The Grid Intelligence Bet

Duke Energy's AI is reading transformer telemetry across 300,000 miles of grid. GE Vernova holds the training data. Your NERC CIP posture doesn't cover it yet.

01

Background

Why utilities are deploying AI for grid monitoring â€” and what the governance architecture looks like

01 / BACKGROUND

The AI Grid Monitoring Market

- Traditional maintenance: fixed inspection cycles â€” miss failures that develop between intervals
- SCADA sensor networks generate millions of data points per substation per day â€” beyond human analysis capacity
- ML systems that process telemetry and surface anomalies are the natural extension of existing sensor investment
- Second driver: outage prediction and storm response â€” SAIDI/SAIFI metrics create direct regulatory incentive
- GE Vernova Grid Solutions serves 900+ utilities globally; ABB Ability covers European and U.S. transmission

01 / BACKGROUND

The Reference Deployments

- Duke Energy: GE Vernova Grid Solutions across transmission operations â€” transformer failure prediction from oil analysis + thermal imaging + load history
- National Grid: ABB Ability ASSET Suite for predictive maintenance across UK and U.S. transmission
- Xcel Energy: AI-assisted wildfire risk modeling in Colorado and Minnesota â€” line segment ignition risk scoring
- Vendor claims: 20â€“30% reduction in unplanned outages in mature deployments
- Performance in broader distribution network applications is less consistently documented than transmission

01 / BACKGROUND

The NERC CIP Compliance Gap

- CIP standards require protection of BES Cyber Systems and Electronic Security Perimeters â€” written before grid AI existed at scale
- AI platforms ingesting SCADA telemetry and feeding control room dashboards operate inside the ESP
- CIP-013 supply chain risk management requirements apply to grid AI vendors â€” most SaaS platforms not initially structured to satisfy
- 73% of utility CIOs cite NERC CIP as primary barrier to accelerating OT AI deployment (Utility Dive/EPRI)
- Several large utilities have delayed planned AI deployments 12â€“18 months to resolve compliance architecture first

02

Decision Required

The deployment is underway or in active procurement. The question is governance architecture.

Who owns the OT data your AI system requires? Without a data portability clause, your switching leverage at renewal is a function of the vendor's accumulated training data â€” not the contract.

Has your compliance team made an explicit NERC CIP categorization decision for the AI platform? "We haven't categorized it yet" is not a defensible answer for a deployment in production for two years.

What is the liability allocation when an operator acts on an AI recommendation that precedes an outage event? Most grid AI vendor contracts disclaim liability for operator decisions â€” the regulatory exposure stays with the utility.

03

Four Options

Option A

Expand current deployment â€” extend to additional asset classes without renegotiating contracts

Captures operational benefit at scale but compounds data ownership and compliance architecture gaps already present

Option B

Recommended

Renegotiate contract before expansion â€” add data portability, CIP compliance, and liability clauses

Adds 60â€“90 days to deployment timeline; resolves governance gaps before scale makes them harder to unwind

Option C

Implement AI recommendation tracking â€” build documented human oversight workflow before extending AI authority

No vendor renegotiation required; instruments the decision chain that regulatory proceedings examine

Option D

Pause expansion pending NERC CIP compliance documentation â€” resolve categorization before extending OT integration

Conservative; adds 4â€“6 months; appropriate for utilities with pending NERC audits or undocumented AI integration

04

Recommendation

Renegotiate the vendor contract before authorizing the next deployment phase â€” data portability clause is the most financially consequential item

Require: specific data format, defined export timeline at contract termination, feature encoding documentation. This costs nothing during normal operations and is worth significant leverage at renewal.

Resolve NERC CIP compliance architecture before expanding OT integration. Make the categorization decision explicitly â€” BES Cyber System, ESP Electronic Access Point, or isolated data diode architecture â€” and document it.

Build the human oversight documentation layer: for every AI recommendation category operators act on, document the verification protocol and how outcomes are recorded against the recommendation.

For vendor evaluation: require production performance data segmented by operating condition â€” normal load, storm events, asset age classes, post-maintenance periods. Aggregate accuracy obscures failure mode distribution.

05

Key Risks

1.

NERC CIP compliance exposure from undocumented AI integration into OT environments

AI platforms with SCADA telemetry access may qualify as BES Cyber Systems. Utilities operating without explicit categorization decisions face audit exposure and mandatory remediation.

2.

Liability gap when operators act on AI recommendations preceding outage events

Grid AI vendor contracts disclaim liability for operator decisions. State PUC proceedings examine AI-assisted decision support in outage root cause analysis. The liability stays with the utility.

3.

Data lock-in from AI vendor training data accumulation

Models trained on your OT data accumulate switching costs over time. Most contracts lack data portability provisions â€” vendor leverage increases at renewal as training data depth grows.

4.

Alert fatigue from false positive maintenance flags eroding operator trust

A system that is 85% accurate on average but generates false positive clusters on specific asset classes will erode operator trust disproportionately in those classes â€” including for genuine at-risk alerts.

5.

Model degradation from grid modernization â€” training distribution drift as DER penetration changes load patterns

AI trained on centralized-generation grid patterns misreads condition signals as solar, storage, and EV charging change load distribution. No retraining schedule tied to modernization milestones means silent accuracy degradation.

06

Questions to Answer Before Next Deployment Phase

1. Does your grid AI vendor contract include explicit data portability provisions â€” format, timeline, and model documentation at contract termination?
2. Has your compliance team made an explicit NERC CIP categorization decision for the AI platform â€” BES Cyber System, ESP access point, or isolated architecture â€” and is it documented for audit?
3. When operators override AI maintenance or switching recommendations, are decisions tracked and outcomes measured? Does your incident documentation capture the AI recommendation before an outage event?
4. Has the vendor provided disaggregated performance data by operating condition â€” storm events, aging assets, post-maintenance periods â€ not just aggregate accuracy?
5. Does your AI retraining schedule include triggers tied to DER penetration thresholds or grid topology changes that would shift the training distribution?
6. If your grid AI platform became unavailable during a major storm response, what is your fallback protocol â€ and has it been tested in a tabletop exercise in the past 12 months?

AI INSIGHT LAB

The Grid Intelligence Bet

The operational case for AI grid monitoring is proven. The governance architecture â€” data ownership, NERC CIP compliance, human oversight documentation â€” is where most utilities are exposed. Build the governance layer before the deployment scale makes it harder.

Read the full memo at aiinsightlab.cloud/memos/duke-energy-grid-ai