AI and Liability Exposures

Managing AI Risk in the Utility Sector

The Fastest Technology Adoption in History

Adoption at unprecedented speed

–From curiosity to board priority in under three years

–“You won’t lose your job to AI — you’ll lose it to someone using AI.” — Jensen Huang, NVIDIA

–Employees use AI with or without policies; vendors embed it by default

–The liability parallel: risk comes from using AI poorly, not from using it

A Utility Scenario: When AI Gets It Wrong

August. A severe storm. Thousands without power.

–AI drafts outage updates, summarizes field reports, prepares customer messages

–An AI-generated update posts an incorrect restoration estimate

–Customers, municipalities and media all rely on it — and it's wrong

–The real question: the AI, or the process around it?

Most AI liability issues aren't technology failures. They're governance failures.

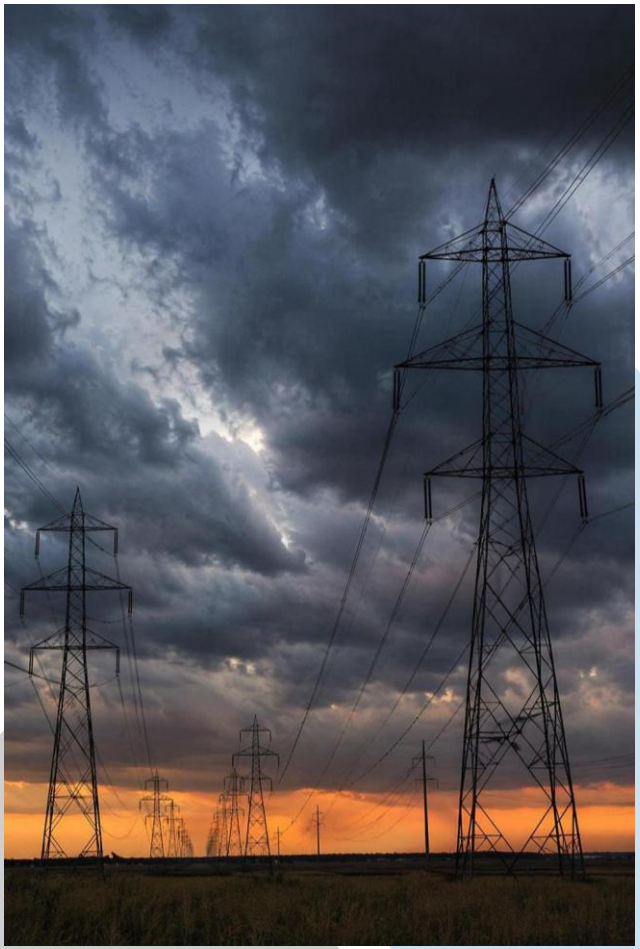

Four Emerging AI Liability Categories

Most AI claims arise from familiar concepts — negligence, discrimination, privacy, misrepresentation. AI just changes how those failures occur.

1 Bias & Discrimination
AI trained on historical data can replicate past hiring and demographic bias.

2 Error & Misinformation
Generative AI can be confidently wrong in customer-facing communications.

3 Privacy & Confidentiality
Employees paste sensitive data into public AI tools, creating exposure.

4 Governance & Regulatory Scrutiny
Regulators ask not whether you used AI, but whether you used it responsibly.

Category 1: Bias & Discrimination

AI learns from data — including human bias

–A utility adopts AI recruitment to speed hiring and screen applicants

–Trained on historical hiring data, it can replicate past demographic patterns

–Older or female applicants may be ranked lower — with no intent to discriminate

–“The algorithm decided” is no defense — the employer owns the decision and must show fairness

Category 2: Error & Misinformation

Confidently wrong

–Generative AI can sound authoritative — and be inaccurate

–During a storm, AI summarizes field reports and drafts customer updates

–One update carries a wrong restoration estimate; customers, partners and media rely on it

–Courts and insurers focus less on the AI's mistake than on whether oversight was reasonable

Category 3: Privacy & Confidentiality

The simplest risk: employees paste data into AI

–An employee uploads customer account data into a public AI tool to analyze trends

–The intent is innocent — but was disclosure authorized? Where is the data stored? Who can access it?

–It may create privacy obligations or a reportable incident

–The greatest privacy risk is often human behaviour, not the technology

Category 4: Governance & Regulatory Scrutiny

The category that connects all the others

–A utility rolls out enterprise AI; usage spreads across HR, customer service and reporting
–An issue emerges — a privacy complaint, a dispute, inaccurate customer information
–Regulators ask: Who approved it? Was there a risk assessment? Training? Oversight?
–The question isn't "Did you use AI?" — it's "Did you use AI responsibly?"

It’s Not the AI. It’s the Governance Around It.

When an incident occurs, the question shifts from “What did the AI do?” to “What did the organization do?” — that’s where liability is determined.

Managing risk **together**

Why AI Liability Is Different

We've always faced human error, discrimination and privacy claims. Three things make AI different.

Scale
One biased manager affects a few decisions; a biased AI can affect every decision at once.

Speed
AI generates and distributes instantly — and mistakes spread just as fast.

Explainability
Staff may not know how AI reached a result, yet the organization must still justify it.

Accountability stays human — even as decisions become machine-assisted.

The Accountability Gap & the AI Risk Spectrum

The Accountability Gap

AI recommended it. The employee relied on it. The organization deployed it. The vendor built it.

So who is responsible? Usually the organization.

AI can assist decisions — but it does not transfer accountability.

Four Places It Goes Wrong

- **The AI** — hallucinations and misleading recommendations
- **The data** — garbage in, garbage out
- **The user** — poor questions, over-reliance, no verification (often the biggest risk)
- **The governance** — no policy, training, oversight or documentation

AI → Employee → Decision → Outcome — liability emerges at Decision and Outcome

When AI Goes Wrong: Four Real Cases

These aren't hypotheticals. Each maps to one of our four liability categories — and in every case, accountability landed on the organization, not the AI.

1 EEOC v. iTutorGroup (2023)

Bias & discrimination. Hiring software auto-rejected 200+ older applicants; the first EEOC AI-bias settlement — $365,000. The employer, not the algorithm, was accountable.

2 Moffatt v. Air Canada (2024)

Error & misinformation. A website chatbot gave a customer the wrong fare policy. A tribunal rejected "the AI did it" and held the airline liable for the output.

3 Samsung ChatGPT Leak (2023)

Privacy & confidentiality. Employees pasted confidential source code into a public AI tool, prompting Samsung to ban it. Human behaviour, not the tool, created the exposure.

4 SEC AI-Washing Cases (2024)

Governance & regulatory scrutiny. The SEC's first "AI-washing" actions fined two advisers for overstating AI use — regulators scrutinize the claims and the oversight behind them.

Same lesson every time: the organization owns the outcome — "the AI did it" is never the defense.

When AI Causes a Loss: The Insurance Puzzle

One AI event can trigger multiple policies — or none exactly as expected.

Type of Claim	Potential Coverage
Employment discrimination	**EPL**
Privacy breach	**Cyber**
Professional advice error	**E&O**
Board oversight failure	**D&O**
Bodily injury / property damage	**CGL**

AI doesn't fit neatly into one category — one event may spark disputes over which policy responds.

The Coverage Gap & the Assumption Gap

The Coverage Gap

Are existing policies keeping pace with AI? Insurers now ask:

–Was the AI approved, and were controls followed?
–Was there human oversight?
–Was use authorized or unauthorized?
–Technology failure, or human reliance?
–Did a third-party vendor contribute?

The Assumption Gap

The bigger risk is assuming —

–that technology involvement means it's insured
–that the vendor is responsible
–that AI reduces the organization's liability

Insurance is no substitute for governance — strong governance shows the organization acted reasonably.

What Good AI Governance Looks Like

The goal isn’t to eliminate risk — it’s to use AI confidently, responsibly and defensibly. Seven practical principles follow.

Managing risk together

Seven Governance Principles (1–4)

1 Know Where AI Exists
You can't govern what you can't see — build an AI inventory and tackle "shadow AI."

2 Classify Risk Before Deployment
A draft agenda isn't a hiring screen — match oversight to low, medium and high risk.

3 Keep Humans Accountable
AI assists, but a person must approve outage estimates, safety summaries and candidate screens.

4 Protect Data Relentlessly
Control what data may be entered, which tools are approved and how compliance is monitored.

Seven Governance Principles (5–7)

5 Document Decisions
Record why AI was used, the risks and controls considered, and who approved it — that creates defensibility.

6 Train, Train, Train
Policies aren't enough — teach staff what AI can and can't do, how to verify outputs, and when to escalate.

7 Make AI a Boardroom Topic
AI governance is enterprise risk governance — boards should know where AI is used and how accountability is maintained.

The Governance Test

If regulators, insurers or auditors arrived tomorrow, could you confidently answer:

1. **Where are you using AI?**
2. **Who is accountable?**
3. **What controls exist?**
4. **How are outputs validated?**
5. **How is risk monitored?**

Good governance isn't a perfect policy — it's demonstrating AI is used thoughtfully, responsibly and with oversight. That may be **the strongest defense** an organization has.

Looking Ahead: Four Trends

What utility leaders should watch over the next three years

1 AI Will Become Invisible
It's embedded into everyday software — you'll adopt AI without realizing it.

2 Regulators Will Expect More
"We didn't know how it worked" is ending — expect AI-specific rules and governance standards.

3 Governance Becomes an Advantage
Like cybersecurity, mature AI governance builds trust with customers, regulators and insurers.

4 The Human Role Grows
As AI gets more capable, judgment, ethics and accountability matter more, not less.

Final Takeaways

- **A governance issue** — AI is not primarily a technology problem.
- **Familiar legal concepts** — most claims arise from negligence, discrimination, privacy and weak oversight.
- **Used poorly, not used at all** — organizations face liability for using AI badly, not for using it.
- **Accountability and discipline** — the best-positioned organizations use AI with the greatest oversight.

AI may change how decisions are made.
It does not change who is accountable for them.

Managing risk **together**