From: Kevin McMorrow
To: Pura.ExecutiveSecretary
Subject: Public Comment, Docket No. 26-10-11
Date: Sunday, November 1, 2026 10:43:56 AM

EXTERNAL EMAIL: This email originated from outside of the organization. Do not click any links or open any attachments unless you trust the sender and know the content is safe.

Public Comment, Docket No. 26-10-11
Petition of the Office of the Attorney General and Office of Consumer Counsel for Review of AI, Data Security, and Cybersecurity

Submitted by: Kevin McMorrow, President, Artificial Intelligence Engineering LLC
Home address: 1615 Marine Parkway, Brooklyn, NY 11234
Email: kmcmorrow@aiengineering.io

I support the petition's request that the Authority open a review and require each utility to report the full scope of its AI systems, their objectives, the level of human oversight, and the policies that govern them. I offer these comments as an engineer who has worked on the operating side of a regulated utility. From 2016 to 2021 I was ICS/OT lead at National Grid, where I helped deliver a transmission-pressure SCADA upgrade for gas service to customers in New York, Massachusetts and Rhode Island and led IT-OT convergence across more than 1,000 PLCs and 700 field locations. I have since led IEC 62443 and NIST CSF assessments, served as a Department of Energy red team lead at Idaho National Laboratory, and written a book, Elements of AI Governance (2026), on governing AI in critical infrastructure.

My central recommendation is that the Authority ask for evidence, not only policy statements. A utility can truthfully say "our AI is not connected to SCADA" or "a human reviews every output." The useful question is how each statement is enforced and how the utility would know if it stopped being true. The following reporting structure would let Staff read every utility's filing the same way, whatever its size.

1. A complete inventory, including vendor AI. Each AI use case should be listed with its business owner, whether it is built in-house or supplied by a vendor, what data it uses, where that data is processed (vendor cloud, private cloud or on premises), and whether it is in pilot or production. The inventory should expressly include AI embedded in vendor products, such as outage management, meter data analytics, work management and customer service platforms, since much utility AI arrives through software updates rather than a deployment decision. This answers the petition's third-party vendor and Advanced Metering Infrastructure concerns directly.

2. Consequence tiers. Not every AI tool needs the same scrutiny. I suggest three tiers. Tier A covers office tools that draft, search or summarize with no customer or system effect. Tier B covers decision support, where AI recommends and a person decides on customer, cost, storm staffing or asset matters, such as load forecasts or outage predictions. Tier C covers any AI whose output feeds grid, pipeline, plant or water operations, or that can act without a person approving each action. Tier A needs data rules. Tier B needs the full set of controls below. Tier C should require independent verification before the utility relies on it.

3. A declared operating envelope for each Tier B and C system. The utility should state what the system may do, what it may never do, and how that boundary is enforced technically. For anything near operations, enforcement should be described in IEC 62443 terms: which zone the system sits in, which conduits it uses, and whether data flows into the control network at all. A boundary enforced by network architecture is stronger than one enforced by a written policy.

4. Defined human sign-off points. "Human in the loop" should be replaced by a list of the specific decisions that require a named role to approve, what information that person sees, whether they can override, and whether they have time to do so in practice, for example during storm restoration.

5. Logging and retention. Inputs, outputs, model versions and the human decisions taken on them should be recorded, retained for a stated period, and available to Staff on request. Without logs, no one can reconstruct why a forecast missed or why crews were staged where they were.

6. Baselines and drift monitoring. Many AI systems do not give the same answer twice, so they should be judged on their behavior over time rather than on single outputs. Each utility should report performance before go-live, what it measures afterward (accuracy, drift, false alarms), how often, and what threshold triggers review or rollback. The petition's concern about AI weather and storm forecasting is exactly this question: how a forecast model's accuracy is measured against actual outcomes.

7. Fallback and incident reporting. Each Tier B and C system should have a documented manual fallback and an incident process, and the Authority should consider a duty to notify it when an AI system fails, drifts outside its envelope, or is implicated in a data or cybersecurity incident.

Standards mapping. The Authority need not invent a framework. These elements map onto NIST AI RMF 1.0 (Govern, Map, Measure, Manage), ISO/IEC 42001:2023 for an organization-wide AI management system, and IEC 62443 for the operational boundary. Asking each utility to state which framework it follows, and to map its filing to it, makes filings comparable and lets Staff see gaps quickly. A missing answer need not be a finding against a utility; it identifies the next information request.

Smaller water and gas companies could answer the inventory, accountability and data questions in a paragraph each, with the fuller requirements applying only where they operate Tier B or C systems.

I would be glad to answer any questions from Staff on these recommendations.

Respectfully submitted,
Kevin McMorrow
President, Artificial Intelligence Engineering LLC