes sits a formal, tracked set of research questions we're deliberately trying to answer, not just deliver against. Each one is included here with why we think the answer matters — not as a claim of proven outcomes, but as the honest case for why the question is worth the cost of answering it properly.

RQ-001 — Does the Identity → Community → Purpose → Stability sequence hold across sectors, not just military?

Hypothesis: it reflects a general human pattern of identity-defining role loss, not a military-specific one. Current confidence: emerging — 22 of 108 planned interviews coded, including non-military comparators (disability, redundancy, addiction, public office). This is why Out in the Bay's LGBTQIA+ cohort and Red Rose Recovery's addiction-recovery cohort matter as much as the veteran-focused pilots: if the sequence only holds for veterans, the framework should stay narrow. If it holds across all three, that's a much stronger signal, worth testing rigorously rather than assuming.

Why it matters: if a single, teachable sequence genuinely applies across very different populations navigating identity-defining transitions, that is a preventative model worth scaling carefully — support that could, in principle, be offered earlier and more cheaply than crisis-stage intervention. That potential is exactly why it needs to be tested honestly rather than assumed.

RQ-002 — Does pre-emptive identity work before transition reduce or prevent identity disruption?

Hypothesis based on a single early case, referenced above as one of our disconfirming interviews. Current confidence: hypothesis only — one data point is not evidence.

Why it matters: this is the most directly preventative of the four questions. If identity work done before someone leaves a defining role measurably reduces the disruption that follows, the case for early, low-cost, pre-transition support — rather than downstream crisis response — becomes considerably stronger. We

treat that as a real possibility worth testing, not yet as a demonstrated saving to any service, public or private.

RQ-003 — Are internally generated regulation strategies more durable than externally anchored ones?

Current confidence: emerging, weakly supported. The general psychological mechanism is well understood in behavioural science — self-efficacy and belief in one's own capability to act are established constructs — but whether internally-generated strategies specifically outlast externally-anchored ones in this population, under this framework, is not yet established.

Why it matters: strategies that don't depend on an external structure (a unit, a role, a routine) staying in place are, by definition, more sustainable — they don't collapse when circumstances change again. If confirmed, this has direct implications for how any future support model, ours included, should be designed: to build portable capability, not dependency on a programme.

RQ-004 — Is non-clinical coaching an appropriate and effective delivery model for identity transition specifically, compared with therapy?

This question is newly formalised from Gap G028 in the Master Evidence Repository, and should be added to the Research Questions Register (sheet 38) accordingly.

Hypothesis: non-clinical, goal-directed coaching is an appropriate and acceptable delivery model for identity-transition support in this population, though not a substitute for clinical care where clinical need exists. Current confidence: partially supported — real, peer-reviewed evidence exists that coaching is a recognised, evidence-based approach for non-clinical, goal-directed populations, that coaching is associated with promising if limited behaviour-change outcomes generally, and that UK veterans find peer-led, non-clinical support acceptable and valuable alongside clinical care. No study directly comparing coaching against therapy for identity transition specifically has been located.

Why it matters: every AURIS pilot draws a firm line — this is coaching, never clinical or therapeutic support; facilitators signpost, they don't treat, and anyone in crisis is routed to appropriate services immediately. That boundary is both an ethical safeguard and, in principle, a scalability advantage: non-clinical, peer-supported models can potentially reach far more people than clinical capacity alone ever could, at lower cost per person supported. Whether that potential holds up under direct comparison is exactly what remains unanswered — and exactly why we're not claiming it does yet.

4. Inside a Pilot: The Six-Week Design

Every AURIS pilot follows the same underlying shape, regardless of partner or population. A Week 0 covers baseline measurement, onboarding, ground rules and safety routing — confirming each participant's treatment and crisis contacts before any group work begins, and establishing explicitly that the pathway runs alongside any existing support, never instead of it. Weeks 1 through 6 then map directly onto the five pillars: Awareness in week one, naming what has actually shifted; Understanding in week two, making sense of the patterns that have emerged since; Regulation in week three, building steady, portable tools before asking anything harder of the mind; Identity across weeks four and five, first rebuilding and then actively applying a grounded sense of self; and Stability in week six, consolidating the practical and psychological

footing — housing, routine, ongoing support, relationships — that makes the new identity sustainable rather than theoretical. A closing review stage repeats the baseline measures, agrees an onward plan, and signposts to further support as needed.

Each individual session follows the same standing internal structure too: a welcome and safety check, where anyone not safe to continue that day is routed out immediately; a review of the previous week alongside any app-based reflection; the core pillar content and group work itself; a practical, stability-focused action; and a closing safety check before the session ends. This consistency matters for evaluation as much as for delivery — without a fixed structure, it becomes very difficult to assess fidelity, one of the five implementation-science measures every pilot is evaluated against.

7.1 Referral, screening and the route-out principle

No one enters an AURIS pilot without a structured referral and stability screening process, and the specific screening criteria are adjusted, deliberately, per partner and population — this is one of the clearest examples of the bespoke wrapper around the fixed core. A veteran-focused pilot screens for recovery-stability and readiness for group identity work; a higher-acuity cohort, such as the veterans-in-recovery pathway prepared with Red Rose Recovery, adds a heavier safeguarding layer by design, including on-site naloxone provision and trained staff before delivery can begin at all, reflecting the elevated overdose and suicide risk documented in that population specifically. In every case, the same underlying principle applies: the group never holds acute risk. Anyone in crisis, active withdrawal, or unmanaged risk is routed to the appropriate treatment or crisis service first, with a clear, non-punitive route back in once they are stable. Relapse, where relevant, is treated as data to learn from, not as automatic exclusion.

5. How Implementation Science Actually Evaluates a Pilot

"Pilot delivery is not proof" is not just a caution we repeat — it is a specific methodological commitment, borrowed from implementation science, that shapes exactly what our pilot evaluations measure. Rather than asking only "did people like it" or "did people attend," each pilot is assessed against five distinct domains.

8.1 The five evaluation domains

- **Acceptability:** whether participants and partner staff find the pathway appropriate, tolerable and worth continuing
- **Feasibility:** whether the pathway can actually be delivered within the partner's real-world constraints — venue, staffing, referral volume
- **Fidelity:** whether delivery matches the Programme Design Map as written, with any variation recorded and explained rather than silently drifting
- **Adoption:** whether the partner organisation takes up and continues to refer into the pathway beyond the pilot cohort itself
- **Participant outcomes:** the pre/post measures described in our evaluation toolkit — wellbeing, recovery capital where relevant, housing status, and an AURIS identity self-rating across the five pillars

Only once a pilot has been assessed across all five domains, honestly, including where it falls short, does its evidence count toward answering any of the four research questions above. A pilot that delivered smoothly but shows poor fidelity to the model, for instance, tells us something important — but it does not yet tell us whether the model itself works, only that this particular delivery of it may not be a fair test of it.

6. Why the Pilots Are Free, and Why That Matters for the Evidence

Every current AURIS pilot is delivered at no cost to the partner charity, with lifetime access to the AURIS app gifted to each participant. This is not simply a generosity decision — it removes cost as a confound in evaluating acceptability and adoption. A charity that continues to refer participants into a free pilot, and a participant who continues to engage with a pathway that costs them nothing, are both giving a cleaner signal about whether the model itself has value than either would if payment were involved from the outset. It also reflects a deliberate staged commercial approach: prove the model through funded, free pilots first; only then move towards paid delivery, licensing and corporate partnerships, once there is real evidence — not assumption — behind the case for them.

6.1 How the funding pathway connects to all four questions

None of these questions get answered by research alone — they get answered by delivery, evaluated properly, and delivery has to be funded. Every AURIS pilot is delivered free to the partner charity, funded by grant income rather than charged to the charity or its community. That funding pathway is not incidental to the research programme; it is how the research programme gets its data. Every funded programme generates the same evaluation evidence — acceptability, feasibility, fidelity, adoption, participant outcomes — that then feeds directly back into answering the four questions above, and into refining the framework itself, its facilitator training, and the app's guided pathway content.

This is also why bespoke programme design matters commercially, not just scientifically: a funder assessing a grant application wants to see that the money produces evidence, not just delivery. A pilot built to the same fixed core but wrapped for a named, specific population, with evaluation built in from the start, is a fundamentally stronger funding case than a generic programme delivered without a research question attached to it.

7. Why We Don't Yet Put a Number on the Savings

It would be easy to attach a headline pound figure to all this — to say a preventative, identity-focused model saves the NHS, or any other system, a specific sum of money. We're not going to do that, because the evidence to support a specific figure doesn't yet exist for AURIS specifically. What does exist is a well-evidenced general picture: identity disruption after major transition carries real downstream costs — in wellbeing, in relationship stability, in long-term engagement with work and community — and untreated, those costs land somewhere. Establishing whether, and by how much, an identity-focused preventative model changes that picture is precisely what our pilot evaluations and the four research questions above are designed to find out. Until they have, any specific savings figure would be a promise the evidence can't yet back.



8. What Comes After the Pilots

Each pilot ends with a collaborative evaluation report, co-produced with the partner organisation and named accordingly — not a private internal readout. Where the evidence supports it, that becomes the basis for scaling, including into the AURIS Identity System membership app itself, whose guided pathway and AI-supported reflection content are directly built from and refined by this research. Where the evidence doesn't support it, that becomes a finding too, and the framework gets revised rather than the result being quietly dropped.