

🏆 TSA 2026 DIGITAL INNOVATION AWARD WINNER — ALCOVE, YOKERU & ENOVATION

PILOT EVALUATION REPORT · MARCH 2026

AI-Enabled Inbound Triage for Alarm Receiving Centres — Pilot Evaluation

A collaborative pilot between Alcove, Yokeru and Enovation

DATE	PREPARED BY	STATUS
March 2026	Alcove, Yokeru & Enovation	Pilot Complete — Scaling Phase

446 TOTAL CALLS TRIAGED BY AI	54.9% CLOSED BY AI WITHOUT OPERATOR	100% CALLS HANDLED CORRECTLY	Zero MISSED EMERGENCIES	300 LIVE PILOT USERS
---	---	--	-----------------------------------	--------------------------------

PARTICIPATING ORGANISATIONS

Alcove

Alcove provides technology-enabled care services to over 20,000 people nationally and operates a TSA QSF-accredited Alarm Receiving Centre delivering 24/7 monitoring and response services.

Yokeru

Yokeru builds AI-powered care infrastructure that enables care providers and Alarm Receiving Centres to safely automate routine interactions, triage demand, and deliver more responsive, personalised support at scale.

Enovation

Enovation provides UMOx, a leading Technology Enabled Care platform designed for 24/7 monitoring centres supporting vulnerable people in their own homes and sheltered accommodation. UMOx connects devices, telephony, monitoring workflows, and care records, enabling secure, reliable, and scalable service delivery.

SECTION 1

Executive Summary

A proven, standards-aligned approach to reducing avoidable demand in UK ARCs

This report evaluates the AI-enabled inbound triage pilot conducted by Alcove, Yokeru and Enovation, with independent oversight from TEC Quality aligned to the Quality Standards Framework (QSF).

The pilot introduced conversational AI triage into the live alarm-handling workflow at Alcove's Alarm Receiving Centre (ARC). When a pendant or device triggers an alarm, the AI voice agent ("Amy") answers the call, determines whether it is a genuine emergency through structured conversational triage, and either safely closes the call or immediately escalates it to a human operator.

The approach builds on prior deployments in the United States—where formal standards are still emerging—and is further informed by staged trials that generated critical operational learnings. These include how AI must interoperate with existing call-handling systems and the essential role of secure fallback to human operators. Together, these insights position the pilot to inform the development of UK and European policy, as well as contribute to emerging global standards for TEC.

HEADLINE RESULTS — AS OF 13 MARCH 2026

446 calls have been presented to the AI triage agent. Of these, 245 (54.9%) were successfully resolved without operator involvement, while 201 (45.1%) were appropriately escalated for valid reasons, including true emergencies, no speech detected, inability to ascertain a false alarm, carer requests, and test calls. **100% of calls were handled correctly, with zero missed emergencies.**

The pilot has now expanded to 300 live users and is scaling towards becoming the default approach for handling alarm activations across Alcove's ARC.

SECTION 2

Industry Context & Challenges

Why Alarm Receiving Centres need to work differently

Alarm Receiving Centres are under sustained pressure from avoidable demand. Despite decades of digital transformation across the sector, the way ARCs operate has not materially changed. Skilled operators still answer every call, regardless of whether it is a genuine emergency or an accidental press.

Over 50% of calls are avoidable More than half of all inbound calls to ARCs are false alarms or test calls that require no operator intervention. In one reported week, Alcove recorded 14% of all calls from falls detectors alone as false alarms.

40% of operator time on false alarms Every false alarm that reaches an operator takes time away from genuine emergencies and proactive work. This creates queue pressure, slows response times, and reduces the capacity available for preventative care.

Rising demand, constrained resources The analogue-to-digital transition is generating more actionable data from connected devices, increasing inbound volume. ARC workloads have more than tripled, yet funding has not kept pace. Some centres are struggling to stay afloat.

This is an opportunity to intelligently partition and route different devices and calls based on priority, including leveraging alternative pathways (such as DMPs) for non-urgent data rather than routing all demand through ARC workflows.

SECTION 3

Why We Did This

Creating immediate relief while building for the future

Alcove recognised that the traditional ARC model, where every alarm activation is answered by a human operator regardless of urgency, was not sustainable. The goal was twofold:

Immediate relief: Remove avoidable inbound calls from the operator queue without compromising the safety of any individual. Free up operator capacity for genuine emergencies and proactive, preventative work.

Industry-wide change: Create an evidence base aligned to the TEC Quality Standards Framework that could inform safe AI adoption across UK ARCs. Establish a robust evidence base to inform the future adoption of AI through the TSA Sector Risk and Innovation Group (SRIG), ensuring alignment with the TEC Quality Standards Framework (QSF)

This approach has already been proven in the United States, where Yokeru's inbound triage has been deployed across comparable ARCs, with approximately 30% of eligible calls handled safely without human intervention. Alcove is the first provider to adopt this approach in the UK.

TEC Quality and TSA provided independent oversight throughout the pilot, emphasising the need for embedded safety and compliance from the outset. This was supported by an initial kickoff meeting and ongoing touchpoints, to ensure that learnings are considered in future development of the QSF and that the pilot was conducted in line with recognised industry standards.

SECTION 4

What We Did

Embedding AI triage into the existing workflow

Rather than introducing a separate system, the pilot embedded conversational AI triage directly into the live alarm-handling workflow using the existing platform infrastructure. This was made possible by bringing together three partners with distinct, complementary capabilities.

Yokeru: Yokeru provides the conversational AI voice agent (Amy) and the capability to rapidly iterate and improve the triage logic. Yokeru's platform powers the AI layer that sits within the existing call flow, handling structured conversational triage in real time.

Enovation UMO: Enovation provides the UMOx connected care platform, which is the integration layer connecting alarms, telephony, sensors and operational records. UMOx manages the call routing, recording and audit trail that underpins the entire workflow.

Alcove: The ARC operator provided operational reality, service user governance, and clinical oversight. Their deep understanding of day-to-day call handling ensured the AI was configured for real-world operational conditions, including variability in call types, user behaviour, and escalation pathways.

Background of learnings:

The pilot was delivered in a structured, phased approach to ensure safety, resilience and operational readiness before scaling:

- **Build & integration:** Embedding AI triage into the existing ARC workflow required integration between Yokeru and Enovation UMOx. A core requirement was eliminating any single point of failure, which introduced additional complexity—particularly in call routing and fail-safe design—and resulted in approximately one month of additional development time.

How the call flow works:

1 An alarm is triggered

AI triage currently handles wearable-triggered alarms (e.g. pendants, falls detectors, GPS trackers, and push-button devices), which represent the majority of false alarms. High-risk alert types such as smoke alarms, fire alarms, and medical monitoring alerts are excluded and always routed directly to operators.

2 The call is routed to Amy (the AI triage agent) via Enovation's UMOx platform

3 Amy answers and conversationally ascertains whether it is a genuine emergency using structured questioning and double confirmation.

4 Resolution or escalation

If the caller confirms they are safe, Amy closes the call, provides reassurance, and records the full outcome and transcript in UMOx. If there is any uncertainty, no response, distress, or a genuine emergency, the call is immediately escalated to a human operator within seconds.

Post-call monitoring:

When a call handled by the AI agent is completed, a near real-time notification is shared with key stakeholders across Yokeru, Enovation and Alcove to enable ongoing monitoring. Each notification includes a transcript, call recording, and a structured summary of the outcome and call disposition.

Possible KPIs:

IMPLEMENTATION PERIOD:

- Daily review of all call transcripts
- Minimum of 50% of calls formally reviewed ('human in the loop')

BAU (BUSINESS-AS-USUAL) PERIOD:

- Ongoing review of 5–10% of calls for continuous learning and quality assurance

KPI considerations for call quality:

- Secondary KPI: target of >10 (to be defined and refined)

Repeat callers who are escalated, even in non-emergency scenarios, should be identified and, where appropriate, routed to alternative pathways rather than repeatedly handled through AI triage.

Secondary KPIs: Response Time Breakdown

Call response time is currently measured as a single end-to-end metric (from alarm activation to call resolution), in line with the QSF. However, to better understand the impact of AI triage, the following secondary KPIs are looked at as the pilot progresses into an established automation.

**The numbers are based on Alcove's performance:

- **AI answer time** — time from alarm activation to Amy picking up the call (AI agent Amy picks up within 3 seconds of alarm activation)
- **AI handling time** — duration of the AI triage interaction (triage is completed in an average of 25–30 seconds)
- **AI-to-operator transfer time** — time between AI transferring the call and an operator picking it up (average of 35 sec)

SECTION 5

How We Did It

A rigorous, phased approach to safe deployment

The pilot followed a structured methodology designed to build confidence at every stage, with clear criteria for progression and the ability to pause or reverse at any point.

1 8-week staff-led testing phase

Before any service user involvement, Alcove's team ran repeatable scenario testing over 8 weeks. This covered false alarms, equipment test calls, non-response events, ambiguity cases, and defined escalation triggers (including GPS "lost" alerts). All test outcomes were logged, reviewed and formally signed off. Claire Aldridge set the threshold: "Continue refining and testing until 100% accuracy is achieved prior to implementation."

2 AI agent refinement and iteration

Where the agent needed to be more predictable, Yokeru reduced variability. The aim was consistent questioning for the same scenario, improving trust and reviewability. Clear boundaries were defined for where the AI must not resolve the call (e.g. GPS "lost" scenarios must always be escalated to a human operator). The design philosophy was to automate only what is clearly safe and escalate everything else.

A key learning during this phase related to the integration between Yokeru and Enovation UMOx. Ensuring there was no single point of failure across the end-to-end call flow introduced additional complexity, particularly in call routing, signalling, and fail-safe design. Early testing uncovered routing edge cases and integration challenges, requiring iterative refinement and resulting in approximately one month of additional development time.

These findings emphasised the importance of robust operational guardrails, including secure fallback to human operators, clearly defined escalation pathways, and reliable call handling through appropriate signalling (e.g. DTMF) and consistent use of preferred routing methods such as SIP.

This phase also highlighted the need to validate not only individual system components, but the interaction between platforms. Effective performance depended on how Yokeru and UMOx operated together within the live call flow, rather than in isolation.

A key implementation learning is that integration pathways should be assessed early and timelines set conservatively, particularly where resilience and fail-proof operation are critical. Confidence in the system was built through structured testing and progressive refinement, rather than assumption.

Key Learnings & Guardrails:

Integration complexity and resilience

Ensuring there was no single point of failure across the integration between Yokeru and Enovation UMOx introduced additional complexity, particularly in call routing and fail-safe design. Resolving integration edge cases and achieving full resilience required iterative refinement and resulted in approximately one month of additional development time.

LEARNING

Integration pathways should be assessed early and timelines set conservatively. Where fail-proof operation is required, additional complexity and unforeseen delays should be expected and planned for.

Call routing and platform interaction

Early testing identified challenges in call routing and platform-to-platform interaction, requiring refinement of routing logic and signalling (e.g. DTMF) to ensure reliable call handling and closure.

LEARNING

It is critical to design and test both routing logic and system interoperability, not just individual platform performance.

Guardrails and safe fallback

Clear guardrails were established to ensure safe operation, including defined escalation pathways and secure fallback to human operators in all uncertain or high-risk scenarios.

LEARNING

AI triage should only automate scenarios that are clearly safe, with all ambiguity or risk escalated to human operators.

Confidence through staged testing

Structured testing and staged rollout were essential in identifying and resolving issues prior to live deployment. This approach built confidence in system performance and safety.

LEARNING

Confidence in AI systems is built through progressive testing and refinement, not assumption.

Call routing and platform interaction:

Early testing identified challenges in call routing and platform-to-platform interaction, requiring refinement of routing logic and signalling (e.g. DTMF) to ensure reliable call handling and closure across systems.

LEARNING

It is critical to design and test both routing logic and system interoperability, including signalling mechanisms, rather than treating platforms in isolation.

SECTION 5A

Built-In Safeguards

Safety architecture underpinning the pilot

Fail-safe fallback: If the AI cannot connect, the call disconnects, or anything unexpected occurs, the call is automatically routed to a human operator. The system is designed with no single point of failure, and incorporates double confirmation to minimize the risk of incorrect classification.

This approach balances efficiency with safety: the AI reduces the burden on staff by filtering routine calls, while maintaining a controlled and minimal risk profile. In contrast, many ARCs today rely on overstretched teams handling repetitive tasks, where errors and delays are already a known issue.

Risk-averse by design: The AI is deliberately configured to escalate in cases of uncertainty. In early stages, this resulted in some calls being escalated that could potentially have been resolved by the AI; however, this approach reflects a safety-first principle, where over-escalation is the preferred starting position.

Exclusion criteria: Service users with safeguarding concerns, smoke detectors, epilepsy/seizure risks, or high-risk categorisation are excluded from AI triage. Smoke and fire detection alerts are always routed directly to a human operator.

Dynamic suitability review: Eligibility for AI triage is continuously reviewed and can be adjusted as user needs and risk profiles change. AI triage is removed immediately if safeguarding concerns arise, the risk of genuine emergency increases, or patterns of escalation indicate that the pathway is no longer appropriate.

While current identification of these patterns is supported by manual review, this establishes the foundation for more advanced, data-driven monitoring over time. The approach ensures that AI use remains adaptive, reversible, and aligned to individual risk.

SECTION 6

Outcomes & Results

Evidenced performance as of 13 March 2026



Of the 201 calls (45.1%) that were escalated to a human operator, all were escalated for appropriate reasons:

Escalation Reason	Description	Assessment
True emergency	Service user indicated they needed help	Correct escalation
No speech	No response detected from the caller	Correct escalation
Unable to ascertain	AI could not confidently determine false alarm	Correct escalation
Carer request	Information requested about carer attendance	Correct escalation
Test call	Equipment test requiring operator confirmation	Correct escalation

Every escalation was appropriate. The AI did not miss a single genuine emergency. The 54.9% resolution rate represents calls where the AI correctly identified and safely closed false alarms, while the 45.1% escalation rate demonstrates the system's deliberate risk-averse design.

SECTION 6A

Earlier Pilot Data

Initial 10-user controlled trial results

The initial controlled pilot with 10 service users provided the first quantified evidence of impact:

- **71% of false alarms resolved without human operator
- 2.56 hrs operator time saved per month (10 users)

**The 71% false alarm resolution rate observed in the initial 10-user pilot reflects a small, highly targeted cohort and should be interpreted as directional rather than representative.

Users were selected based on high-frequency false alarms, resulting in a concentration of low-risk, AI-suitable calls, alongside more controlled pilot conditions.

This led to higher resolution rates than would be expected at scale, but provided clear proof of concept that AI triage can safely resolve a meaningful proportion of calls.

SECTION 6B

Expanded Cohort Impact

Results across the 300 service user cohort

As the pilot expanded to 300 live service users, the additional operator capacity freed up by the AI was reinvested into proactive and preventative monitoring. This resulted in a 22% increase in onward proactive referrals. Operators were able to spend more time analysing trends, reviewing frequent callers, and identifying patterns of decline.



SECTION 6C

Scalability Projections

What this means at scale

If these results were expanded across Alcovē's total number of ARC customers, the projected impact on operator capacity and workforce efficiency is significant:

Metric	Value
Operator hours saved per customer per month	0.256 hours
Estimated total operator hours saved per year	Up to 17,000+ hours
Estimated workforce efficiency savings per year	Over £300,000

These savings would be reinvested into faster emergency response and more proactive, preventative, data-led support, without additional staffing or funding.

SECTION 7

Next Steps

From pilot to standard operating procedure

The pilot has demonstrated that AI triage can safely and effectively reduce avoidable demand within a UK ARC. The evaluation and expansion to 300+ service users are complete. The next phase focuses on three priorities:

- 1

Roll out across the full Alcovē ARC
Scale AI triage to become the standard approach for handling eligible alarm activations across Alcovē's full ARC. The graduated approach has provided the evidence and confidence needed to make this the default for how alarm activations are handled. Rollout will continue to follow the same consent processes, monitoring rigour, and dynamic suitability reviews that underpinned the pilot.

Early transcript and reporting analysis indicates that AI triage is already reducing overall call handling time in a subset of cases, particularly for non-urgent and false alarm scenarios. Currently, call time is measured as a single end-to-end metric, from alarm activation through to the call being answered and handled.

SECTION 8

Summary

A responsible, proven model for safe AI adoption in UK ARCs

This pilot demonstrates that conversational AI triage can be deployed safely and effectively within a UK Alarm Receiving Centre, reducing avoidable demand without compromising care quality or putting service users at risk.

The results speak clearly: 446 calls triaged, 54.9% resolved without operator involvement, 100% handled correctly, and zero missed emergencies. Every escalation was appropriate. Every safeguard worked as intended.

The approach taken, with structured testing, phased rollout, service user consent, continuous monitoring, and independent TEC Quality input and TSA, provides a replicable model for safe AI adoption across the sector. It is not experimental. It is a working system producing measurable results with real service users.

What this means for the TEC community:

FOR SERVICE USERS

Routine alerts handled quickly and calmly. Genuine emergencies get faster human attention.

FOR ARC OPERATORS

Fewer avoidable calls. More capacity for emergencies and proactive, preventative work. Better use of skilled workforce.

FOR COMMISSIONERS

Outcomes-driven services where more time is spent meeting genuine community needs, while remaining cost-effective.

FOR THE SECTOR

A responsible model for AI-assisted call handling that demonstrates how AI can be adopted safely and transparently, and which will contribute to standards development in partnership with TEC Quality as the industry's own standards body.

Alcove, Yokeru and Enovation look forward to continued collaboration with TEC Quality, partners and other providers to share learnings and best practice, and to bring the future of AI in care to the sector safely and responsibly.

Contacts:

Claire Aldridge

Head of Client Development & Clinical Oversight, Alcove
claire.aldridge@alcove.care

Monty Alexander

CEO, Yokeru
monty@yokeru.io

Paul Shead

Managing Director, Enovation
paul.shead@enovationgroup.com