

Grounded AI and the Path to Measurable Utility Value

EPRI Open Power AI Domain Intelligence Working Group
(also open to the OPAI Implementation Working Group)



Apurba Sakti, Ph.D.
Electric Power Research Institute
asakti@epri.com

Monthly Working Group Meeting

16 Jul 2026

GROUNDING DOES NOT MEAN GUARANTEED



AI SAFETY SHARE

Citations and grounding improve reliability — but important decisions still require verification.



WHY THIS MATTERS



Grounded answers can still be incomplete



Retrieved content may be outdated or misleading



Untrusted content can influence model behavior



1. CHECK THE SOURCE

Is the cited source authoritative for this decision?

Prefer official, current, and context-relevant materials.



2. CHECK THE CONTEXT

Does the cited passage actually support the conclusion?

Look for missing nuance, outdated details, or misinterpretation.



3. CHECK THE CONSEQUENCE

If the output affects safety, compliance, cost, or operations, verify before use.

Use human review and established controls.



TAKEAWAY

Use grounded AI outputs to strengthen judgment — not replace it.



Grounding helps. Verification completes the process.




Sources: NIST AI 600-1 Generative AI Profile; RAGTruth (ACL 2024); OWASP Top 10 for LLM Applications; UK NCSC guidance on prompt injection.

Housekeeping

A few norms to keep today's discussion useful, interactive, and publicly shareable


1



Participate actively

Please mute upon entry. Use raise hand or chat for questions and discussion.


2



Keep it publicly shareable

Please share only material suitable for public release.

3



No automated notetakers

Please do not add automated notetakers to this meeting.

4



Recording notice

This session will be recorded. Participation acknowledges consent.

5



Goal for today

Practical discussion and shared learning. Product pitches should be avoided or kept at a minimum.



Powered by EPRI

OPEN POWER AI CONSORTIUM

Our Work Groups

Domain Intelligence



Enabling domain-aware AI performance for power sector applications

Use Case Sandbox



Identify, prioritize, and evaluate real-world applications

Implementation



Share methods & best practices for AI deployment and scaling



Harnessing AI's transformative potential for a resilient, cost-effective energy future through collaboration

Open-source data and AI models for industry-wide value



Google



ORACLE



Learn More:
www.openpowerai.org



Domain-Specific Model
to include knowledge from

> 15,000

EPRI Reports



250+

use cases identified for consortium prioritization and development

Join us in Shaping the Future of AI and Energy

Today's throughline: from performance evidence to grounded, utility-relevant AI

How evaluation, trusted knowledge, and practical utility adoption connect to more reliable AI performance

1

Benchmarking reveals the gaps

Framing the session

- Identify where general-purpose models perform well – and where they struggle
- Provide repeatable evidence to guide improvements

2

Ground AI in trusted knowledge



Justin Thibault, P.E.
Product Manager, EPRI.ai

Show how retrieval, grounding, and search design connect AI to trusted organizational content and improve answer quality.

3

Turn AI into business outcomes



Josh Murray
AVP of AI, Innovation, and Emerging Technology, Alliant Energy

Highlight practical utility use cases, lessons learned, and how AI initiatives can be scaled around measurable value.

4

Connect the pieces

Working Group Discussion

- Discuss how evaluation, grounding, and real-world adoption fit together
- Explore what enables more reliable and trusted AI for power-sector applications




Takeaway

Evidence shows where the gaps are; grounded systems and practical utility adoption show how to start closing them.


The electric power ecosystem sees value from AI; scaling needs include overcoming data, security, and integration challenges

OPAI Ongoing *Member Representative Committee* Survey Results


Where AI is delivering value



Top barriers to scaling



Where expansion likely in 1-3 years



Number of respondents vary from 131-133 with strong representation from generation, transmission, distribution, and AI solution providers

Implication

To scale responsibly, **assurance evidence** (performance + controls + monitoring) is needed

Scaling AI responsibly requires documented, verifiable evidence

External assurance frameworks converge on demonstrating performance, controls, and operational readiness

What leading frameworks expect

NIST AI RMF

Govern-Map-Measure-Manage

Measure focuses on **TEVV** + monitoring.

Other relevant frameworks

EU AI Act

Demonstrate robustness/accuracy + human oversight

ISO/IEC 42001 & ISO/IEC 23894


Monitoring, measurement, and auditability

IEC 62443 & NERC CIP

Security assurance: verified controls + monitoring

Assurance Evidence Loop

A practical way to organize and continually refresh the evidence



Assurance evidence = demonstrated performance + verified controls + validated operation

Takeaway

Frameworks set expectations; a repeatable evidence loop helps move AI from pilots toward production

WattWorks.EPRI.Com

The Power Sector's AI Benchmarking Hub

Update since last month

- Newer models including Claude Fable 5 benchmarked
- Future version will allow model developers to submit their model for evaluation
- A subset of questions will be made public
- Ongoing efforts at agentic benchmarks to be released in the future

The Power Sector's AI Benchmarking Hub

How Today's Leading LLMs Perform on Power-Sector Questions

EPRI benchmarks leading large language models on **real-world, utility-relevant questions** to help the power sector understand where today's AI systems are strong, **where they fall short**, and how performance evolves over time.

● Live Benchmark — Updated as new models are evaluated

Download White Paper →

Watch Benchmarking Discussion →



Multiple-choice evaluations



Open-ended short answers



Multi-run repeatability



Domain-augmented tools

INTERACTIVE BENCHMARK

Explore how selected models perform across EPRI benchmark datasets, including **short-answer** and **multiple-choice questions**, **model-only** and **web-search modes**, and breakdowns by **difficulty** or **topic**.

DATASET ⓘ

☒ Multiple-Choice Questions (MCQ)

☐ Short-Answer Questions (SAQ)

MODE ⓘ

☒ Search Off (Model Only)

☐ Search On (Web Search)

VIEW ⓘ

☒ Overall

☐ By Difficulty

☐ By Topic

MODELS ⓘ 5/10 selected Select default 10 Clear all

☐ Grok 4

☒ GPT-5

☐ Gemini 2.5 Pro

☐ Claude Sonnet 4.5

☐ Llama 4 Maverick

MCQ — Overall

Weighted Accuracy (%)

Search Off (Model Only)




About these results

Scores reflect difficulty-weighted accuracy across EPRI's multiple-choice benchmark set. Higher scores indicate better performance. Error bars show 95% confidence intervals where available. Search Off uses model knowledge only, without web search.

Use-case benchmarking in action: Knowledge retrieval

Chauncey™ as an example workflow: measuring grounded answers & citation quality for research-backed insights



Domain Intelligence Working group

**Future Discussion Topics – Please
sign up**

<https://forms.office.com/r/AfcAdPDWgX>



Today's throughline: from performance evidence to grounded, utility-relevant AI

How evaluation, trusted knowledge, and practical utility adoption connect to more reliable AI performance



**Takeaway**

Evidence shows where the gaps are; grounded systems and practical utility adoption show how to start closing them.


How EPRI's Chauncey Builds a Response for You

EPRI.ai RAG and Query Orchestration




Justin Thibault
EPRI.ai Product Manager

Typical Retrieval Augmented Generation (RAG) Model




RAG Example Using EPRI's OAI RAG DB on Hugging Face




EPRI Public Reports = ~2/3rds of EPRI's 50 Years of Published Research

Chauncey Query Orchestration





TOGETHER...SHAPING THE FUTURE OF ENERGY®

A high-angle photograph of a wind turbine's nacelle and hub during sunset. Two technicians in safety gear are visible on the nacelle. The background shows a field with several other wind turbines under a golden sky.

Josh Murray

AVP – AI, Innovation, and Emerging Technology

Alliant Energy: At a Glance

Our purpose: Serve customers and build stronger communities



>3 GW of
renewable
resources



1 million electric customers
430,000 gas customers



~3,000 dedicated
employees



28% of distribution
is underground



Top 5 largest regulated
wind owner-operator



Top 5 largest regulated
solar owner-operator



~44% of energy from
renewable resources
in 2024

AI enablement and training

2026 numbers to date

250 People trained in
Copilot Workshops

4 Microsoft Copilot In-Person
Workshops in CR and MSN

11 Alliant Innovation (AI)
Academy Workshops

Bi-weekly schedule will
continue

2 Champion Groups
launched

Travero and Finance
Champion groups launched



3 Agentic AI
Workshops

3 Additional Agent workshops
scheduled for Q3 2026




Generative AI Tools Usage Dashboard

GenAI tools becoming a part of daily activity



Monthly Active Users — Adoption Trend



Approved AI Platforms

- Microsoft Copilot (Builder and Studio)
- Snowflake
- Azure AI Foundry
- AWS SageMaker
- Abacus
- Python

Approved AI Tools:

- Microsoft Copilot
- Microsoft Agents
- AutomationAnywhere Agents
- Snowflake Agents
- Abacus

Our AI Governance Journey

Background:

As we progress and mature we knew we needed; centralized oversight, policies with appropriate fairness controls, and manage evolving third-party risks.

At the same time, we took this as an opportunity to strengthen our Data Governance practice as well.

The AI teams and Cybersecurity are also working together on several initiatives to address cyber and supply chain risks

Objectives:

Establish an AI Management framework and oversight model to ensure accountable decision-making – this involves developing an AI Policy, designing a repeatable AI lifecycle, developing risks and controls and defining technical monitoring and human-review procedures for ongoing control

Outcome:

Alliant Energy now has a clear, enterprise-wide AI Management Guardrails foundation that strengthens oversight, manages risk and enables scalable, responsible AI adoption



AI Management Guardrails Model

Developed Alliant Energy's shared understanding of AI governance roles, decision rights, and cross-functional responsibilities through structured stakeholder workshops

AI Management Framework & Risk Management

Developed Alliant Energy's foundational governance framework built on core pillars and delivered an AI Governance Committee Charter, AI Policy, and Risk and Control Matrix

Asset & Use Case Management

Established Alliant Energy's centralized approach to identifying, tracking, and managing AI assets and use cases across the enterprise

Monitoring & Auditing Design

Developed Alliant Energy's risk-tiered framework for monitoring and auditing AI systems to support ongoing performance, compliance, and risk oversight

Vendor Risk Assessment & Contractual Controls

Implemented AI-specific vendor risk assessment and contractual controls to strengthen third-party governance, accountability, and risk management across the AI lifecycle under the TPRM program with input from the AI Management Committee

Program Roadmap

Outlined a phased AI Management Guardrails roadmap that prioritizes near-term actions, sequences mid-term enhancements, and plans longer-term initiatives to support scalable, enterprise-wide adoption

Agentic AI enabled by cyber transformation

Key Cyber Transformation Initiatives will enable faster AI Adoption

NEED TO HAVE Foundational



DLP

Data Loss Prevention

Prevents sensitive data exfiltration by AI agents



IAM

Identity & Access Management

Controls who and what agents can access



Threat Detection & Response

Detects and contains malicious agent behavior

NICE TO HAVE Strengthening enhancements



PAM

Privileged Access Management

Governs elevated permissions for agents



Data Protection & Encryption

Secures data at rest and in transit

Foundational

Governance

AI Security Governance – Risk Management – Incident Response

Security capabilities mapped to AI risks

How each security capability addresses specific AI risk categories

Capability \ Risk	Data Privacy & Security	Ethical & Bias	Regulatory & Compliance	Operational Risk	Financial Risk	Reputational Risk	Technical Risk
DLP  Data Loss Prevention	✓	✓	✓	✓	✓	✓	
IAM  Identity & Access Mgmt	✓		✓	✓			✓
Threat Detection & Response 	✓		✓	✓	✓	✓	✓
PAM  Privileged Access Mgmt	✓		✓	✓		✓	✓
Data Protection & Encryption 	✓		✓		✓	✓	✓

✓ Capability mitigates this risk ○ Not directly applicable


Wind Failure Analysis

Machine Learning & Analytics

- **Problem Statement:** Wind sites repairs are extremely costly and there is room to optimize both timing and cost of maintenance and repairs
- **Estimated Benefit to Alliant Energy:** Data-driven decision-making for identifying and planning major component replacements at wind sites resulting in \$3 million to \$5 million saved annually
- **Next Steps:** Component failure prediction & proactive component replacement strategy

Wind Failure Analysis

Machine Learning & Analytics



AI usage in call centers



Nice CX1

Alliant Energy procured Nice CX1 for call transcription; the tool can also detect voice recognition and perform sentiment analysis



AI Enhances Agent Support (Future)

AI provides real-time guidance to agents, surfacing relevant knowledge and automating tasks to reduce handle time and improve consistency



Improved Quality Monitoring

AI enables continuous analysis of call interactions to assess customer experience and evaluate agent performance



Data-Driven Insights

AI analyzes call data to identify issues, pain points, and performance trends, enabling training & process optimization and improvements

Data analytics for Energy Markets

Gas Dispatch Weather Monitoring

- Analytic tool delivered to support the EMO team
- Monitors weather forecasts to ensure alignment with forecasted demand
- Integrates weather data and seasonal phase descriptions to generate alerts

IAILMNIWI				Location	
				All	
6/22/2026	6/23/2026	6/24/2026	6/25/2026	6/26/2026	6/27/2026
Summer Normal Conditions	Summer Normal Conditions	Summer Normal Conditions	Summer Normal Conditions	Summer Normal Conditions	Summer Normal Conditions
					
Conditions	Conditions	Conditions	Conditions	Conditions	Conditions
65.2 °F	67.7 °F	68.3 °F	66.8 °F	68.3 °F	72.8 °F
Average	Average	Average	Average	Average	Average
77 °F	78 °F	82 °F	79 °F	79 °F	84 °F
High	High	High	High	High	High
6/28/2026	6/29/2026	6/30/2026	7/1/2026	7/2/2026	7/3/2026
Summer Normal Conditions	Summer Phase 1	Summer Phase 1	Summer Phase 1	Summer Normal Conditions	Summer Normal Conditions
					
Conditions	Conditions	Conditions	Conditions	Conditions	Conditions
78.2 °F	78.1 °F	75.6 °F	75.5 °F	75.3 °F	75.1 °F
Average	Average	Average	Average	Average	Average
89 °F	90 °F	86 °F	87 °F	88 °F	88 °F
High	High	High	High	High	High

AI for utility operations

Powered AI – Microsoft Copilot Studio



Use Case

Copilot Agent for Support Questions

Instant answers for field crews, engineers and operations staff

AI-powered copilot trained on utility knowledge bases – resolves internal support queries in seconds, reducing escalations and improving staff productivity

Faster Resolution

Dramatically reduces time spent searching for procedures and answers

Always Current

Trained on live SOPs, safety codes & rate schedules

Multi-Audience

Field crews, engineers, dispatchers, back-office staff, IT Support

Governed & Auditable

Responses grounded in approved enterprise content

9:16 AM

• Is there a org for when a valve is not found for inspections

Last read

Op Ex Support Agent AI generated 9:16 AM

 Yes. There is guidance for when a valve is not found during inspections, and it's built into the EWAM workflow. Here's the key information:

How to Handle Valve Not Found

- In **MWM Mobile**, when performing a valve inspection and you **cannot locate the valve**:
 - Select "**VNF**" (**Valve Not Found**) on the route.
 - Provide a reason and comments (e.g., "Valve retired," "Obstructed by pavement").
 - Complete the rest of the route as usual.
- **What Happens Next:**
 - The original route step is marked as completed.
 - **Maximo** automatically creates a new follow-up work order for that valve, with the same due dates, so it can be revisited later.
 - A prerequisite (**prereq**) may appear on the original WO, and you will see the additional WO in Maximo under **Related Records**.
- **When the valve is truly retired:**
 - Cancel that additional work order with appropriate reason code (e.g., "Retired Valve").
 - If unsure, escalate to the lead for confirmation.

 **Reference:**

- [EWAM Knowledge Share – Handling Not Found VnS](#)
- Training LT: "Notes on CGA and Not Found" 

 **Best Practice Tip:** Always add detailed comments when using VNF, as missing details can trigger compliance follow-ups.

Do you want me to pull a quick reference one-pager (**ORG-style**) summarizing VNF process with screenshots for field techs and back-office reviewers? Or create a decision tree for "Not Found" scenarios (Valve vs CP Test Point vs Residential/Building)?



AI and robotics – Utility work of the future

~13~



