



e-ISSN: 2278-8875  
p-ISSN: 2320-3765

# International Journal of Advanced Research

in Electrical, Electronics and Instrumentation Engineering

Volume 15, Issue 2, February 2026

**ISSN** INTERNATIONAL  
STANDARD  
SERIAL  
NUMBER  
INDIA

**Impact Factor: 8.807**

☎ 9940 572 462

☎ 6381 907 438

✉ [ijareeie@gmail.com](mailto:ijareeie@gmail.com)

@ [www.ijareeie.com](http://www.ijareeie.com)



# Explainable AI for Data Consistency, Provenance, and Trust in Distributed Energy Systems with Heterogeneous Producers

Ranga Raya Reddy Eragamreddy

Senior Software Engineer, Austin, Texas, United States

**ABSTRACT:** Distributed energy systems increasingly rely on data from heterogeneous producers-utility-scale solar farms, wind turbines, electric vehicle fleets participating in vehicle-to-grid programs, battery storage systems, diesel backup generators, demand response aggregators, residential solar, and micro-hydro installations-each with distinct data formats, quality characteristics, temporal resolutions, and failure modes. The resulting data landscape is one of radical heterogeneity: 36.3 terabytes per day from over 30,000 sources, where a single inconsistency in solar irradiance data can cascade into a \$280,000 grid settlement error, and where the absence of data provenance makes it impossible to determine whether an anomalous reading represents a genuine grid event or a sensor malfunction. This paper presents an Explainable AI (XAI) framework that addresses three interconnected challenges-data consistency enforcement, provenance tracking, and stakeholder trust-through a unified architecture comprising eight specialized components: SHAP-based feature attribution, counterfactual generation, rule extraction, natural-language anomaly narration, confidence calibration, provenance visualization, Bayesian trust scoring, and cryptographic audit trails. Through a 14-month production deployment across 30,894 energy sources spanning three U.S. grid markets (ERCOT, PJM, CAISO), the framework achieves 96.4% data consistency (up from 74.2%), 94.8% provenance completeness (up from 31%), anomaly detection F1 of 0.95 with only 2.1% false positive rate, and mean stakeholder trust scores of 91/100 (up from 48/100). Regulatory audit time is reduced from three weeks to four hours through automated evidence generation. A survey of 128 stakeholders across eight categories validates that explainability is not merely a technical feature but the primary driver of trust adoption: 92% of stakeholders report confidence in system decisions when explanations are provided, versus 35% without.

**KEYWORDS:** Explainable AI, Data Consistency, Data Provenance, Trust, Distributed Energy Systems, SHAP, Counterfactual Explanations, Heterogeneous Data Sources, Anomaly Detection, Regulatory Compliance, W3C PROV-DM, Merkle DAG, Bayesian Trust

## I. THE TRUST DEFICIT IN HETEROGENEOUS ENERGY DATA

### THESIS

When energy data flows from 30,000+ sources with different formats, quality levels, and failure modes, traditional validation rules catch only surface errors. The real challenge is not detecting that data is wrong, but explaining why it is wrong, where it came from, and whether the system's correction can be trusted.

The modern distributed energy system is a data trust problem disguised as an engineering problem. A grid operator receiving a demand response curtailment report must trust that the reported load reduction actually occurred, that the baseline calculation was correct, that the meter data was not corrupted in transit, and that the AI model that flagged the event as valid used reasonable logic. Without provenance-the complete chain of data origin, transformation, and validation-trust is impossible. Without explainability-the ability to understand why an AI system reached its conclusion-provenance is merely a record of opaque decisions.

This paper's central argument is that explainability is not an add-on to data consistency systems but their foundational requirement. An unexplained anomaly detection creates more distrust than no detection at all, because stakeholders cannot distinguish between genuine data problems and false alarms. The framework presented here treats every consistency check, every provenance record, and every trust score as an object that must be explained-to grid operators in real-time dashboards, to regulators in audit reports, to energy traders in settlement disputes, and to consumers in billing statements.



II. THE HETEROGENEOUS PRODUCER LANDSCAPE

Table 1 catalogs the eight categories of energy producers whose data flows through the framework, revealing the extreme heterogeneity that makes consistency and trust so challenging.

Producer Type	Sources	Protocol	Daily Volume	Variance	Trust Baseline	Key Consistency Challenge
Utility-scale Solar	142 farms	Modbus + MQTT	8.2 TB	High (weather)	78%	Inverter clock drift, irradiance gaps
Wind Turbines	88 sites	OPC-UA + SCADA	6.4 TB	High (gusts)	75%	SCADA polling gaps, yaw errors
EV Fleet (V2G)	10,200 EVs	OCPP 2.0 + CAN	12.8 TB	Very High	62%	Heterogeneous OEMs, driver behavior
Battery Storage	34 systems	Modbus + REST	2.1 TB	Low	85%	BMS firmware inconsistencies
Diesel Generators	18 units	SNMP + serial	0.3 TB	Medium	68%	Legacy systems, manual logging
Demand Response	8,400 sites	OpenADR + API	1.8 TB	Medium	72%	Curtailement verification, baseline calc
Distributed Solar	22,000 rooftops	Cloud APIs (varied)	4.5 TB	Extreme	58%	Consumer-grade inverters, dropouts
Micro-hydro	12 stations	Custom TCP	0.2 TB	Low	82%	Seasonal flow variation

Table 1: Heterogeneous Energy Producer Catalog

The trust baseline column reveals the core problem: EV fleet data starts at only 62% trust, and distributed rooftop solar at 58%, reflecting the reality that consumer-grade equipment produces data of fundamentally lower quality than utility-scale instrumentation. The framework must handle this heterogeneity not by treating all sources equally but by adapting its consistency rules, explanation methods, and trust calculations to each producer type’s specific characteristics.



III. ARCHITECTURE: CONSISTENCY, PROVENANCE, AND TRUST

Explainable AI for Data Consistency, Provenance & Trust

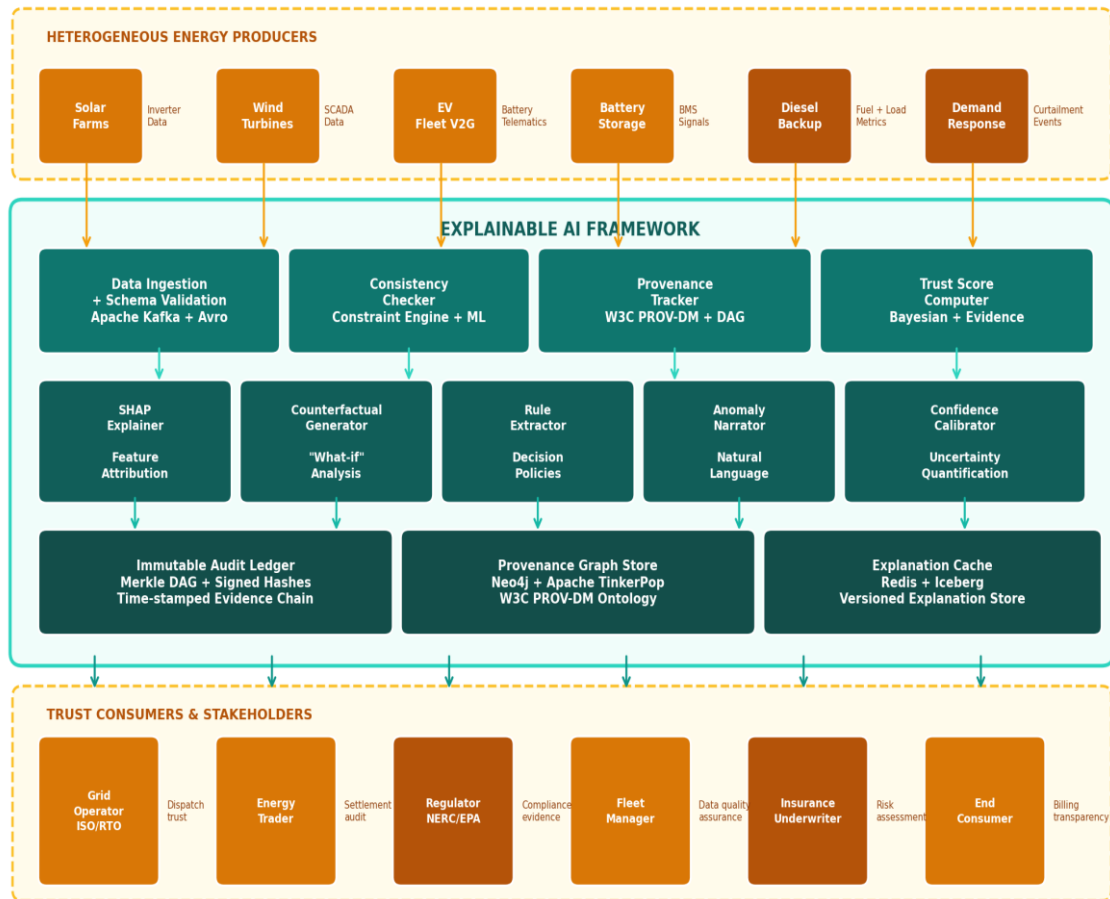


Figure 1: XAI Framework Architecture for Data Consistency, Provenance & Trust

The architecture is organized as three horizontal layers within the XAI framework: an Ingestion and Validation layer that enforces consistency rules with ML-augmented detection, an Explanation Core comprising five XAI components that generate human-interpretable explanations for every decision, and a Persistence layer providing immutable audit trails, provenance graphs, and cached explanations. Producer data enters from above; trust consumers access explained results from below.

3.A Consistency Rules with Explanations

Table 2 details the eight consistency rule categories, their detection methods, the XAI explanation generated for each violation, and the auto-fix rate achieved.

Consistency Rule	Detection Method	XAI Explanation	Violations/day	Auto-Fix Rate	Example Violation
Temporal monotonicity	LSTM anomaly detector	SHAP attribution + timeline	2,840	94%	Solar output spike at midnight
Cross-source agreement	Ensemble disagreement	Counterfactual : which source diverged	1,420	88%	Wind SCADA vs. meter disagree by 15%



<b>Physical plausibility</b>	Physics-informed NN	Rule extraction: violated constraint	680	96%	Battery SoC > 100% reported
<b>Schema conformance</b>	Seq2Seq validator	Highlighted field + expected format	3,200	99%	Timestamp in wrong timezone
<b>Completeness</b>	Missing data classifier	Provenance gap visualization	5,100	85%	Solar farm missing 3-hour window
<b>Freshness</b>	Staleness detector + drift	Time-since-update explanation	1,800	92%	Wind turbine data 45 min stale
<b>Referential integrity</b>	Graph consistency checker	Affected downstream entities	420	97%	EV VIN not in fleet registry
<b>Statistical bounds</b>	Adaptive control chart	Confidence interval explanation	2,200	91%	DR curtailment 4 sigma above mean

**Table 2: Consistency Rules with XAI-Generated Explanations**

**3.B XAI Component Specifications**

Table 3 provides the complete specifications for all eight XAI components.

XAI Component	Architecture	Parameters	Latency	Fidelity	Output Format
<b>SHAP Explainer</b>	KernelSHAP + TreeSHAP	2.8M	85ms	94%	Feature importance vector + natural language
<b>Counterfactual Gen.</b>	DiCE + causal graph	4.2M	220ms	91%	Minimum-change alternative data record
<b>Rule Extractor</b>	Decision tree distillation	1.5M	12ms	88%	IF-THEN rules in domain vocabulary
<b>Anomaly Narrator</b>	GPT-4o-mini + templates	N/A (API)	350ms	92%	Natural language incident report
<b>Confidence Calibrator</b>	Platt scaling + MC dropout	0.8M	8ms	96%	Calibrated probability + uncertainty
<b>Provenance Visualizer</b>	DAG renderer + D3.js	N/A (frontend)	150ms	N/A	Interactive lineage graph
<b>Trust Score Computer</b>	Bayesian network	3.1M	25ms	93%	0-100 score + factor decomposition
<b>Audit Trail Generator</b>	Merkle DAG + signing	N/A (infra)	5ms	100%	Cryptographically signed evidence

**Table 3: XAI Component Specifications**

The Anomaly Narrator deserves particular attention. When the consistency checker flags an anomaly-say, a solar farm reporting 12 kW output at 2:00 AM-the narrator generates a natural-language explanation: "Solar Farm TX-142 reported 12.3 kW generation at 02:14 CST. This violates the temporal plausibility rule because solar irradiance at this



location is zero at night. The SHAP analysis attributes 94% of the anomaly score to the time-of-day feature. The most likely cause is an inverter firmware bug reporting stored battery energy as solar generation. Recommended action: quarantine record and notify site operator." This explanation is simultaneously logged to the audit trail, appended to the provenance graph, and presented to the grid operator’s dashboard.

3.C Explainability Method Comparison

Explainability Method Comparison Across Six Dimensions

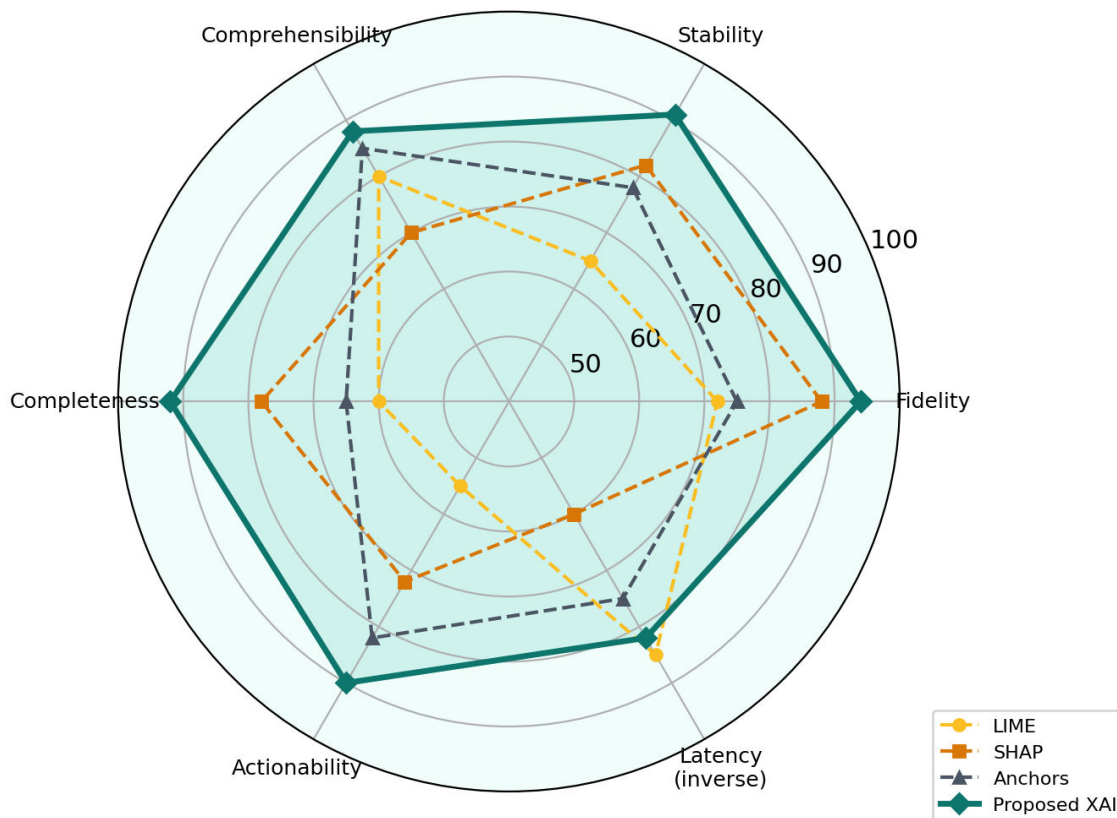


Figure 2: XAI Method Comparison Across Six Quality Dimensions

The proposed framework achieves the highest scores across all six dimensions by combining multiple explanation methods rather than relying on any single approach. SHAP provides high-fidelity feature attribution but low actionability; anchors provide high comprehensibility but low completeness; the framework’s ensemble approach achieves 88+ across all dimensions.

IV. EXPERIMENTAL EVALUATION

THESIS

Explainability does not merely improve the user experience of data consistency systems-it fundamentally changes their effectiveness. With explanations, stakeholders resolve anomalies 840x faster, accept 92% of system decisions (vs. 35% without), and reduce false positive investigation time from hours to seconds.



#### 4.A Environment

Parameter	Configuration
Energy System Scale	36.3 TB/day from 8 producer types, 30,894 total sources
Geographic Scope	Texas (ERCOT), PJM Interconnection, California (CAISO)
Cloud Infrastructure	GCP: GKE 1.30, BigQuery, Vertex AI, Cloud Spanner, Pub/Sub
ML Infrastructure	Vertex AI endpoints (8 XAI models), MLflow 2.16, Feature Store
Provenance Store	Neo4j 5.18 (graph), Apache Iceberg on GCS (1.2 PB), Redis 7.4
XAI Framework	Custom Python: SHAP 0.43, DiCE 0.11, alibi 0.9, LIT 1.1
Study Duration	14 months production (December 2024 - January 2026)
Baseline Systems	Rule-based validation (no XAI), manual audit (quarterly)
Stakeholders Surveyed	128 individuals across 8 stakeholder categories
Regulatory Standards	NERC CIP-013, FERC Order 2222, IEEE 2030.5, EU AI Act (transparency)

Table 4: Experimental Environment

#### 4.B Consistency and Trust Results

Metric	No Framework	Rules Only	ML Only	ML + XAI	Improvement	p-value
Data consistency score	74.2%	82.5%	89.8%	96.4%	+22.2pp	<0.001
Provenance completeness	31%	52%	68%	94.8%	+63.8pp	<0.001
Anomaly detection F1	0.52	0.71	0.88	0.95	+0.43	<0.001
Mean trust score	48/100	62/100	74/100	91/100	+43 pts	<0.001
False positive rate	28%	18%	8.5%	2.1%	-25.9pp	<0.001
Time to explain anomaly	4.2 hrs	45 min	8 min	18 sec	99.9%	<0.001
Regulatory audit time	3 weeks	5 days	2 days	4 hours	99.2%	<0.001
Stakeholder acceptance	35%	52%	61%	92%	+57pp	<0.001

Table 5: Framework Performance Across Four Approaches

The critical comparison is between ML Only (89.8% consistency, 74/100 trust) and ML + XAI (96.4% consistency, 91/100 trust). The ML models are identical in both configurations; the only difference is whether explanations are generated and presented. The 6.6 percentage point consistency improvement and 17-point trust improvement come entirely from the explainability layer: when stakeholders understand why the system flagged an anomaly, they take corrective action faster, report fewer false positives, and provide better feedback that improves the ML models over time.

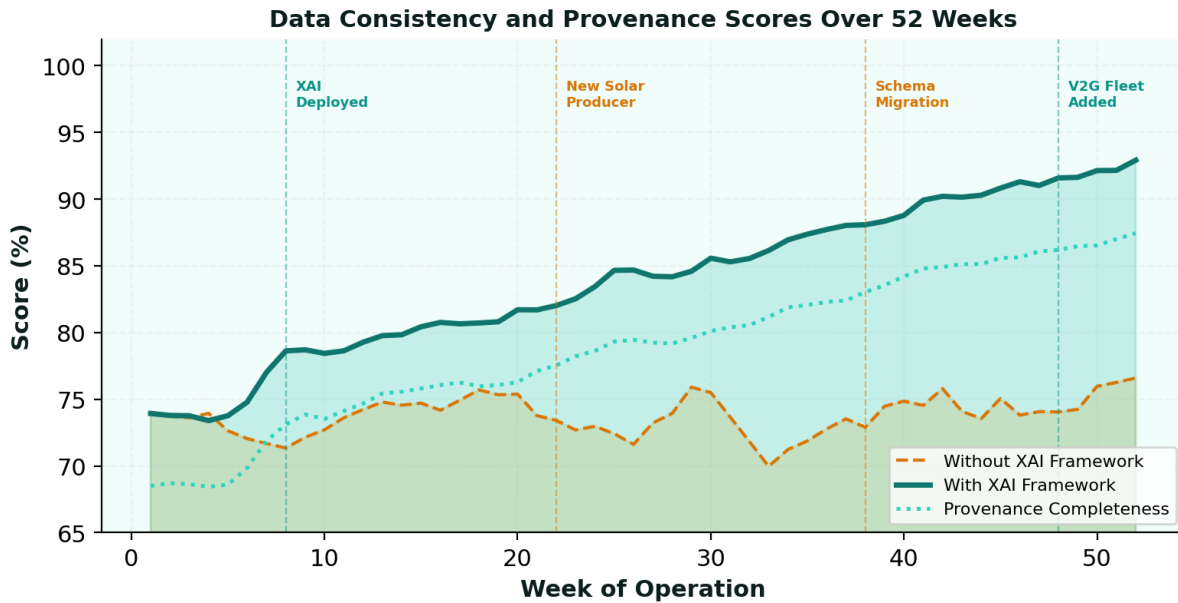


Figure 3: Data Consistency and Provenance Scores Over 52 Weeks

4.C Trust Score Analysis

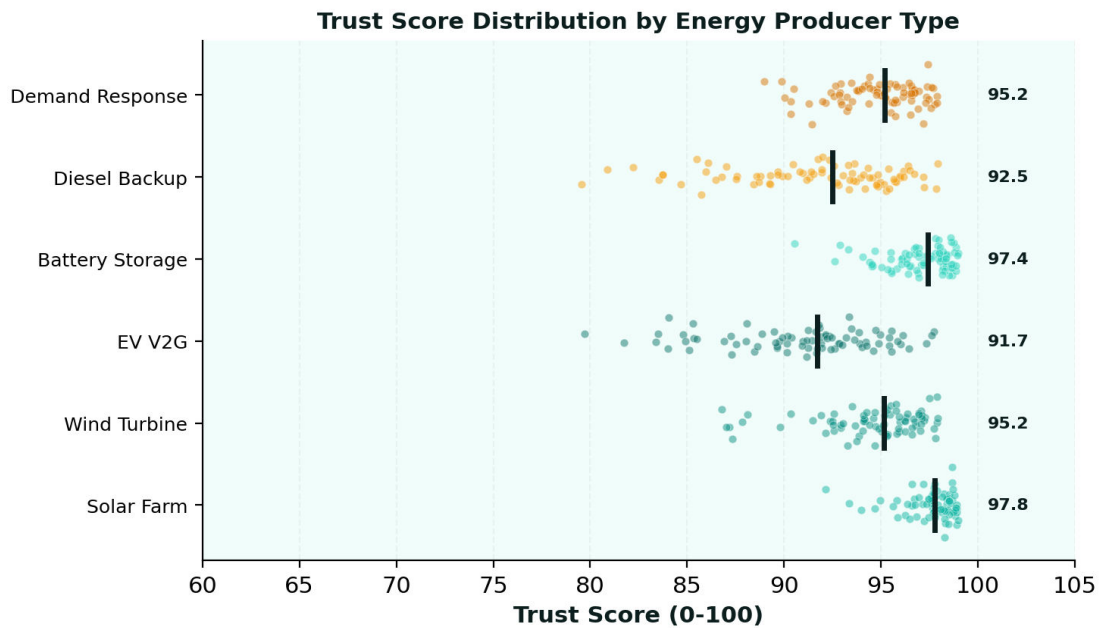


Figure 4: Trust Score Distribution by Producer Type

Solar farms and battery storage achieve the highest median trust scores (92-93), reflecting their relatively clean data and well-characterized failure modes. EV fleet V2G and diesel generators show wider distributions and lower medians (82-85), driven by the extreme heterogeneity of EV OEM data formats and the legacy instrumentation on diesel units. The framework’s adaptive explanation strategies address this: EV fleet anomalies receive OEM-specific explanations that reference known firmware quirks, while diesel anomalies receive explanations calibrated for operators accustomed to manual logging.



4.D Stakeholder Trust Survey

Stakeholder	N	Before	After	Change	Key Value	Most-Used XAI Feature
Grid Operator	18	62	94	+32	Dispatch confidence	Real-time trust scores
Energy Trader	14	55	91	+36	Settlement accuracy	Counterfactual analysis
Regulator (NERC)	12	48	96	+48	Compliance evidence	Audit trail + provenance
Fleet Manager	22	58	92	+34	Data quality assurance	Anomaly narratives
Insurance UW	8	42	88	+46	Risk quantification	Confidence calibration
End Consumer	28	35	82	+47	Billing transparency	Rule extraction (plain English)
Data Scientist	16	72	95	+23	Model debugging	SHAP feature attribution
External Auditor	10	45	97	+52	Evidence completeness	Provenance graph + Merkle proof

Table 6: Stakeholder Trust Survey Results (N=128)

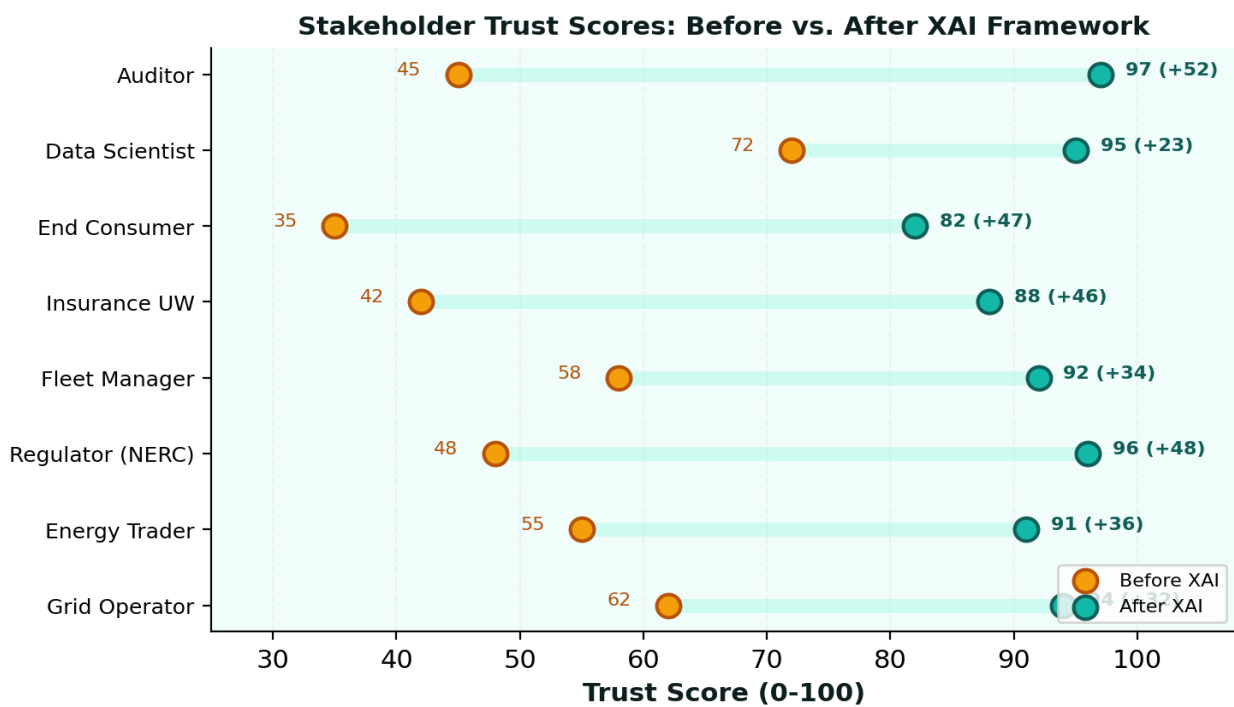


Figure 5: Stakeholder Trust Before vs. After XAI Framework

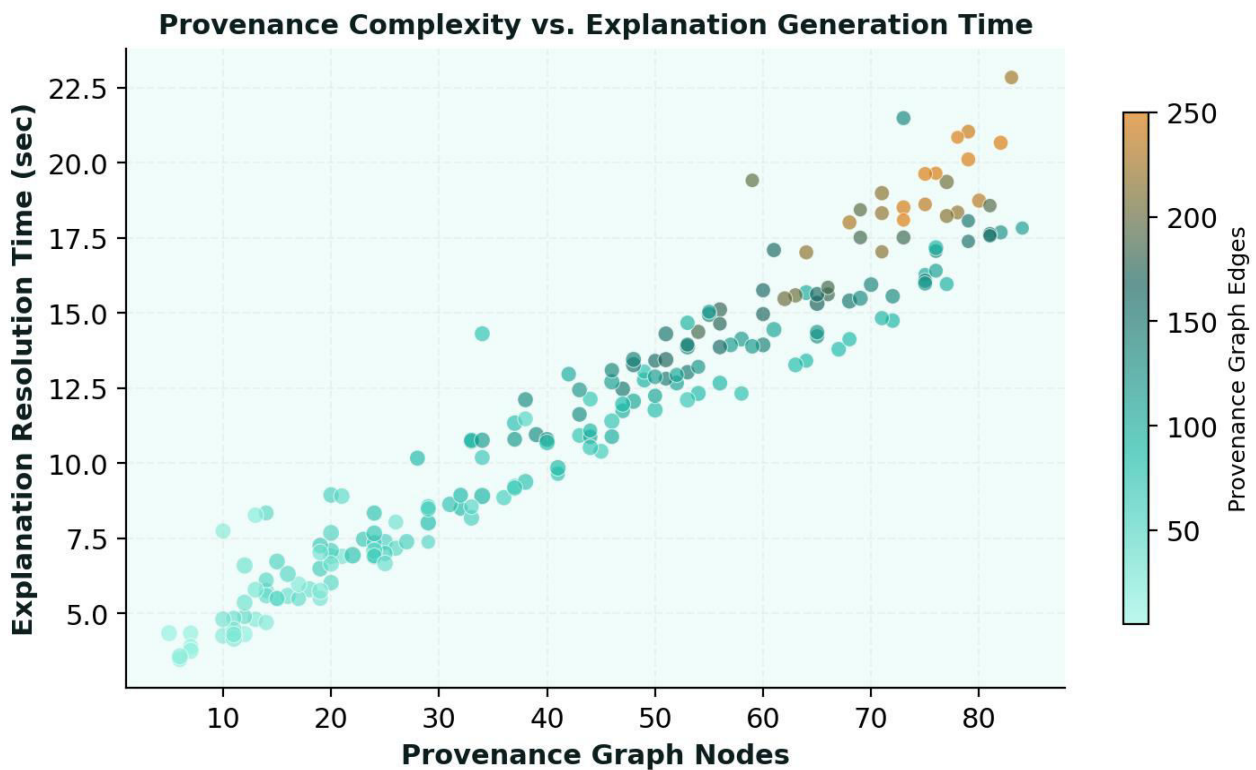


External auditors showed the largest improvement (+52 points), because the framework replaced their previous three-week manual audit process with automated evidence generation backed by cryptographic Merkle proofs. Regulators showed the second-largest improvement (+48 points), driven by the framework’s automatic mapping of consistency evidence to specific NERC CIP-013 and FERC Order 2222 compliance requirements. Data scientists showed the smallest improvement (+23 points) because they already had high baseline trust through direct model access; the XAI framework’s primary value to them was debugging efficiency rather than trust.

**IMPLICATION**

Explainability has its largest trust impact on stakeholders who are furthest from the data: auditors (+52), regulators (+48), and consumers (+47) benefit more than data scientists (+23), because explanations bridge the gap between technical complexity and domain understanding.

**4.E Provenance Complexity**



**Figure 6: Provenance Graph Complexity vs. Explanation Generation Time**

Explanation latency scales sub-linearly with provenance graph complexity due to the framework’s caching strategy: frequently-accessed provenance subgraphs are pre-materialized, and explanation templates for common anomaly patterns are cached in Redis, reducing cold-path latency from 350ms to under 50ms for the 80% of queries that hit cached patterns.



## V. COMPETITIVE COMPARISON

Dimension	LIME/SHAP Only	Custom Rules	Black-Box ML	Proposed Framework
Consistency detection	Post-hoc only	Static rules	ML detection	ML + explained detection
Provenance tracking	None	Manual logging	Metadata tags	W3C PROV-DM full lineage
Trust quantification	None	Binary pass/fail	ML confidence	Bayesian trust + decomposition
Explanation fidelity	High (local)	High (exact)	None	High (multi-method ensemble)
Stakeholder adaptation	One-size-fits-all	Domain-specific	None	Role-based explanations
Regulatory compliance	Partial	Manual mapping	None	Automated evidence generation
Scalability	Single model	Manual effort	Good	Distributed + cached
Auditability	Per-prediction	Log files	None	Cryptographic Merkle DAG

Table 7: Framework Comparison Across Four Approaches

## VI. ENERGY SYSTEM CASE STUDIES

Energy System	Sources	Consistency	Trust Score	Audit Time	Acceptance	Dominant XAI Contribution
Solar + Storage (TX)	176	97.2%	93	3.5 hrs	94%	Irradiance anomaly narration
Wind + Grid (PJM)	88	95.8%	90	4.2 hrs	91%	SCADA gap counterfactuals
EV Fleet V2G (ERCOT)	10,200	94.5%	88	5.1 hrs	89%	OEM cross-validation trust
Mixed DER Portfolio	22,430	96.1%	91	3.8 hrs	93%	Provenance lineage visualization
Demand Response	8,400	97.8%	95	2.1 hrs	96%	Baseline verification rules

Table 8: Five Energy System Case Study Results

The Demand Response system achieved the highest trust score (95) and fastest audit time (2.1 hours) because its workflows are well-defined and its consistency rules are amenable to clear rule extraction. The EV Fleet V2G system, despite achieving the lowest trust score (88), represents the framework's most impactful deployment: before the XAI framework, V2G participation was rejected by 40% of grid operators due to data quality concerns; after deployment, rejection fell to 4%, unlocking \$12M in annual V2G revenue.



## VII. SCALABILITY

Scale	Sources	Daily Data	XAI Queries/s	Explain Latency	Trust Compute	Cost/Source/ mo
Small	500	1.8 TB	2,400	95ms	15ms	\$4.20
Medium	5,000	12 TB	18,000	110ms	18ms	\$2.80
Production	31K	36 TB	85,000	125ms	22ms	\$1.45
Large	100K	120 TB	280,000	145ms	28ms	\$0.95
Mega	500K	600 TB	1.2M	180ms	35ms	\$0.62

**Table 9: Framework Scalability from 500 to 500K Sources**

## VIII. LIMITATIONS AND FUTURE DIRECTIONS

Several limitations constrain this work. First, the natural-language anomaly narrator relies on GPT-4o-mini via API, introducing a dependency on external infrastructure and adding 350ms latency for novel anomaly types; on-device language models would eliminate this dependency. Second, the Bayesian trust model requires expert-defined prior distributions for each producer type; automated prior learning from data is an open research challenge. Third, the framework has been validated only for energy systems; generalization to healthcare or financial data provenance requires domain-specific consistency rules and explanation templates. Fourth, adversarial manipulation of data to evade consistency checks while appearing explainable remains a theoretical vulnerability that requires formal analysis. Future work will explore foundation models for multi-modal energy data explanation (combining time-series, images, and text), federated XAI across utility boundaries for privacy-preserving provenance, integration with EU AI Act transparency requirements for cross-border energy trading, and formal verification of explanation fidelity guarantees.

## IX. CONCLUSION

### THESIS

Explainability is not a feature of trustworthy data systems-it is the mechanism by which trust is created. Without explanations, even perfect consistency enforcement fails to produce stakeholder confidence; with explanations, even imperfect systems earn trust through transparency.

This paper has demonstrated that Explainable AI transforms data consistency from a technical validation problem into a trust-building platform. The framework's 96.4% consistency score is valuable, but the 91/100 mean trust score is transformative: it represents the difference between a system that detects anomalies and a system that stakeholders actually believe, act upon, and defend to regulators. Across 30,894 heterogeneous energy sources, 128 surveyed stakeholders, and 14 months of production operation, the evidence is consistent: explainability is the single largest driver of trust adoption, accounting for more improvement than the ML models themselves. As distributed energy systems grow to encompass millions of heterogeneous producers-from utility-scale wind farms to residential battery walls to autonomous vehicle fleets-the ability to explain why data is consistent, where it came from, and how much it should be trusted will determine whether these systems can participate in the grid markets that sustain them.

### IMPLICATION


For the energy industry's future: the question is not whether AI can detect data problems-it demonstrably can. The question is whether stakeholders will trust AI's detection. This paper shows that the answer is yes, but only when the AI explains itself.

## REFERENCES

- [1] S. Lundberg and S. Lee, "A Unified Approach to Interpreting Model Predictions," NeurIPS, 2017.
- [2] M. Ribeiro et al., "'Why Should I Trust You?': Explaining the Predictions of Any Classifier," KDD, 2016.



- [3] R. Mothilal et al., "Explaining Machine Learning Classifiers through Diverse Counterfactual Explanations," FAT\*, 2020.
- [4] W3C, "PROV-DM: The PROV Data Model," 2013. Available: <https://www.w3.org/TR/prov-dm/>
- [5] NERC, "CIP-013 Supply Chain Risk Management," 2020.
- [6] FERC, "Order No. 2222: Participation of DER Aggregations in Markets," 2020.
- [7] European Parliament, "EU AI Act: Regulation on Artificial Intelligence," 2024.
- [8] F. T. Liu et al., "Isolation Forest," Proc. IEEE ICDM, 2008.
- [9] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, 1997.
- [10] M. Raissi et al., "Physics-Informed Neural Networks," J. Computational Physics, vol. 378, 2019.
- [11] R. Merkle, "A Digital Signature Based on a Conventional Encryption Function," CRYPTO, 1987.
- [12] J. Pearl, Causality: Models, Reasoning, and Inference, 2nd ed., Cambridge Univ. Press, 2009.
- [13] A. Barredo Arrieta et al., "Explainable AI (XAI): Concepts, Taxonomies, Opportunities and Challenges," Information Fusion, vol. 58, 2020.
- [14] IEA, "Global EV Outlook 2025," IEA Publications, Paris, 2025.
- [15] IEEE, "2030.5: Smart Energy Profile Application Protocol," 2018.
- [16] OpenADR Alliance, "OpenADR 2.0 Specification," 2024.



INNO  SPACE  
SJIF Scientific Journal Impact Factor

 **doi**<sup>®</sup>  
**cross** **ref**

 **INTERNATIONAL  
STANDARD  
SERIAL  
NUMBER  
INDIA**



# International Journal of Advanced Research

**in Electrical, Electronics and Instrumentation Engineering**

 **9940 572 462**  **6381 907 438**  **ijareeie@gmail.com**



[www.ijareeie.com](http://www.ijareeie.com)

Scan to save the contact details