

Retail Energy Code Company

info@retailenergycode.co.uk
www.retailenergycode.co.uk

By email only: AIPolicy@ofgem.gov.uk

20 March 2026

RECCo response to: AI Technical Sandbox Consultation

We welcome the opportunity to respond to this consultation. Our non-confidential response represents the views of the Retail Energy Code Company Ltd (RECCo) and is based on our role as operator of the Retail Energy Code (REC). RECCo is a not-for-profit, corporate vehicle ensuring the proper, effective, and efficient implementation and ongoing management of the REC arrangements. We seek to promote trust, innovation and competition, whilst maintaining focus on positive consumer outcomes.

RECCo has an interest in this agenda from two connected perspectives. First, through the REC and the data and market services delivered under it, we are directly engaged in how AI and automation may interact with shared industry processes and datasets. Second, as the Code Manager of the REC, which itself contains sandbox and related derogation arrangements, we have a practical interest in how innovation can be tested safely, governed appropriately, and, where successful, translated into enduring code or service change.

While the appendix to this letter provides our response to each consultation question, we highlight the following points:

- **Support for the AI Technical Sandbox pilot:** We welcome Ofgem's proposal for a 12-month pilot to enable safe, controlled testing and evaluation of AI and automation in the energy sector, with sector-wide learning.
- **This is a live issue for RECCo and REC Parties:** We have received applications for derogations and sandbox trials to enable robotics/automated processing to interact with REC services (including SDES/SDEP), illustrating the need for clear guardrails for machine-to-machine use.
- **Safeguards for automated interactions should be explicit:** In addition to ethics, the sandbox should test operational resilience (performance/load), security, auditability, and the ability to suspend/rollback automated accounts where issues arise.
- **Clear pathways to adoption are essential:** Each use case should set out how learnings translate into implementation (including any required licence/code changes), with early engagement of code bodies and market service providers to avoid successful trials stalling.

We are happy to discuss any of the points raised in this response.

Yours sincerely,

Jon Dixon
Director, Strategy and Development

Appendix: RECCo response to consultation questions

<p>Q1. Do you agree with the proposed eligibility criteria for lead Participants (licensees, market participants, and operators of essential services) and the encouragement of partnerships with technology providers, academia, and other innovators? Please explain your reasoning.</p>
<p>RECCo broadly agrees with the proposed approach that lead Participants should be licensees, market participants, and operators of essential services, with partnerships encouraged. This balances innovation enablement with clear accountability and enforceability.</p> <p>However, we also encourage Ofgem to make explicit in the participation guidance that code bodies and market service/data providers (including any relevant service providers) are expected partners for relevant use cases. Many high-value AI applications may depend on code-governed processes, shared services and market data flows.</p> <p>RECCo has already received applications for derogations and sandbox trials to enable robotics/automation to interact with REC services (including SDES/SDEP). This experience underlines the need for clear responsibility for third-party AI agents and for specific controls for automated accounts (identity and access management, audit logging, rate limiting, and a kill switch/suspension mechanism).</p> <p>We also recommend that Ofgem also consider mechanisms to encourage participation from smaller market entrants and non-traditional innovators, ensuring the sandbox is accessible to a broad range of stakeholders, including those who may lack extensive resources.</p>
<p>Q2. Are the proposed use case selection criteria (including commercial neutrality, innovation, sector impact, regulatory uncertainty, testability, governance, and data access) appropriate and sufficient to ensure a fair and transparent process? Are there other criteria, safeguards, or considerations we should include?</p>
<p>We consider that the proposed selection criteria are appropriate and provide a strong basis for a fair and transparent process.</p> <p>We suggest strengthening the criteria and associated safeguards to reflect emerging machine-to-machine use of market services and to support consistent decision-making:</p> <ul style="list-style-type: none"> • Mandatory impact assessment: expected consumer outcomes, sector benefits, costs and risks (including distributional impacts). • Operational resilience and performance: testing for load/performance impacts on shared services, failure modes, rollback, incident response, and the ability to suspend automated access where required • Data governance: clear plans for representative, anonymised and/or synthetic datasets; data minimisation; security controls; and lawful/ethical use of data. • Implementation pathway: from the outset, identify any required changes to licences, industry codes and/or market service arrangements (with early engagement of relevant code bodies). • Model governance: clarity on model versioning, change control, monitoring for drift and re-validation where models update over time.
<p>Q3. Is the proposed approach for the AI Technical Sandbox clearly distinct and complementary to other initiatives such as Ofgem’s AI Reg Lab, Energy Regulation Sandbox, Future Regulation Sandbox, UKRI-funded and SIF/NIA initiatives, NESO, FCA regulatory sandbox experience, and DSIT AI Growth Lab? Are there other relevant initiatives or examples of best practice that Ofgem should consider, and if so, which ones?</p>

We agree that the proposed AI Technical Sandbox is distinct from, but complementary to, Ofgem's other initiatives. In particular, we see a clear difference between: (i) a technical environment for designing and testing AI systems; (ii) the AI Reg Lab's policy and ethical work; (iii) the Energy Regulation Sandbox, where time-limited live regulatory flexibilities may be needed; and (iv) the Future Regulation Sandbox, where evidence points to the need to trial potential rule changes before wider adoption. We also agree that alignment with NESO, UKRI, SIF and NIA activity is important so that datasets, tooling, evaluation approaches and lessons learned can be reused rather than recreated.

To make that complementarity work in practice, Ofgem should adopt a more joined-up coordination and routing approach across these mechanisms. In particular, there should be a simple front-door or triage process, supported by a common core information set, so that applicants can be directed to the right route at the right stage and are not required to duplicate materially the same application, evidence or governance information across multiple initiatives. This would reduce friction for innovators and support a clearer progression from technical testing, to live trial where necessary, to any subsequent regulatory or code change.

We also suggest that Ofgem considers the role of sandboxes, trial arrangements and derogation mechanisms operated by code bodies such as RECCo and Elexon. In some cases, an applicant may need access to code-governed datasets, industry processes or operational services provided through those bodies before a use case can be tested meaningfully. RECCo's earlier response to the call for input noted that code managers may need to support parallel code sandbox activity, facilitate access to relevant datasets and processes, help develop anonymised or synthetic ("dummy") data, and ensure that test environments reflect real market arrangements. This is particularly important given the wider challenges around fragmented, proprietary and restricted data access identified in the government's recent call for evidence on data for AI in the energy system. Ofgem should therefore ensure that coordination extends beyond its own initiatives to relevant code bodies and data holders, so that applicants can access the right route, the right data and the right governance framework in a joined-up way.

Q4. Does the proposed governance structure (steering group, working groups, open forums) provide sufficient oversight, transparency, and opportunities for stakeholder engagement? Are there other mechanisms or safeguards that should be included to ensure effective governance and knowledge sharing?

We welcome the proposed governance structure and agrees that a Steering Group, supported by working groups and open forums, provides a strong foundation for oversight, transparency and stakeholder engagement. We agree in particular that a representative Steering Group can play an important role in advising on governance and sector engagement, reviewing applications and monitoring progress, while working groups and open forums provide a practical route for wider technical, operational and regulatory input. This is consistent with the consultation's wider emphasis on transparent processes, regular reporting and publication of non-confidential learning.

To strengthen that framework further, the Steering Group could also undertake part of the coordination and routing role described in response to Q3. In practice, the AI Technical Sandbox will need to interact with a number of adjacent initiatives and delivery bodies as use cases mature, including the AI Reg Lab, the Energy Regulation Sandbox, the Future Regulation Sandbox, and relevant innovation programmes. A Steering Group with visibility across those routes could help ensure proposals are directed to the most appropriate mechanism, that learnings are transferred efficiently between initiatives, and that applicants are not asked to duplicate materially the same information across multiple processes.

We also agree that robust conflicts of interest provisions will be essential. However, those provisions should be designed in a way that is proportionate and workable in practice. Given the proposed model, the organisations best placed to bring relevant operational, regulatory, code, consumer, technical, data and cyber expertise are also likely, at some stage, to be direct or indirect participants in trials, potential applicants, or partners to applicants. We therefore support clear declarations of interest, case-by-case assessment, and recusal where a specific conflict arises, rather than rules that would unnecessarily prevent knowledgeable organisations from participating altogether. This would be consistent with the consultation's proposal that members act independently and that detailed recusal provisions are set out in the Terms of Reference and participation agreements.

It may also be appropriate for the Steering Group, or another suitable cross-sector body, to identify and bring forward priority models or use cases for testing where there is a clear sector-wide question but no single organisation is well placed, or sufficiently incentivised, to propose the trial itself. In some cases, waiting for a prospective implementer to apply may delay useful learning, particularly where commercial sensitivities or intellectual property considerations could discourage firms from bringing forward a model on an openly shareable basis. For example, there may be value in testing how AI could support or improve auto-switching arrangements, with the learning captured and published in a non-proprietary form for wider sector benefit. Any such role would, however, need to operate within clear conflict of interest, confidentiality and recusal arrangements.

Q5. Are the proposed next steps for developing and launching the pilot clear, and is there anything further we should consider as we refine the timeline?

RECCo considers the proposed next steps are broadly clear and supports Ofgem's intention to launch the pilot by Autumn 2026, followed by a 12-month pilot with regular reporting and evaluation.

To support timely delivery and reduce avoidable delay, Ofgem could:

- **Provide greater upfront clarity and support** through a pre-application clinic, early triage, clear milestones and early publication of standard templates.
- **Use a phased and agile delivery approach** with cohort windows, iterative review cycles, and enough lead time for data access, procurement and integration with shared market services.
- **Align the pilot with wider regulatory and code processes** so applicants are routed to the right mechanism early and successful trials can move efficiently into service, code or regulatory change.

That would give participants greater certainty, support faster learning, and improve the prospects of successful trials translating into practical outcomes.

Q6. Does the consultation and proposed pilot sufficiently address ethical considerations (fairness, transparency, responsible use, consumer trust) in line with Ofgem's AI guidance? Are there further steps we should take to embed ethics and safety in the sandbox?

We agree that the consultation and proposed pilot address the main ethical issues, including fairness, transparency, responsible use and consumer trust. In particular, the proposed ethics and safety checkpoints,

together with Ofgem's intention to embed ethical assurance throughout design, testing and evaluation, provide a sound foundation.

To strengthen this further, we suggest:

- requiring a proportionate AI impact assessment for each use case at the outset, covering consumer harm, bias, explainability, accountability and operational risk;
- setting clear data governance, privacy and security requirements, including alignment with UK GDPR and appropriate controls over access, retention and auditability;
- including mechanisms for consumer input where relevant, particularly for use cases that may affect customer outcomes or communications;
- providing for periodic independent ethical review, or equivalent challenge, for higher-risk or more novel trials; and,
- where relevant, seeking alignment with emerging sector trust frameworks, including RECCo's Consumer Consent Solution (CCS).

As RECCo's CCS is being developed as a central, standardised and trusted mechanism for managing and validating consumer consent, supported by governance, assurance and technical controls, we would expect sandbox trials to align with those principles where possible, rather than cut across or override them. Participation in the sandbox should not displace underlying trust, privacy, consent or security requirements, and trial participants should remain responsible for compliance with applicable legal and regulatory obligations.

These steps would build on Ofgem's proposed pre- and post-test ethics reviews and help ensure that the sandbox generates evidence not only on technical performance, but on the conditions needed for safe and trusted deployment in the sector.

Q7. Do you have suggestions for how Ofgem can best support stakeholders throughout the pilot and beyond?

We believe that Ofgem can best support stakeholders by providing clear guidance, practical support and opportunities to share learning. In particular, it could publish simple guidance and templates covering application requirements, ethical and technical issues, governance expectations, and how proposals will be assessed. A dedicated contact point or helpdesk would also help applicants resolve technical or regulatory questions quickly, especially where use cases involve shared services, data access or code-governed processes. Finally, regular workshops or forums should be used to share lessons learned, common challenges and emerging good practice, so that learning from the pilot can benefit the wider sector.

Q8. Do you have any other comments, suggestions, or concerns regarding the proposed pilot, the consultation process, or the expected outcomes? Please provide evidence, examples, or reasoning to support your responses wherever possible.

We support the proposed pilot. To maximise its value, Ofgem should ensure there is a clear route from trial evidence to implementation, so successful projects can progress into service change, code change, regulatory guidance or, where relevant, the Future Regulation Sandbox. This is particularly important where AI and automation may interact with shared market services and code-governed processes, as those use cases are likely to require agreed controls around machine-to-machine access, monitoring, capacity management and incident response.

We would also support baked-in progress reporting and structured publication of lessons learned, robust evaluation metrics covering both technical performance and consumer outcomes, and explicit consideration of the longer-term regulatory implications of successful use cases. The pilot should be designed so that smaller market participants are able to engage, whether directly or through partnerships, to ensure the diversity of innovation is not lost. These steps would be consistent with the consultation's focus on published learning, clear evaluation, wider participation and the use of pilot evidence to inform future regulatory development.