



Powered by EPRI  
**OPEN POWER  
AI CONSORTIUM**



Learn More:  
[www.openpowerai.org](http://www.openpowerai.org)



**Our Work Groups**

**Domain-specific model**



Domain-specific model trained on energy-specific datasets

**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**



Domain-Specific Model Includes Knowledge from

**>10,000  
EPRI reports**



**250+**

use cases identified for consortium prioritization and development

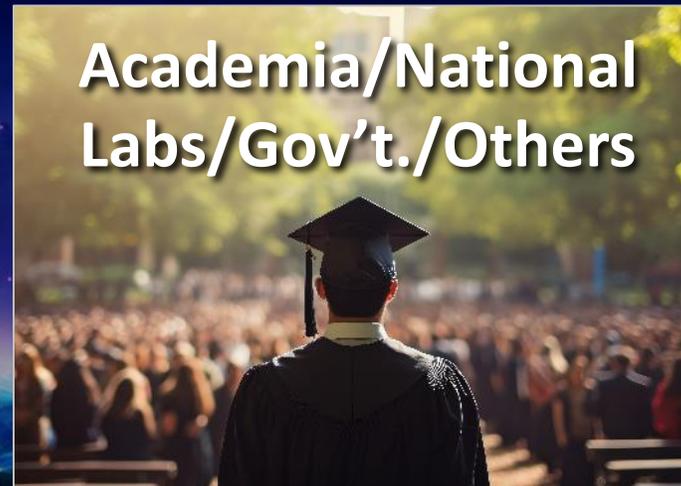
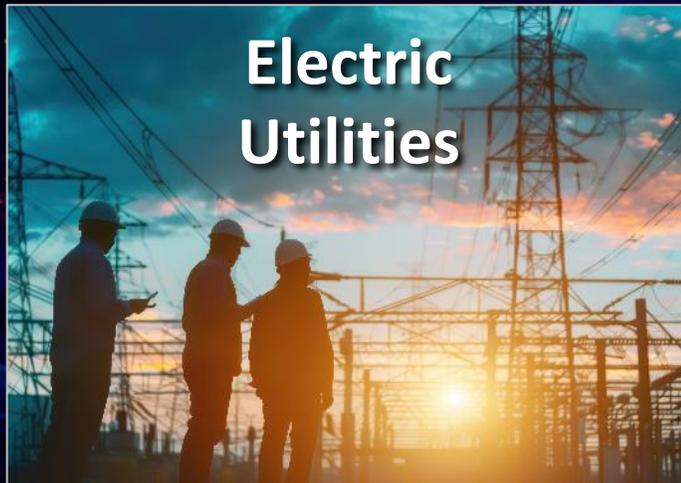
**Join us in Shaping the Future of AI and Energy**

# EPRI Launched the Open Power AI Consortium

## Bringing Enduring Value to the Global Electric Sector

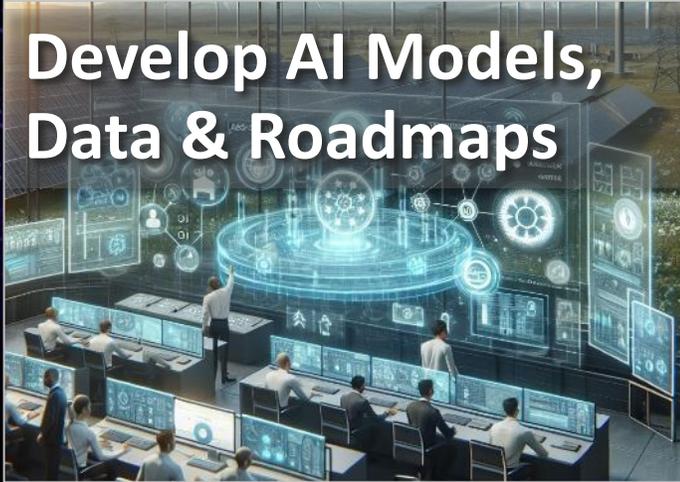


**OPEN POWER  
AI CONSORTIUM**

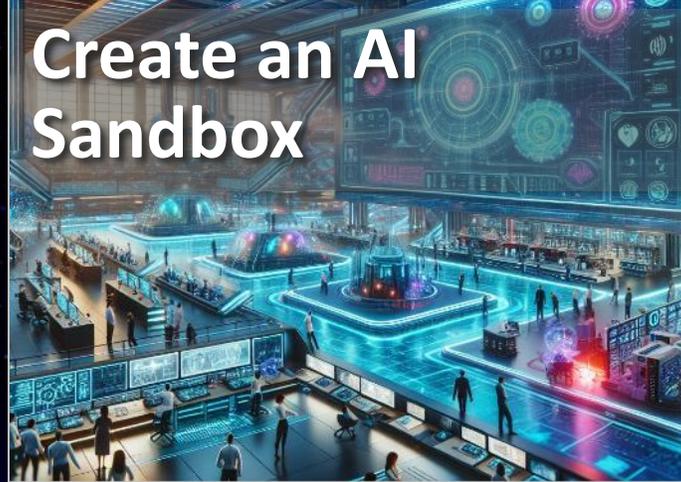


# Open Power AI Consortium – Key Objectives

**Develop AI Models,  
Data & Roadmaps**



**Create an AI  
Sandbox**



**Implementation and  
Lessons Learned**



**Create an ecosystem for stakeholders to identify, develop, validate and deploy AI solutions to transform the electric sector**

# OPA1 Participants

## AI and Energy Technology Partners



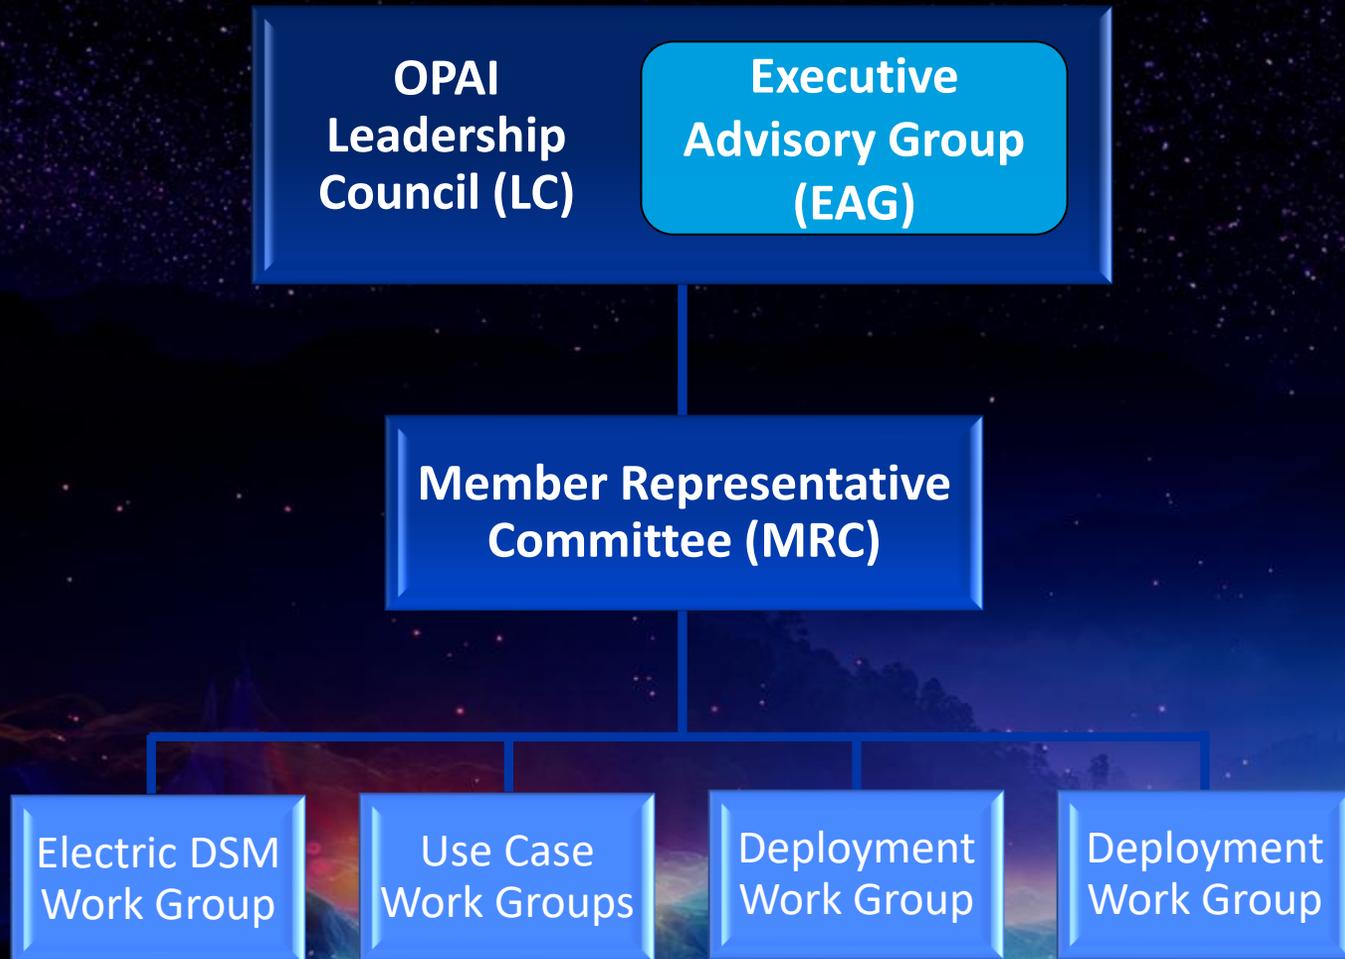
## Academia & Other Strategic Partners



## Energy Partners



# OPAI Consortium Member Engagement



## Consortium Engagement

Many companies continue to join the consortium

**>200** Confirmed Members

**>100** MOU/Requests In-Process

# Regional OPAI Consortium Events

Reuters Energy Live  
Houston, December 2025

DistribuTECH  
San Diego, Feb. 2026

PowerGen International  
San Antonio, January 2026

Enlit Europe  
Bilbao, November 2025

EPRI-Microsoft Event  
Munich, September 2025

EPRI Gulf Event  
Kuwait, April 2025

EPRI-Microsoft Event  
Japan, September 2025

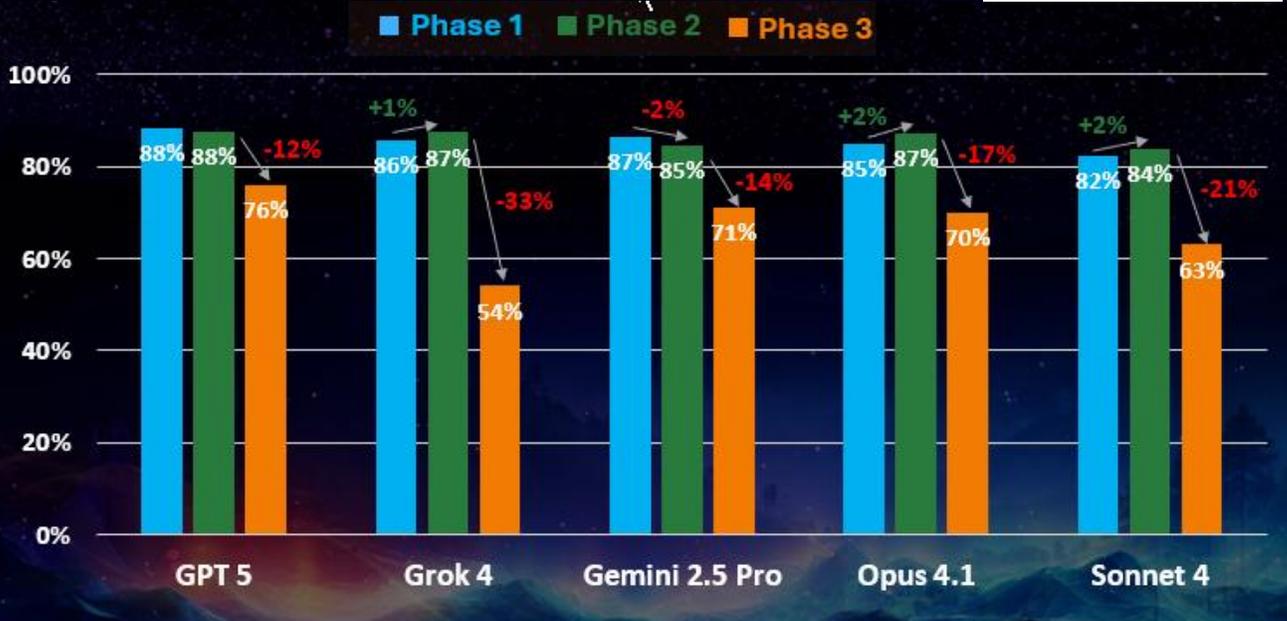
EPRI-Microsoft Event  
Australia, November 2025



# Recent OPAI Deliverables

## Benchmarking Performance of LLMs on Power Industry Questions:

- Model Accuracy depends on the model, available tools, and question type
- Models score higher on multiple-choice versus open-ended questions



2025 White Paper

AI Readiness in Utilities: Turning Data into Strategic Advantage



The first in a series of guidance documents to accelerate AI implementation

# OPAI Use Case Survey – Example Format



## Utility Survey



## Technology Provider Survey

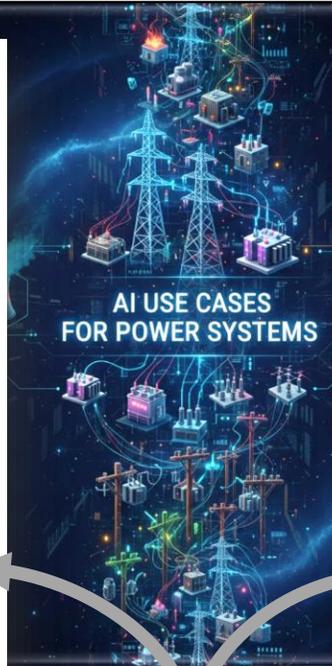
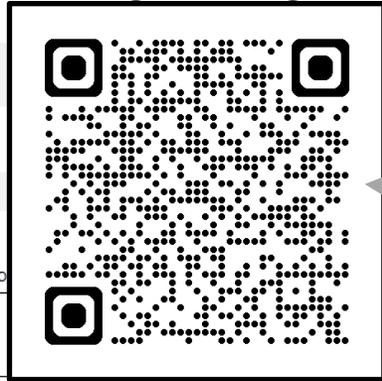
Same Use Case – Different Response

### 4. Advanced Metering Infrastructure (AMI) Data Validation

AI validates smart meter data by detecting anomalies, missing values, or implausible readings. Models compare real-time data against historical patterns and neighboring meters. Errors are flagged and corrected or imputed before billing or analysis. This ensures accurate consumption records and strengthens customer trust.

	Limited	Minimal	Moderate	Significant	Extensive
Impact - Company or Customer Cost Saving	<input type="radio"/>				
Impact - Reliability of Resilience	<input type="radio"/>				
Impact - Safety or Security	<input type="radio"/>				
Implementation Complexity	<input type="radio"/>				
Data Availability	<input type="radio"/>				
Risk of Misoperation	<input type="radio"/>				

Do you have an existing project for this use case in your company?...if so

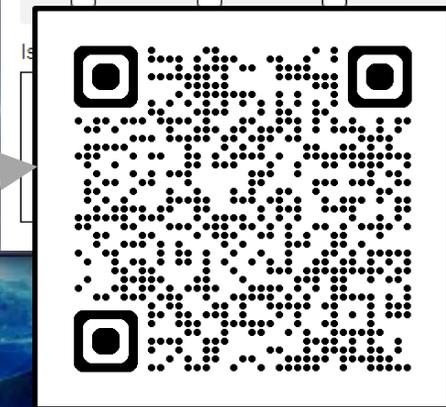


### 4. Advanced Metering Infrastructure (AMI) Data Validation

AI validates smart meter data by detecting anomalies, missing values, or implausible readings. Models compare real-time data against historical patterns and neighboring meters. Errors are flagged and corrected or imputed before billing or analysis. This ensures accurate consumption records and strengthens customer trust.

TRL 1: Basic principles observed and reported	TRL 2: Technology concept formulated	TRL 3: Proof of concept	TRL 4: Laboratory validation	TRL 5: Relevant environment validation	TRL 6: Demonstration in relevant environment	TRL 7: Prototype demonstration in operational environment	TRL 8: System completed and qualified	TRL 9: Actual system proven in operational environment
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Is there a demo of your technology for this use case? If so please elaborate



Click or Scan

<https://www.surveymonkey.com/r/OPAI-Use-Cases-Utilities>

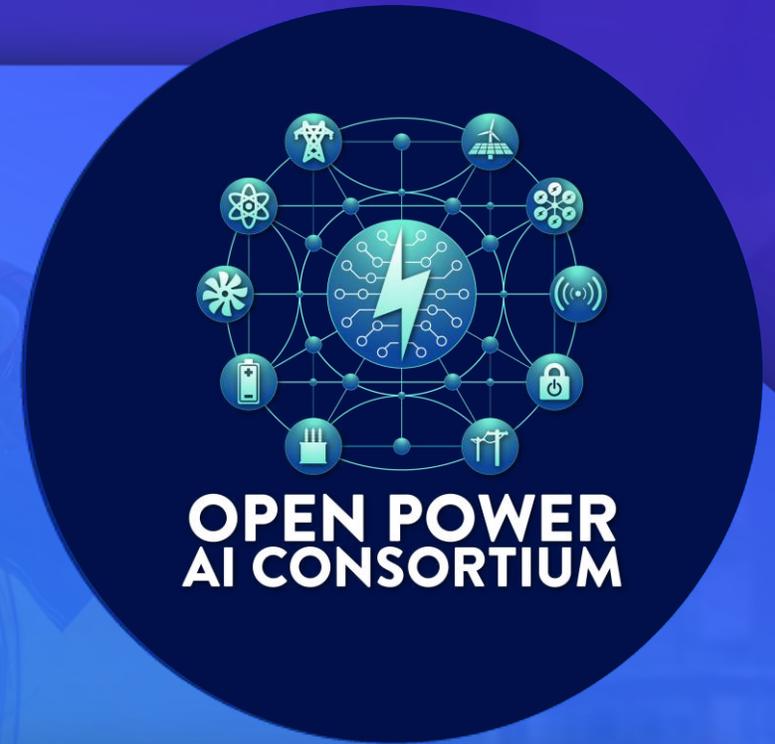
<https://www.surveymonkey.com/r/OPAI-Use-Cases-Tech-Providers>

# How to Get Involved?

- Send an email to [OpenPowerAI@epri.com](mailto:OpenPowerAI@epri.com) with the following information:
  - Intent to participate
  - Name a point of contact for the organization - typically the organization's AI lead
  - Desire to be part of the logo zoo



# Collaboration to Enable AI Solutions for the Electric Sector



D-TECH 2026  
San Diego

2026 February 05

# OPAI Key Activities

## Domain-Specific Model Working Group



## Use Case Working Group



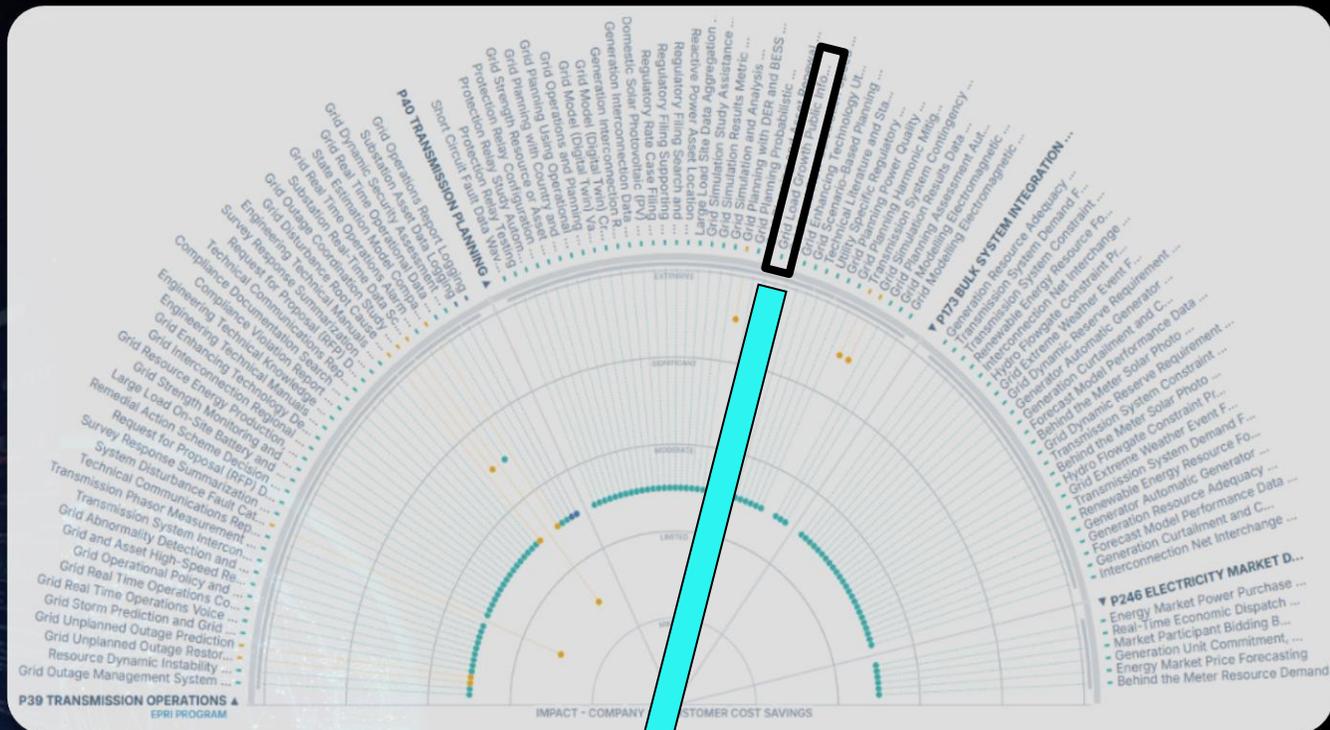
## Deployment Working Group



# Use Case Database

>250 Use Cases

## AI USE CASES FOR POWER SYSTEMS



## Grid Load Growth Public Information Search & Knowledge Retrieval

Information retrieval systems (such as LLMs) automatically scan public records, planning document, and news sources for data on anticipated load growth. AI extracts and synthesizes relevant information, such as new developments or industrial expansions, feeding it into grid planning models. This enhances the accuracy of demand forecasts and investment decisions. The process reduces manual research time and ensures planners have up-to-date-intelligence

# Use Case Survey By the Numbers – So Far!



Utilities

Technology Providers

53

*75+ Entities Submitted at  
Least 1 Response*

22

261

*Use Cases Evaluated by at  
Least 1 Person*

150

289

*338+ People Evaluated*

49

10875

*11258+ Data Entries (Evaluation,  
Maturity or Comment)*

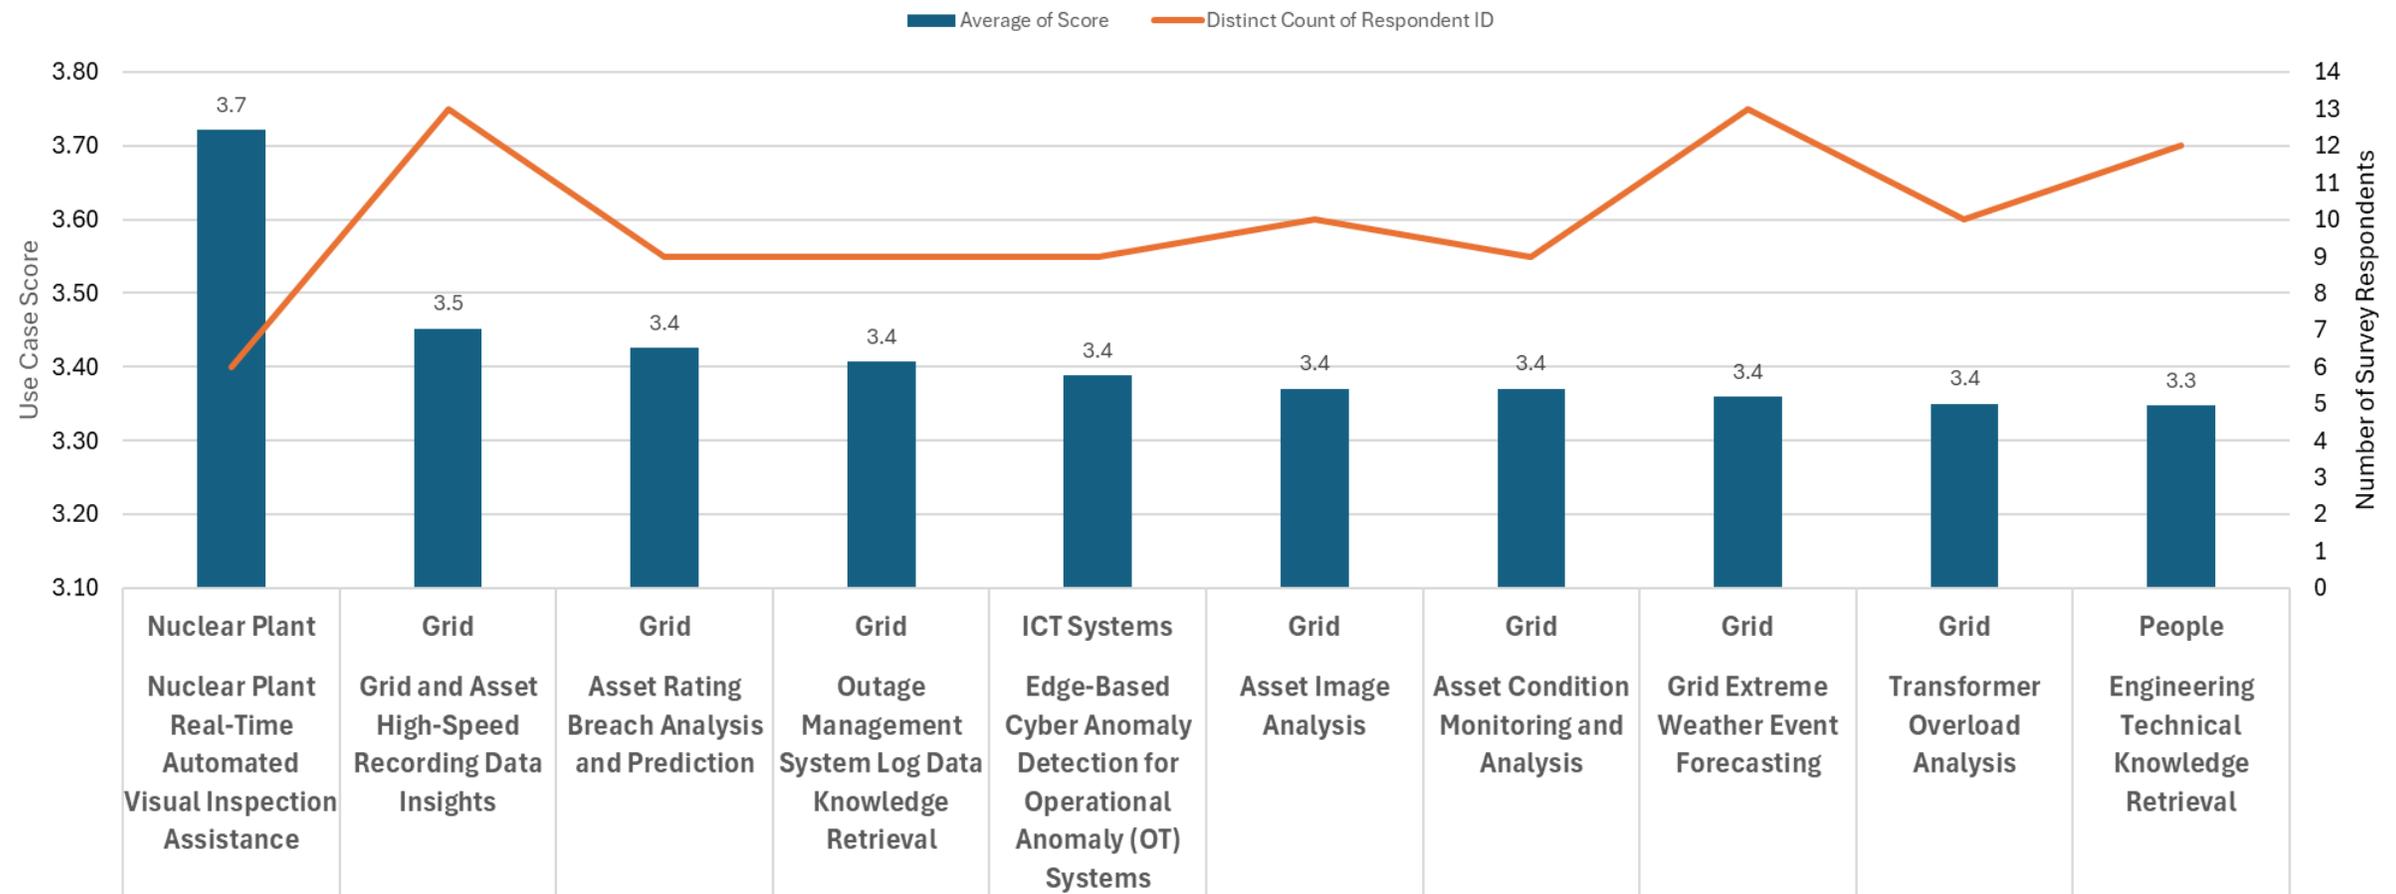
383

*Assuming ~20 Seconds Entry – Cumulative Total of ~62 hours of data entry  
Thank You!*

# Top 10 Use Cases – All Functions (6+ Respondents)

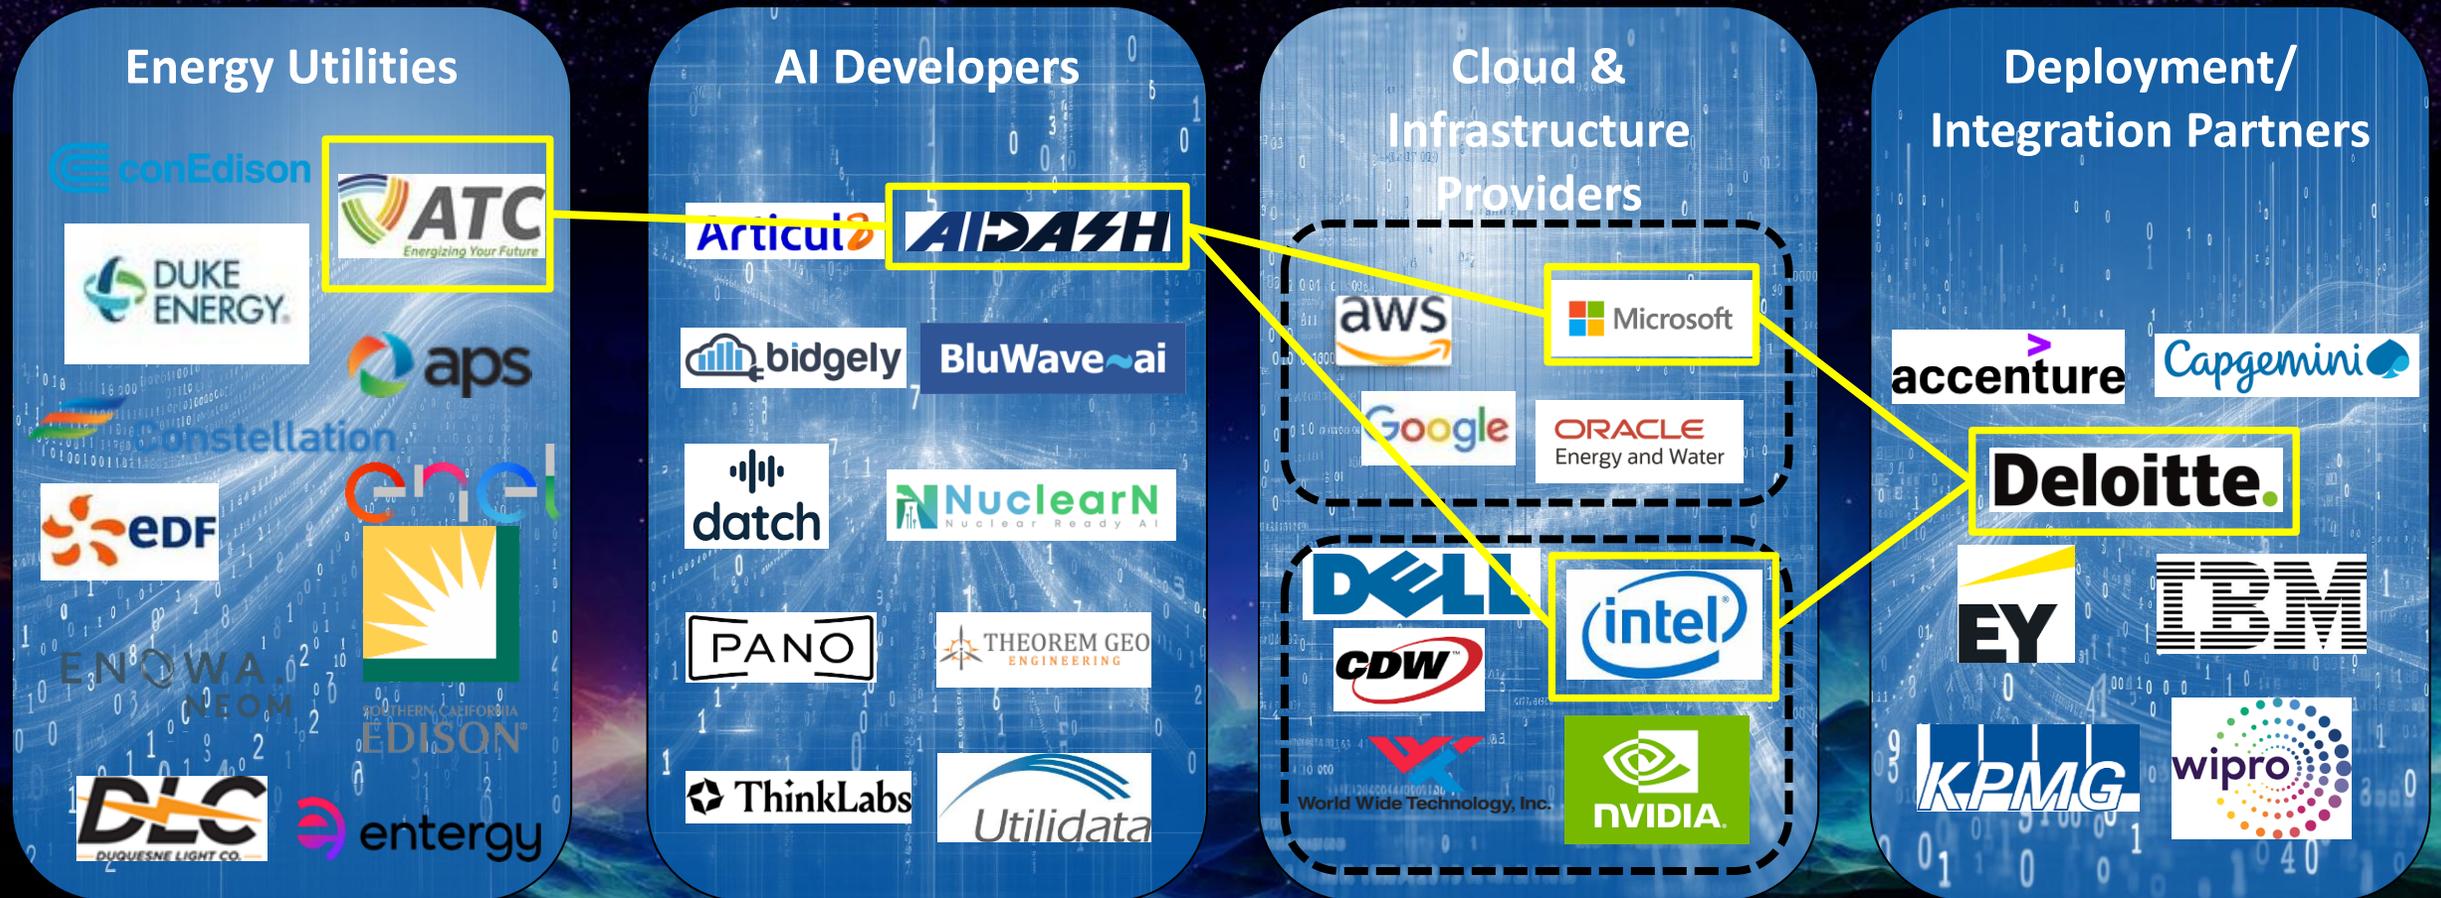
Nuclear use cases is top identified, but grid-based use cases are 7 of top 10.

Use Case Score (Out of 5 LHS Axis) with Number of Respondents (RHS Axis)



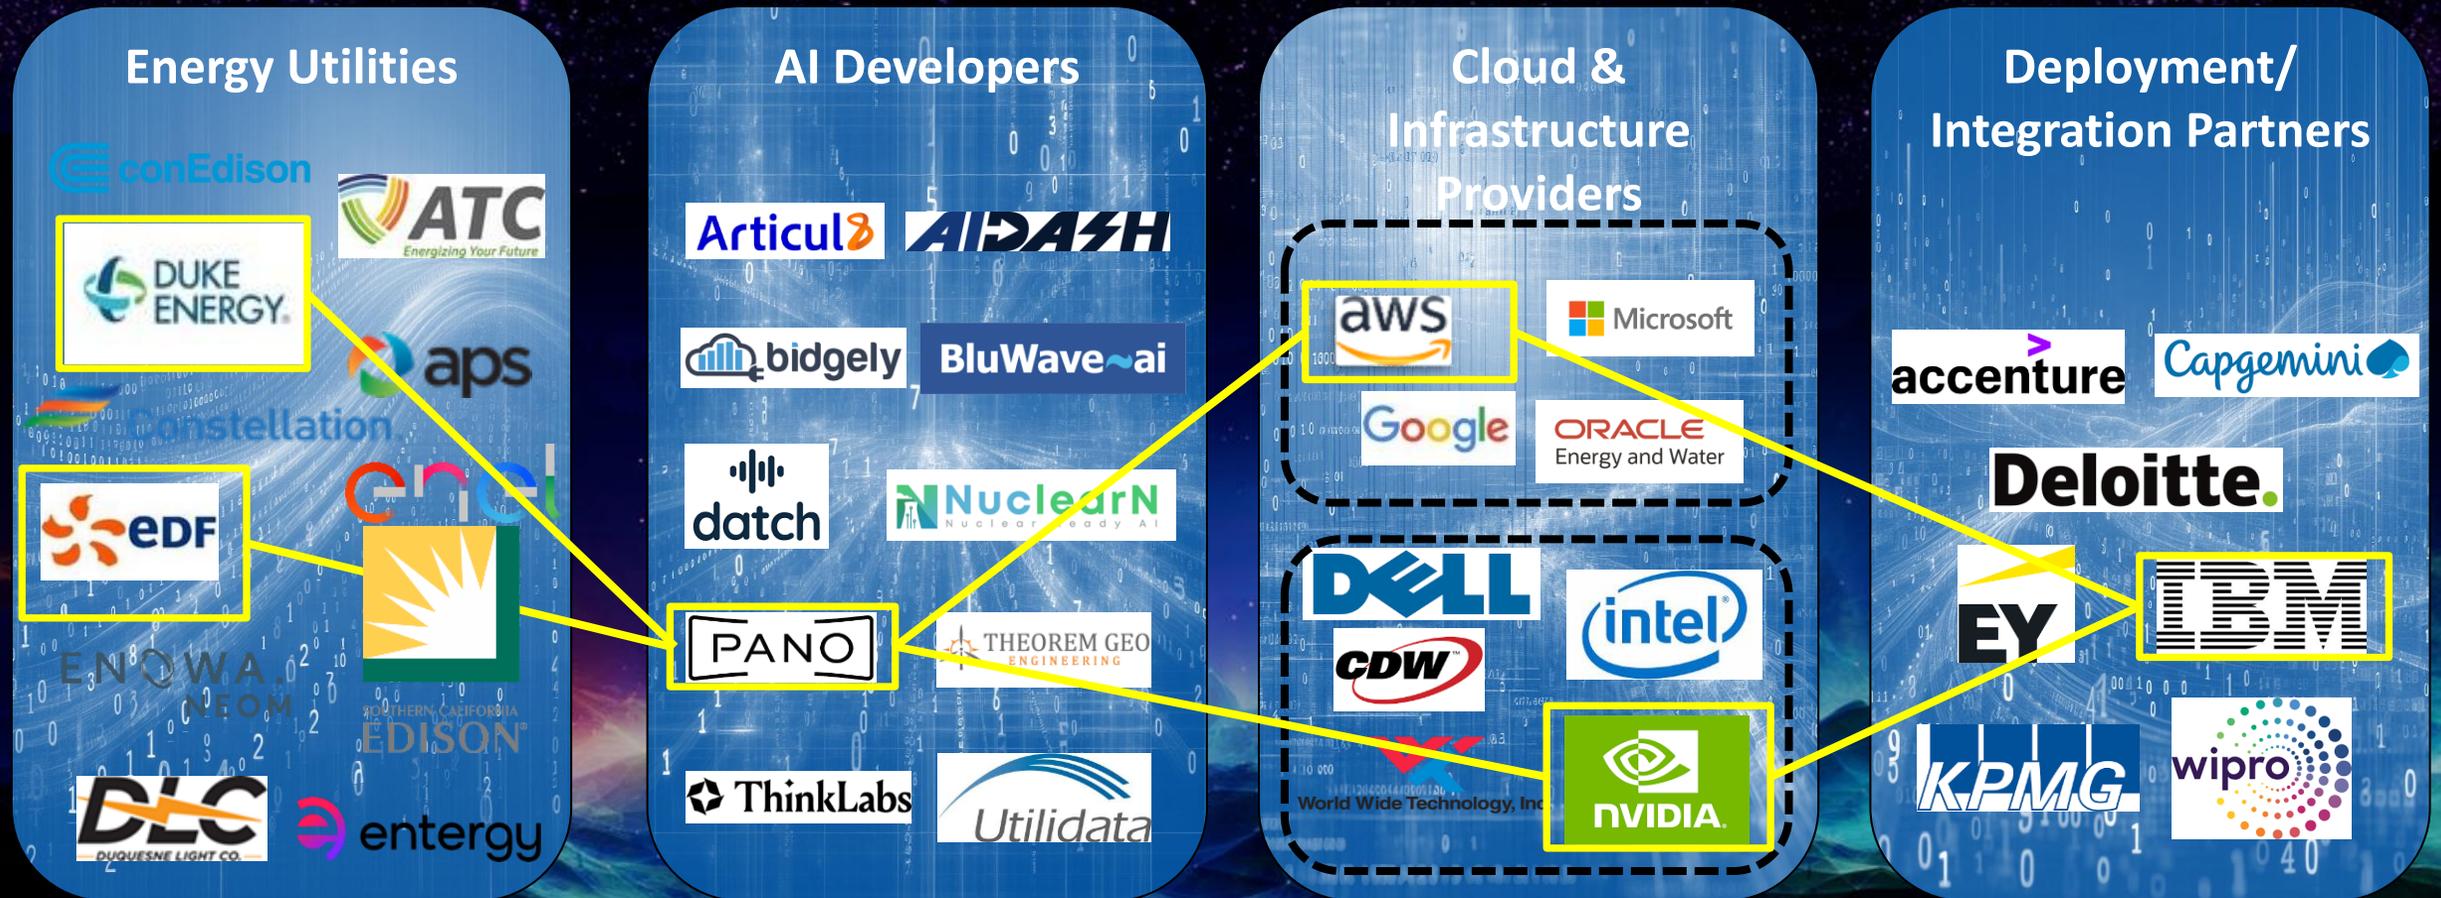
# Co-Developing and Deploying AI Solutions with Partners

Priority Use Case #1: Interested OPAI Members commit to developing, piloting, and (potentially) deploying solutions (hypothetical example for Use Case #1)



# Co-Developing and Deploying AI Solutions with Partners

Priority Use Case #2: Interested OPAI Members commit to developing, piloting, and (potentially) deploying solutions (hypothetical example for Use Case #2)



# Example AI Projects Informed through OPAI

**Generator Reliability  
Expert Model**



Generator Reliability



EPRI aLLarMa

**AI-Powered Alarm  
Management &  
Ops Data Tool**

**Transmission Planning  
& Protection Study  
Assistant**



Raphson



NUCLEAR LLM

**Nuclear Industry  
Virtual Assistant (NIVA)  
(Multi-Org Collaboration)**



# AI Enabled Power Transformer Health: Better answers?

Sampling

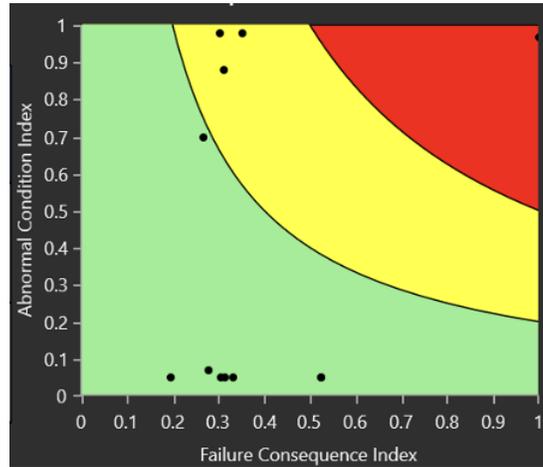


On-line



Measurements

**EPRI**  
**PTX**  
Algorithms  
(AI-based)



Fleet Health Assessment +

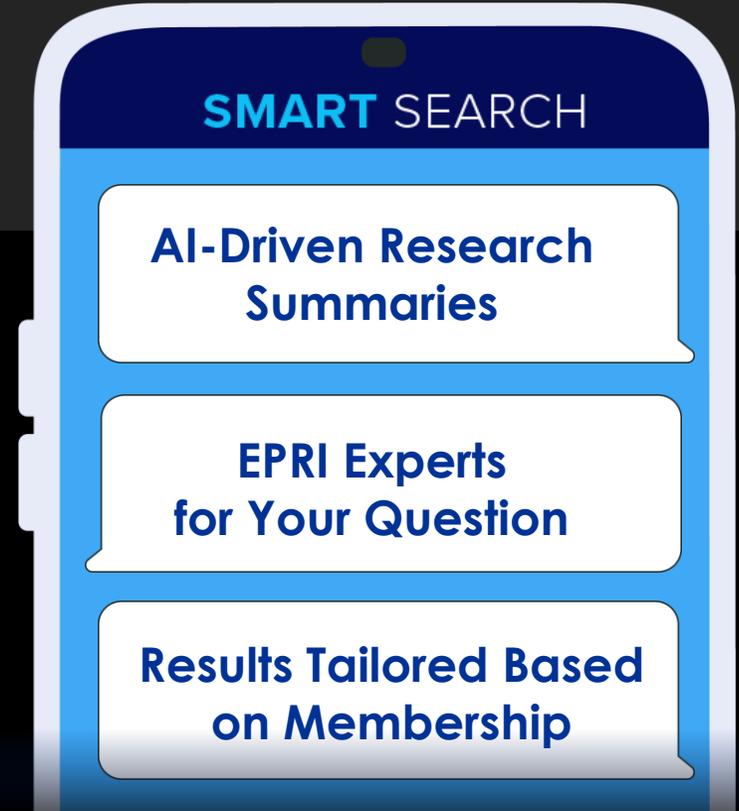
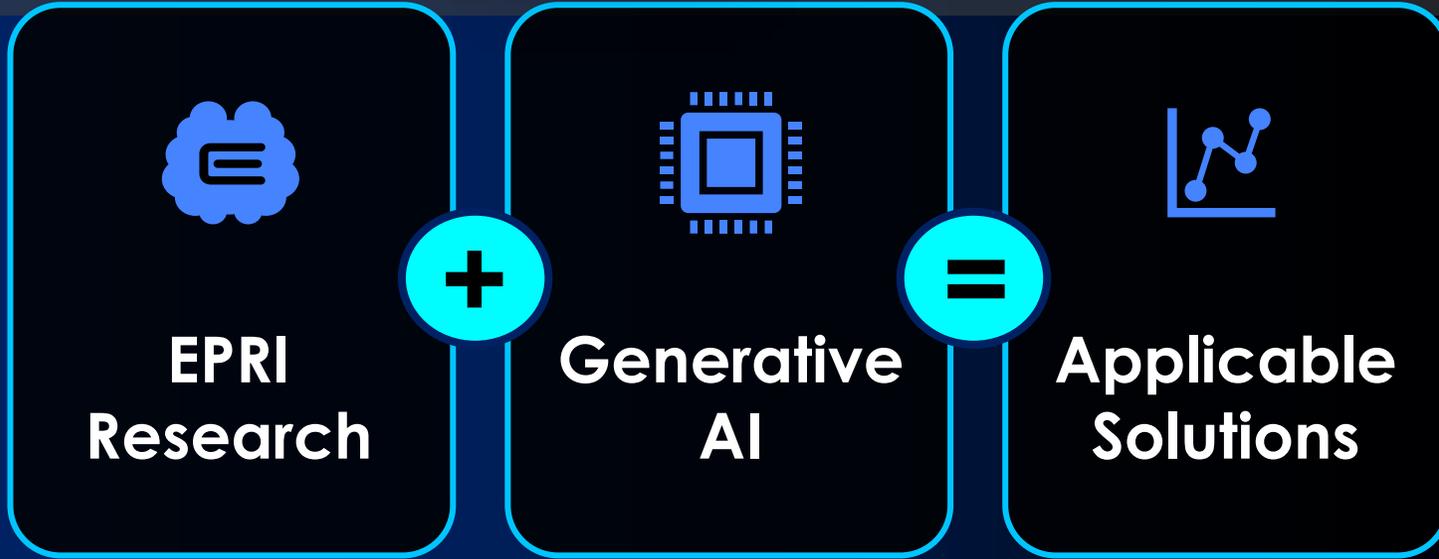
Belief of Fault Conditions



# SMART SEARCH\*

Powered by  EPRI.AI

Leveraging EPRI Research and Generative AI to Transform Member Experience



## Progress to Date

- Internal AI Research Tool
- Smart Search Research Tab
- Smart Search People Tab

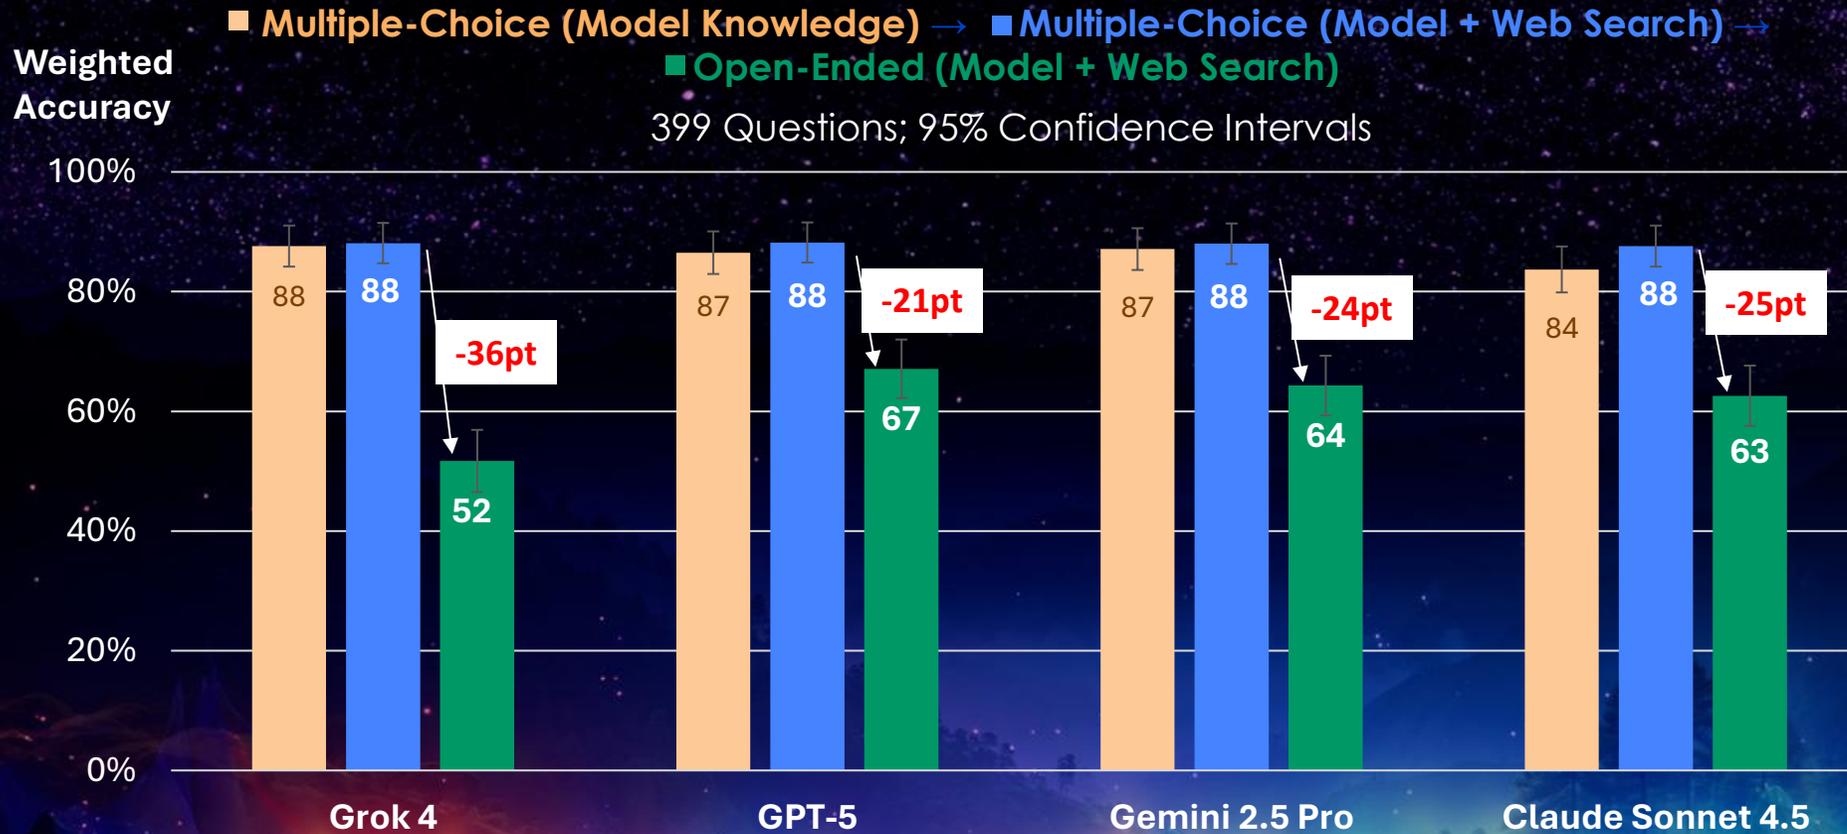
## What's Next in 2026

- Limited Member Beta Test
- Additional Smart Search Tabs
- Integrate into EPRI.com

\*Platform product name subject to change

# Model Accuracy depends on the model, available tools, and question type

## EPRI's Power-Systems Benchmarking – from multiple choice to open-ended questions



Additional LLM Details: Grok 4: grok/grok-4-0709; GPT 5: openai/gpt-5-2025-08-07; Gemini 2.5 Pro: google/gemini-2.5-pro; Claude Sonnet 4.5: anthropic/claude-sonnet-4-5-20250929



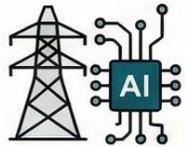
Accuracy of 52-67% on single prompt, open-ended power system benchmark questions further **highlights** the need for **use-case benchmarking**

# Utility Veritas: AI Governance and Alignment for the Power Sector

## Why Utilities Need This



Safety & Reliability



OT-Aware AI Use

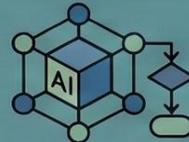


Regulatory Confidence

## UTILITY VERITAS OPERATING MODEL



**Data**  
Classify & Protect



**Models**  
Document & Align

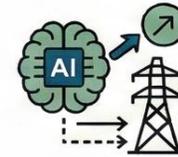


**Testing**  
Validate & Monitor

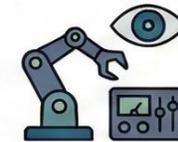


**Assurance**  
Trust & Compliance

## What Utility Veritas Aligns



AI Behavior with Grid Outcomes



Autonomy with Human Oversight



Innovation with Risk Tolerance

Scales by Risk



Applies Across Utility AI Use Cases



Planning & Forecasting



Asset & Maintenance



Operations & Restoration



Customer & Enterprise

Inspired by MAS Veritas (Monetary Authority of Singapore)  
<https://www.mas.gov.sg/schemes-and-initiatives/veritas>



# The Intelligent Grid at Scale

The AI Factory for the Utility of the Future

Rich Kropfl | Area Vice President, Field CTO | World Wide Technology

EPRI | DTECH 2026

AI Session – Open Power AI Consortium

# The Pace of AI Innovation Has Fundamentally Changed

AI progress is no longer linear—it is accelerating simultaneously across infrastructure, models, and systems..



## Infrastructure Is Advancing at Unprecedented Speed

Ampere → Grace Hopper → Blackwell → Vera Rubin

Order of magnitude gains per generation



## AI Capabilities Are Advancing Faster Than Organizations Can Absorb

Chat → RAG → Tool Use → Agentic → Reasoning → Physical AI

Each step increases autonomy and operational impact

---

## The Result: A Compound Acceleration

Infrastructure efficiency, model capability, and system integration are advancing together—not independently

---

**AI evolves in months.**

**Grid modernization evolves in decades**

---

The question is no longer **if AI will advance**,  
It is whether the grid **can absorb that pace safely**



CORE CHALLENGE

# The Central Tension: AI Velocity vs Grid Reliability

AI innovation is accelerating.

Models, compute, and capabilities are advancing at an unprecedented pace.

The grid must operate deliberately.

Safety, reliability, and regulatory trust are non-negotiable.

The challenge is absorption.

How do we integrate breakthrough capabilities into systems that cannot afford to fail?

---

***Critical infrastructure cannot adopt AI at the speed of hype.***

INFLECTION POINT

# The AI Efficiency Inflection Point

AI progress is no longer driven by isolated breakthroughs

It is compounding across the entire technology stack

Three forces are converging — and together they change what is operationally possible for the grid::



## Compute efficiency

Accelerated architectures are delivering step-function gains in performance per watt



## Model efficiency

Smaller, domain-specific models increasingly outperform brute-force approaches in real-world use cases.



## System efficiency

End-to-end pipelines outperform isolated models and one-off deployments in production.

**The result:** AI capabilities that were previously impractical are now operationally viable at scale

# From Episodic AI Projects to a Continuous AI Factory

*Utilities have always delivered change through projects. What's changed is the pace, coupling, and compounding nature of AI innovation.*

**AI Factories turn innovation into repeatable, industrial outcomes — at scale — through an integrated, secure, continuously improving flywheel.**

## BEFORE

### Episodic AI Projects

- Discrete pilots and implementations
- Optimized for one-time outcomes
- Heavy reintegration per use case

## AFTER

### The AI Factory

- A **continuously expanding**, governed system
- Designed for **repeatable, incremental change**
- Tight **feedback loops** across data, models, and operations
- **Security, governance, and safety embedded by design**
- Built for **sustained production**, not one-off success

This is not a shift from “projects” to “platforms.” It is a shift to a **continuously improving, governed AI flywheel** — designed for scale, trust, and longevity.

# Hybrid AI Is How the Grid Operates — Not a Deployment Choice

Cloud, on-prem, and edge operate as execution domains within a single AI Factory.



**Hybrid is not a compromise**

**It is the operating model required to absorb AI safely across the grid**

# Hybrid AI Factory Reference Architecture for the Grid

*A production system for building, deploying, and operating AI **safely and repeatedly** at scale*

## Data & Signal Fabric

Secure ingestion and governance of OT, IT, edge, and external signals across the grid..

## Compute Fabric

Purpose-built execution across cloud, on-prem, and edge environments.

## Model Lifecycle Platform

Training, validation, deployment, and continuous improvement across the model lifecycle.

## Security, Compliance & Trust Plane

Policy, explainability, lineage, and regulatory controls.

## Operations, Automation & Agents

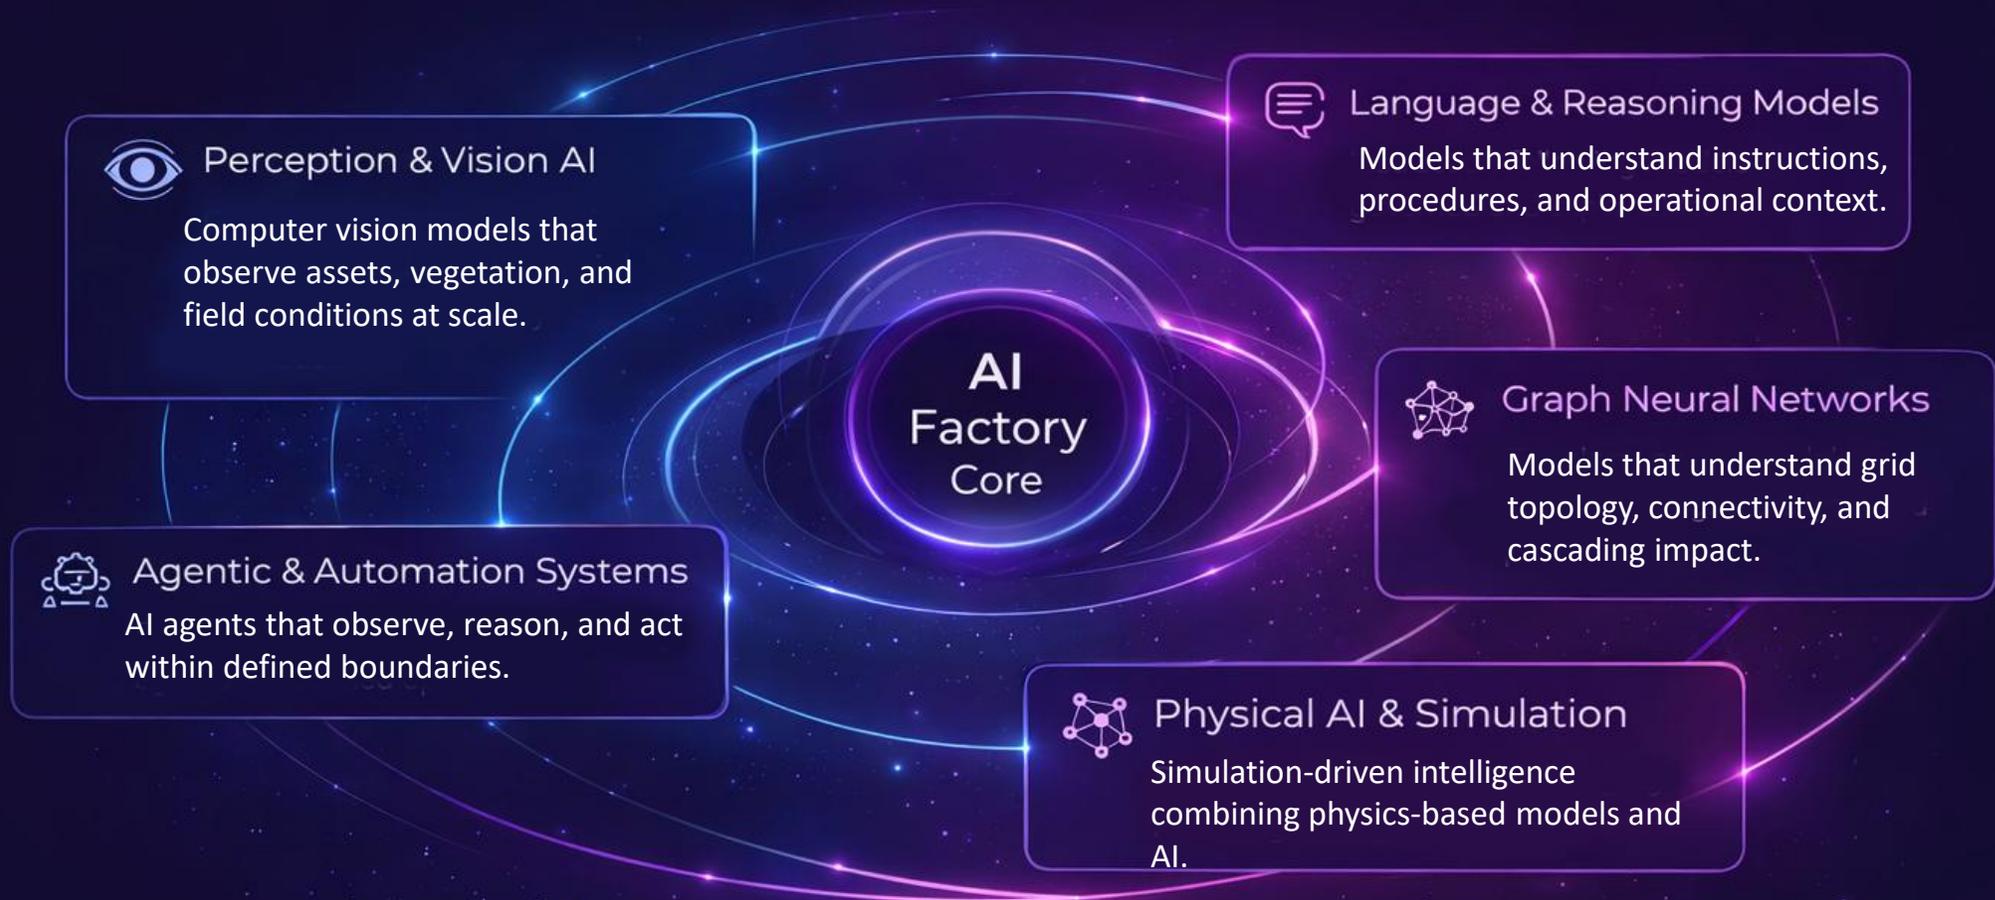
Monitoring, orchestration, and human-in-the-loop, AI-assisted operational workflows.

---

*Architecture is how innovation is governed.*

# Where Advanced AI Lives in the Factory

Multiple forms of intelligence, operating together inside one governed AI Factory



**The factory is not optimized for one model—but for continuous, multi-modal intelligence.**

# The AI Factory Is a System — Not a Deployment

The AI Factory is how utilities absorb continuous AI innovation without destabilizing operations.



## Stable Foundations

- Repeatable deployment patterns across use cases
- Faster path from research to production



## Always On Lifecycle

- Models evolve without re-architecting platforms
- Operational feedback loops into engineering



## Governed Autonomy

- Policy-constrained experimentation
- Clear separation between AI autonomy and human authority

---

**The result:** Innovation accelerates. Operations remain stable.

---

# The AI Factory Is How Utilities Scale AI — Safely

Turning continuous AI innovation into operational capability for the grid.

AI innovation is now continuous — not episodic. Utilities can no longer restart architecture, governance, or operations for every new model.

---

The grid requires hybrid, governed execution by design. Cloud, on-prem, and edge are execution domains within one integrated system — not competing choices.

---

The AI Factory converts velocity into reliability. Faster outcomes without compromising safety, trust, or operational control.

---

This is how AI scales in critical infrastructure.





# AI for Energy

Driving innovation, collaboration, and sustainable impact



## **NVIDIA**

Pioneering Accelerated Computing and AI

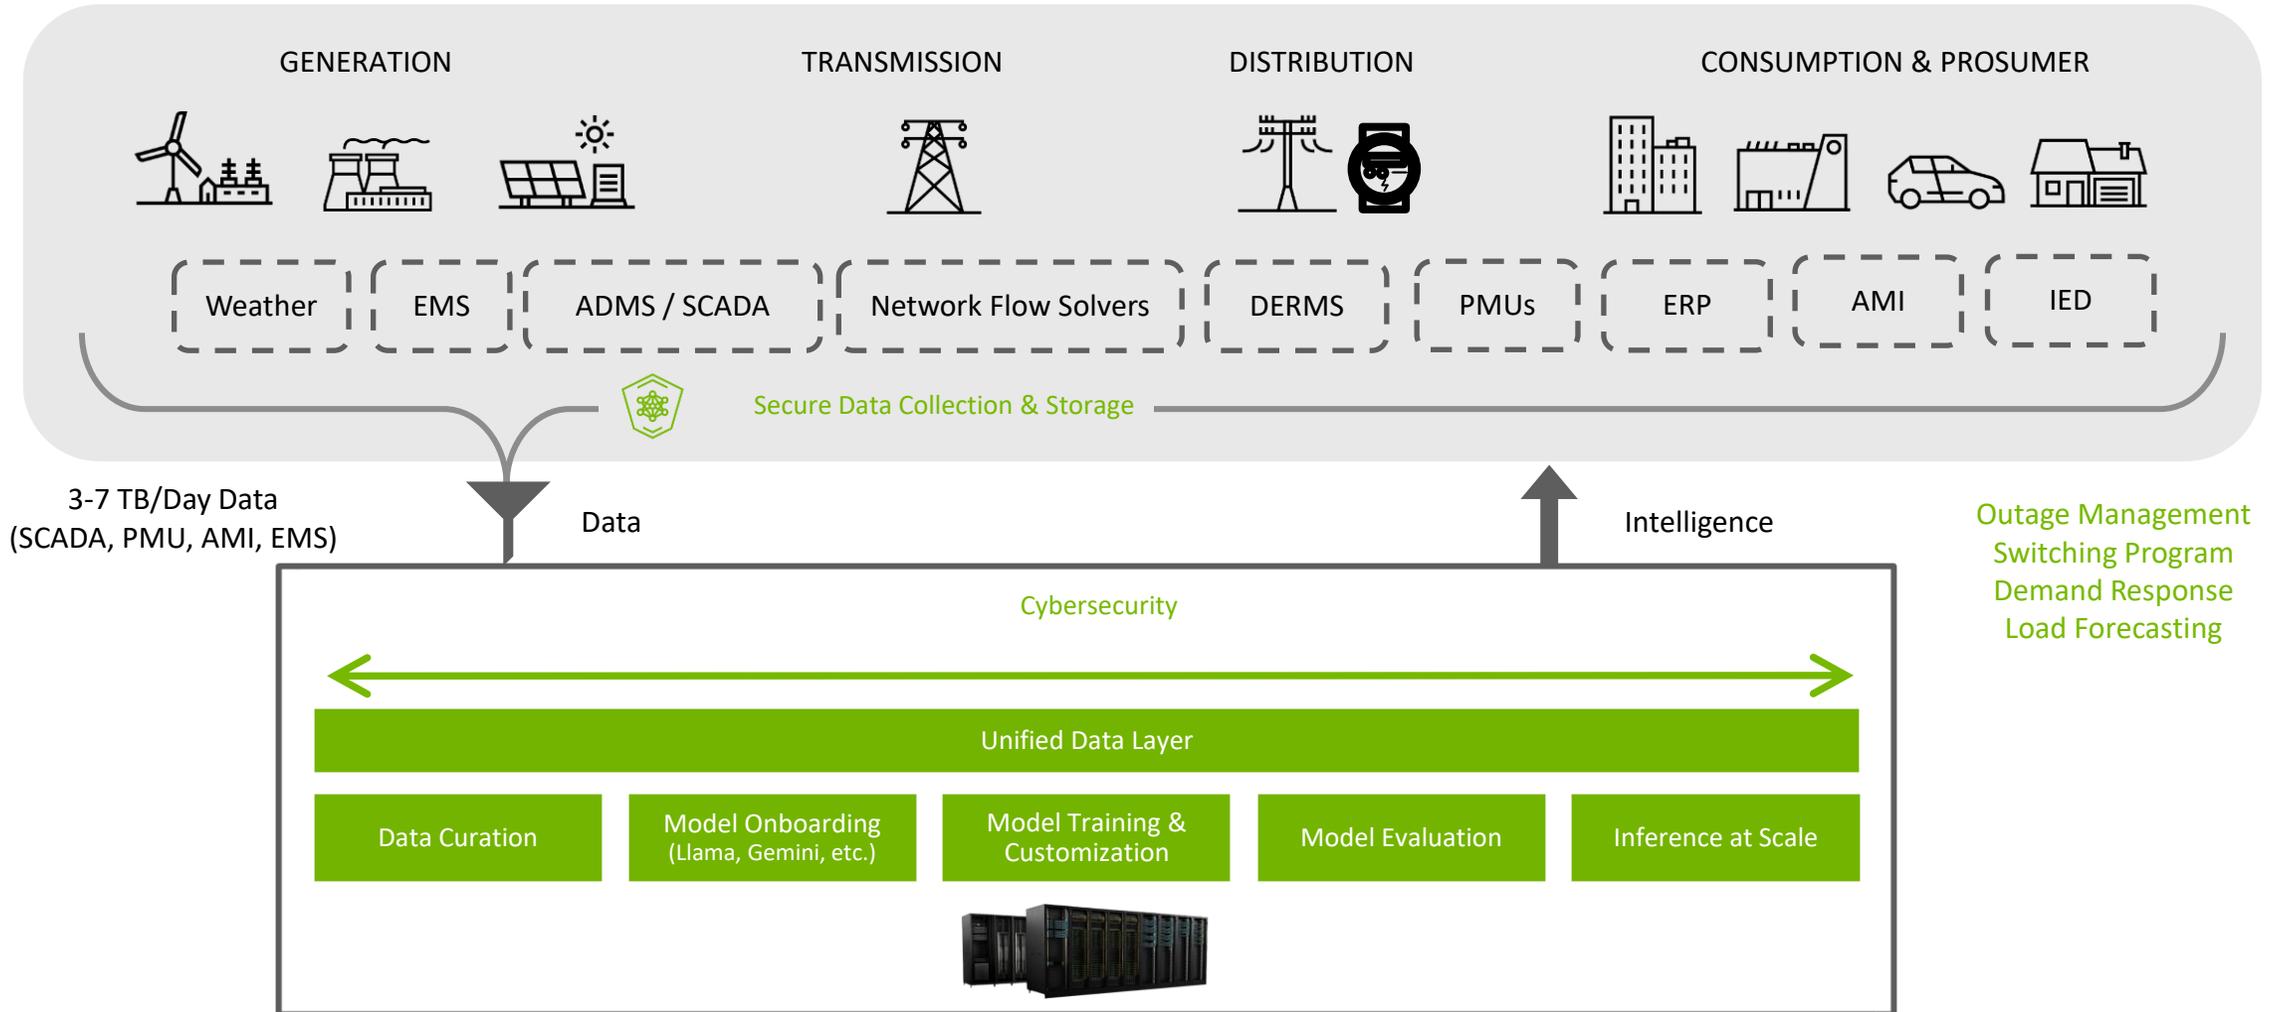
Full stack – hardware, networking, software

From cloud to data center to edge

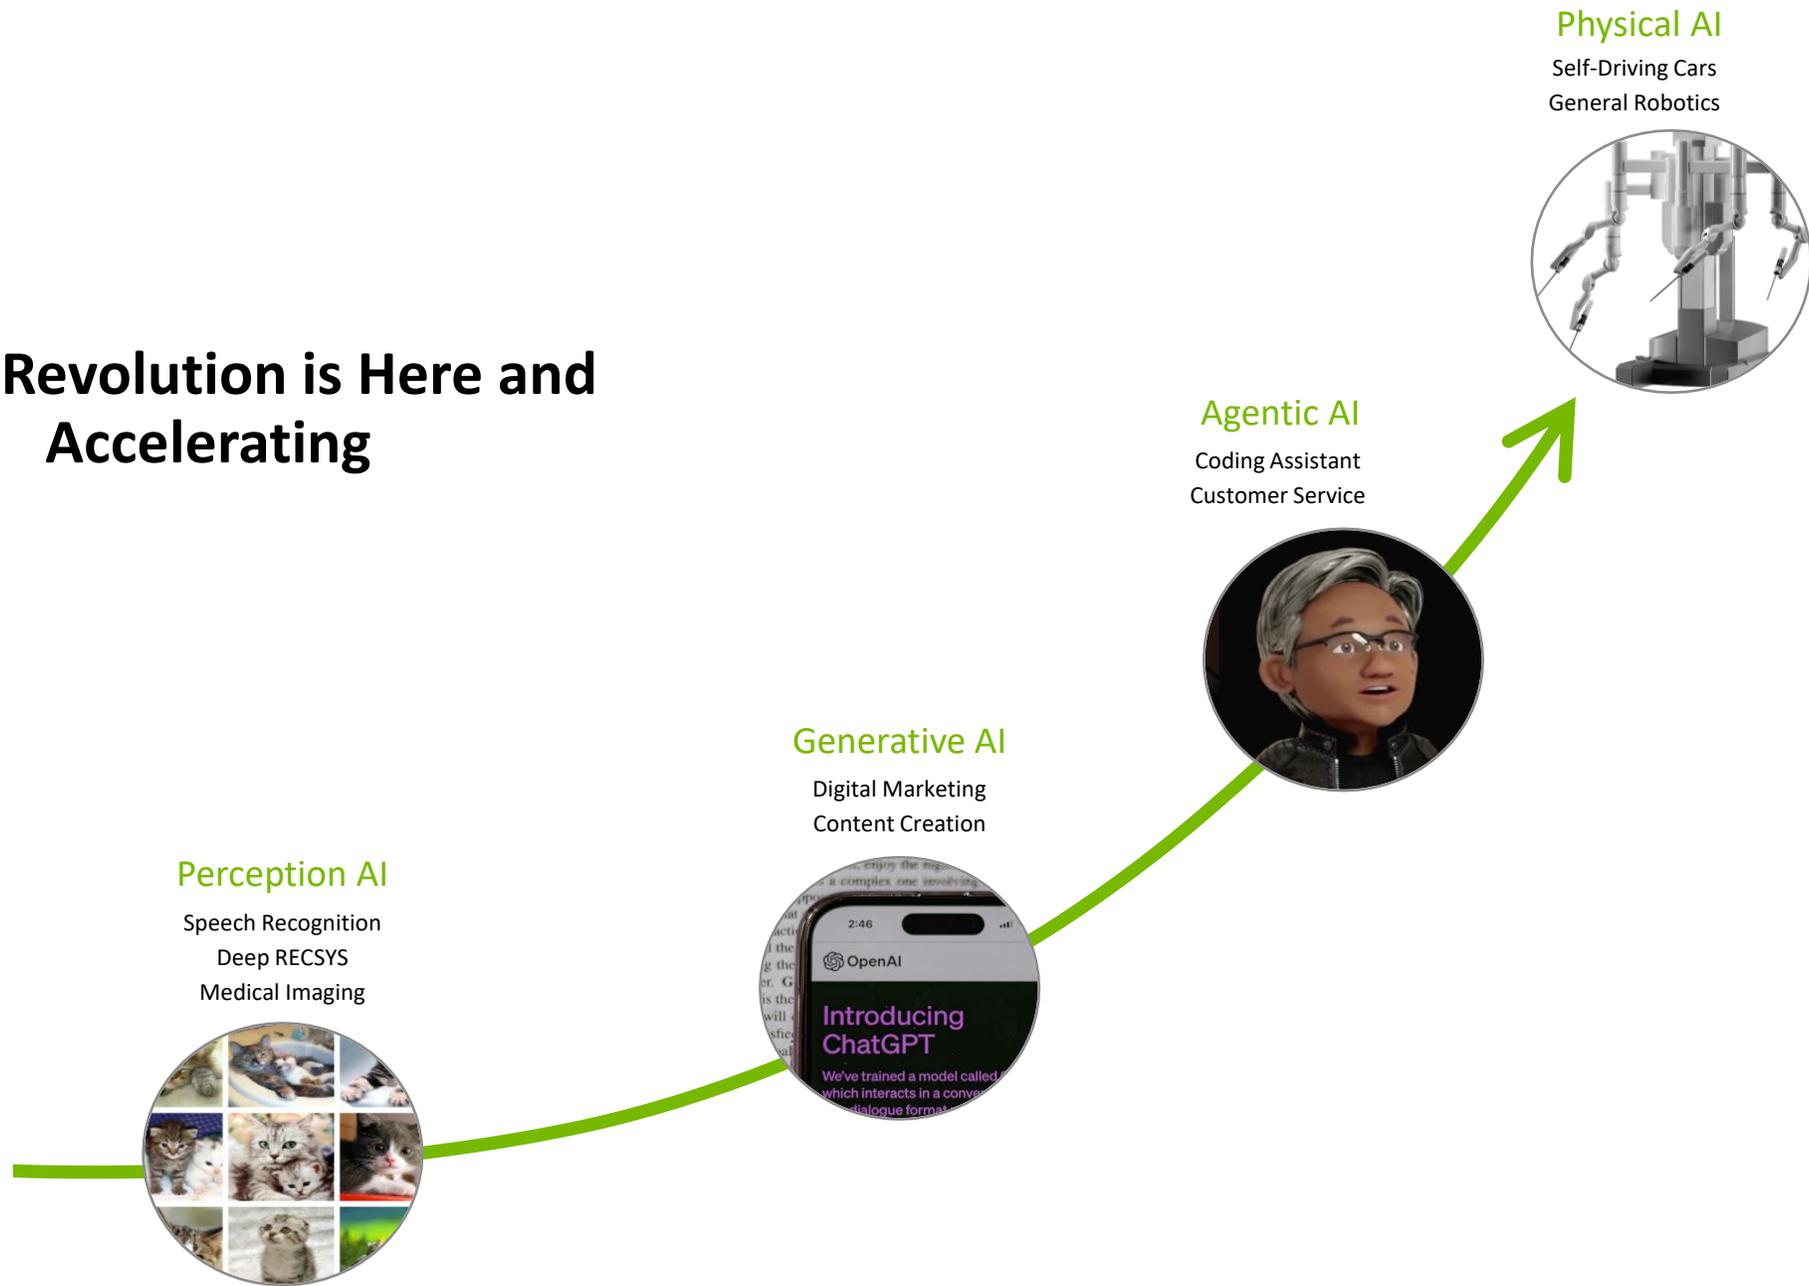
40,000+ organizations run NVIDIA AI Platforms

6M developers, 27K startups in ecosystem

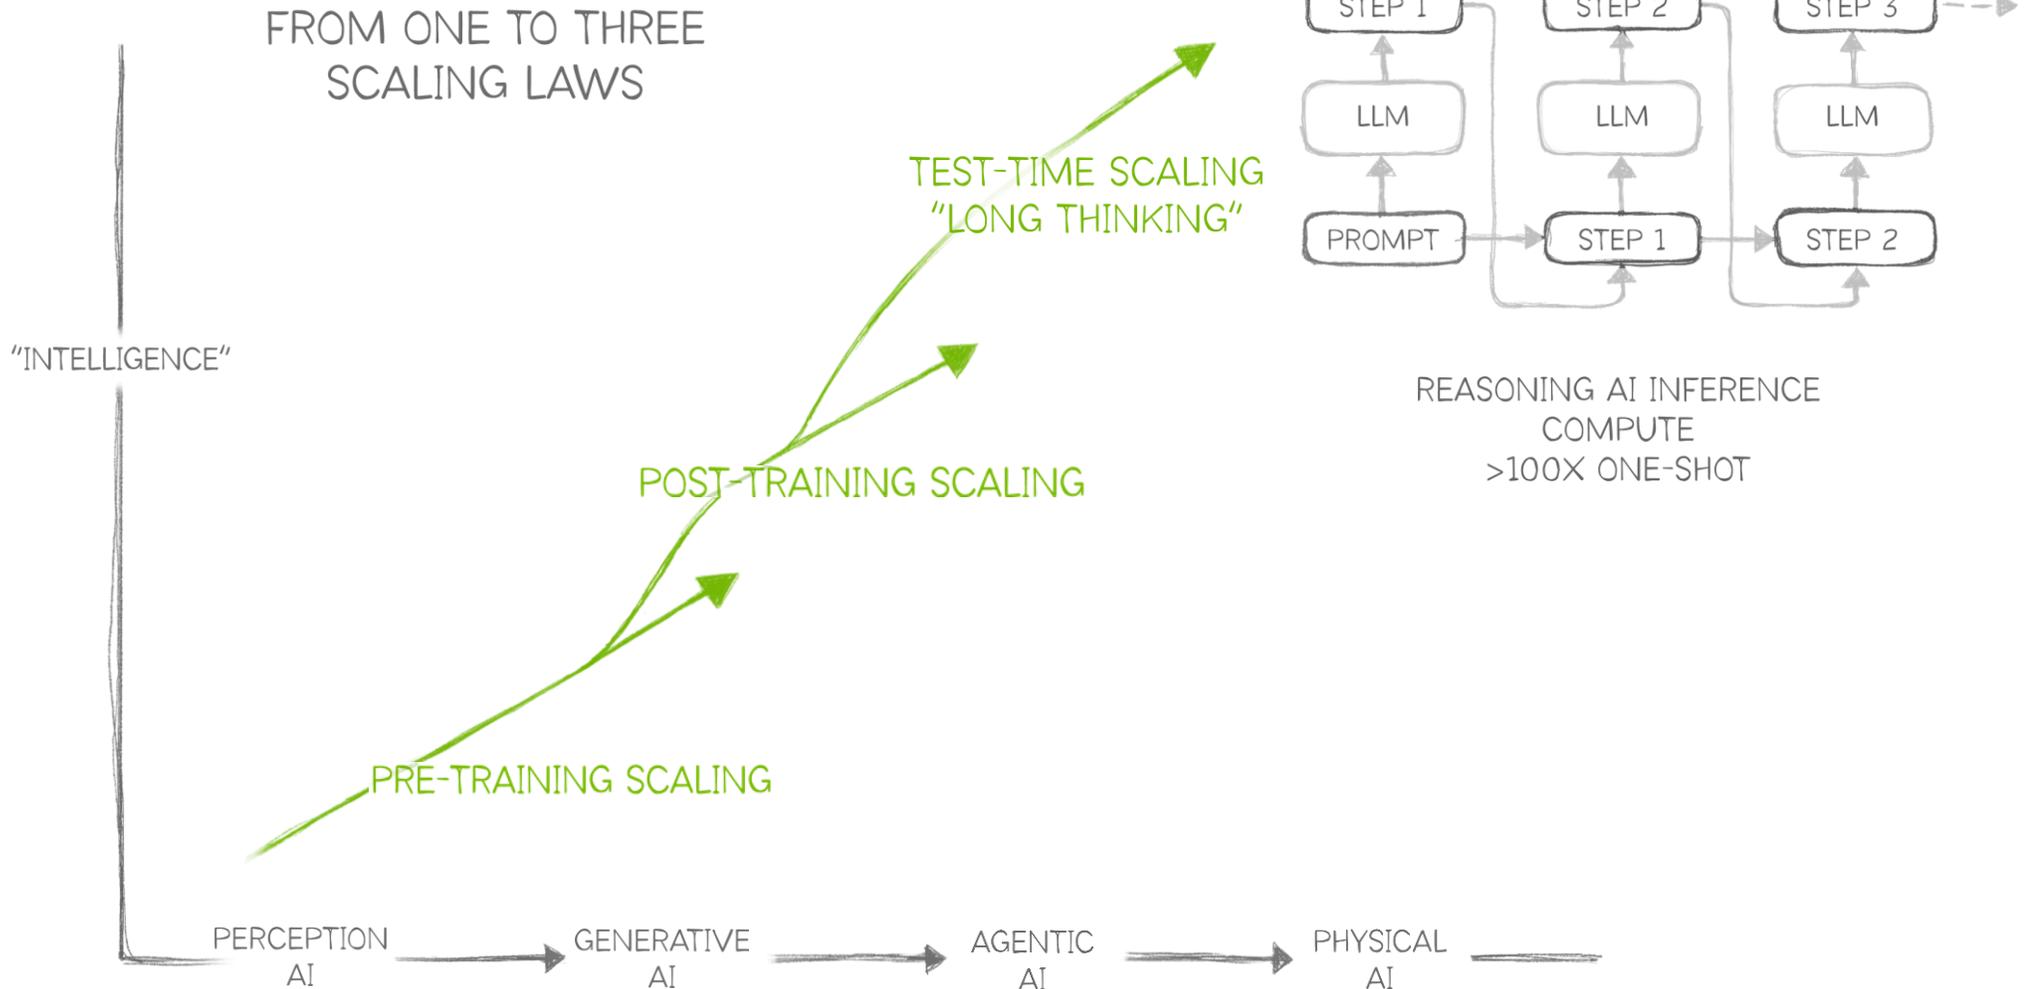
# AI Factories Transform Energy Data into Actionable Intelligence



# The AI Revolution is Here and Accelerating



# AI Scaling Laws Drive Exponential Demand for Compute



# Evolution of the Electric Grid

## From Centralized to Decentralized



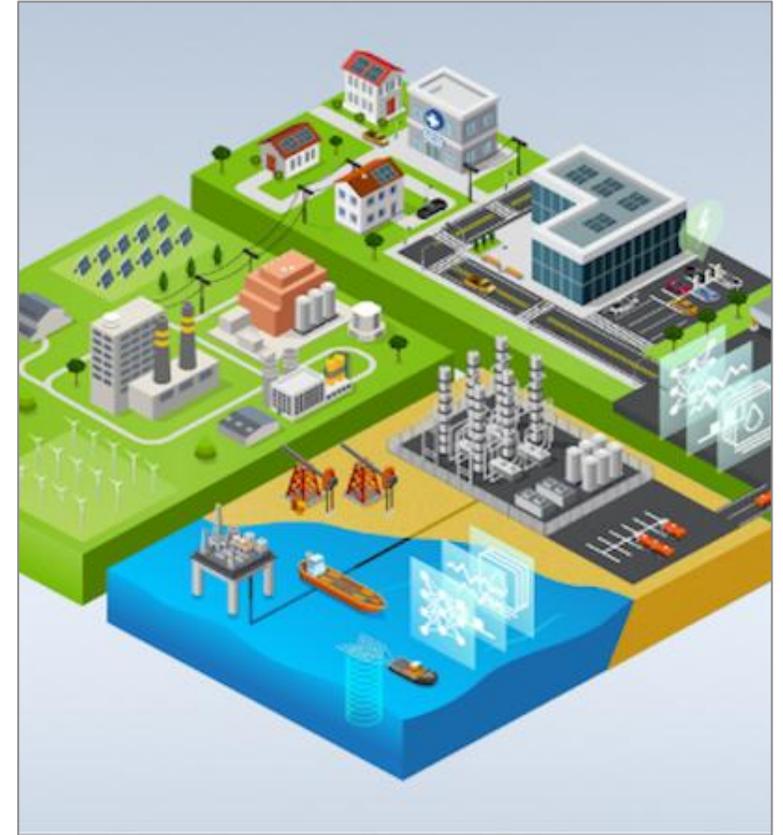
Traditional Grid ( Legacy Power System)

One-way power flow + Manual operations



Modern Grid ( Distributed and Digital)

Two-way power and data flow + Responsive infrastructure

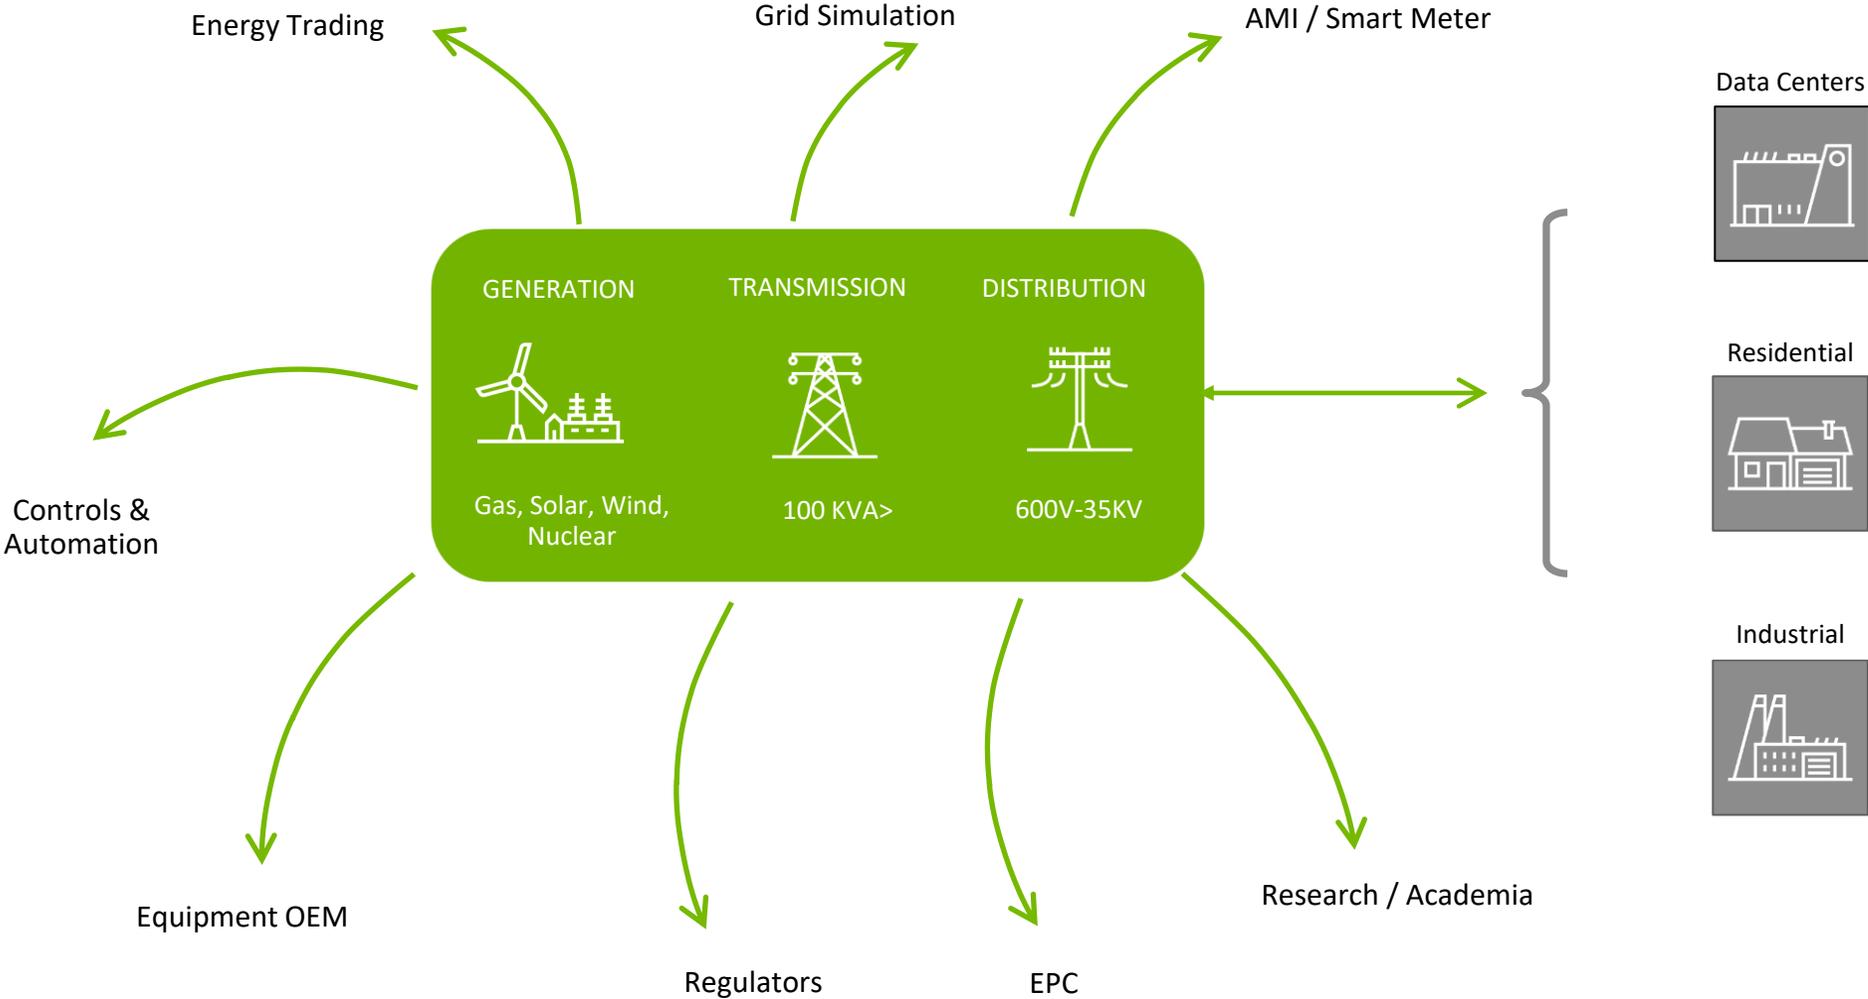


Autonomous Grid

Autonomous decision-making + AI-driven optimization

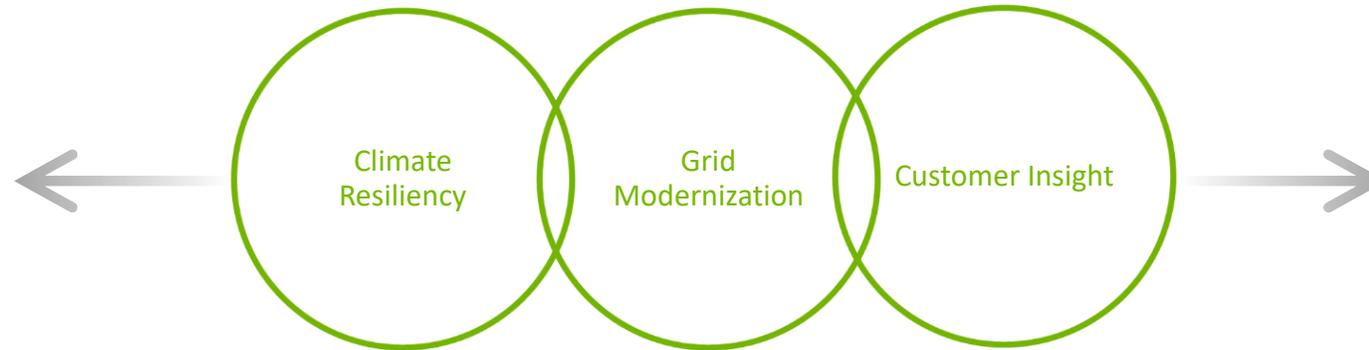
# NVIDIA Enabling Chip-2-Grid Ecosystem

Data Center as Flexible Resource



# How AI is Being Leveraged Across Power and Utilities

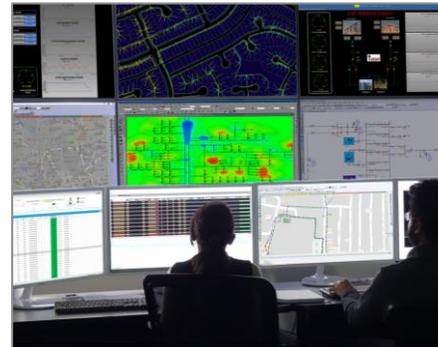
Becoming a software defined energy provider



Generation



Predictive Maintenance & Inspection



Load Planning and Optimization



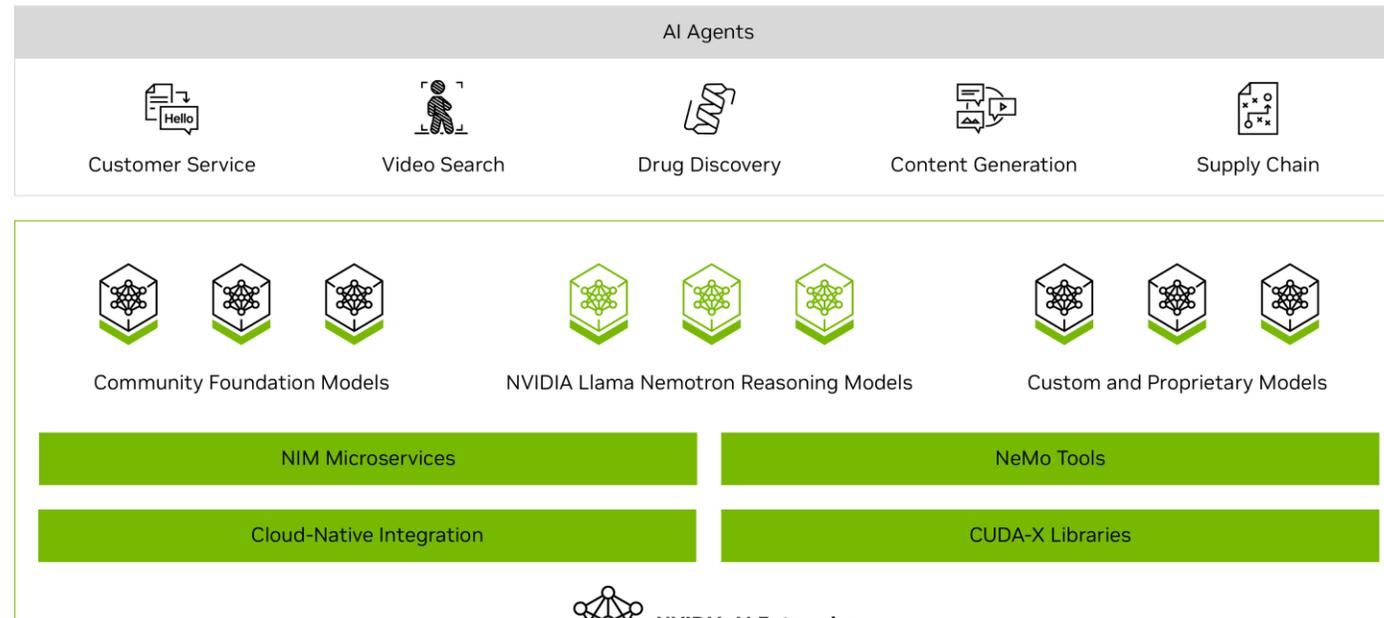
Grid Optimization



Cybersecurity

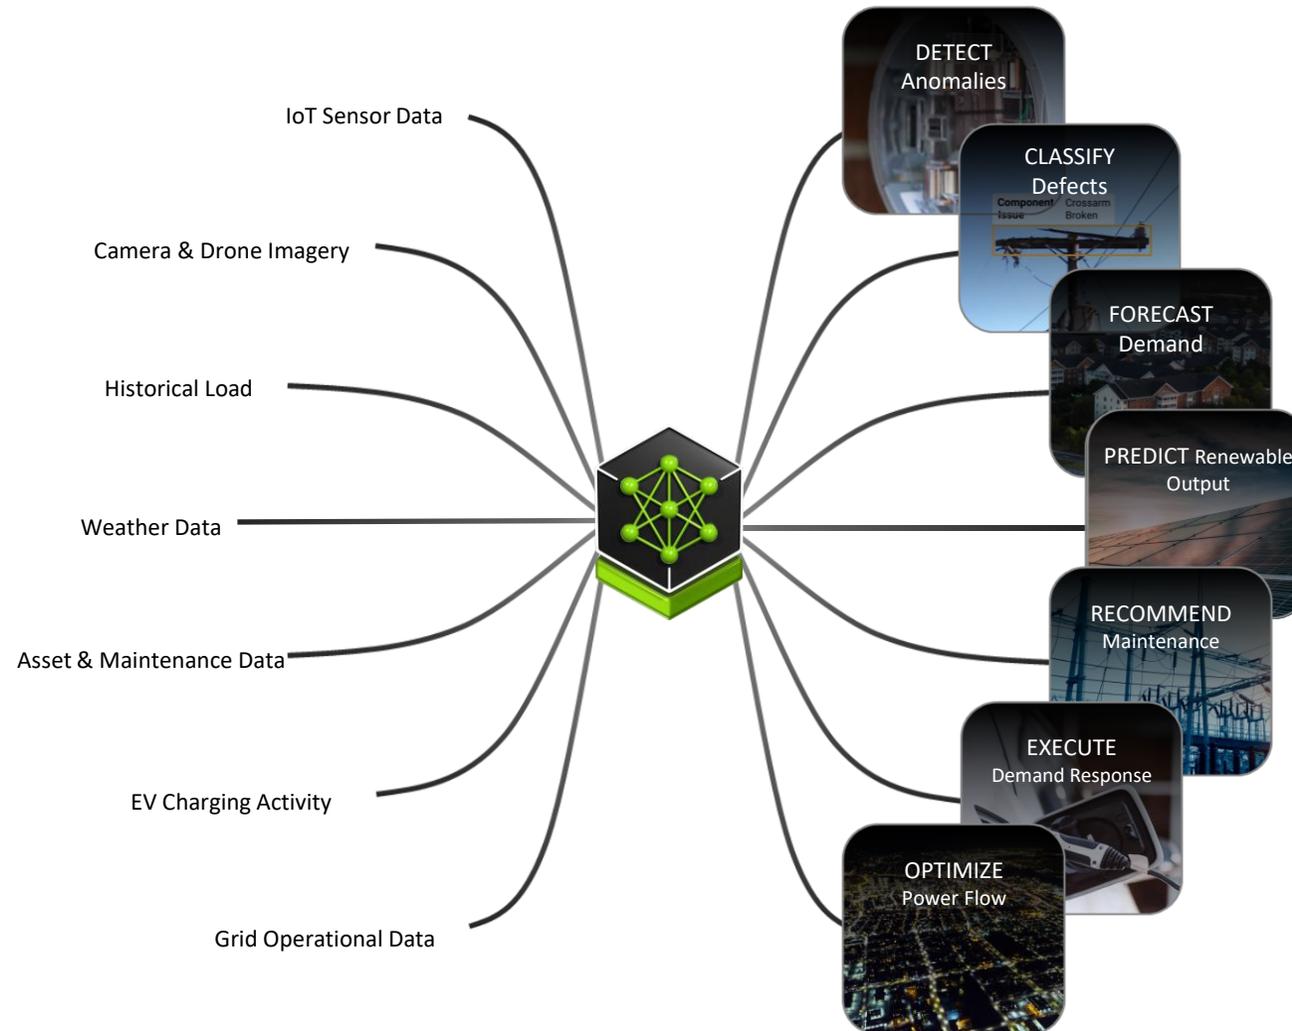
# AI-Powered Utilities Use Cases

## Powering Energy Transition



# The Power of Domain-Specific AI Models

Transform vast amounts of energy data into actionable intelligence



# Leading with Energy Ecosystem Partners

ISV-Accelerated  
with NV and Scale  
with Partner



- Co-innovating scalable AI solutions for grid modernization.

Light House  
Engagements



- Defining use cases jointly with utility stakeholders

Scale through  
Mutual  
Collaboration



- Tapping a global ecosystem to co-create AI transformation

40+ GTM Partners | Co-developed Industry Use Cases| Global Reach

# NVIDIA Accelerated Computing Platform



NVIDIA NIM & AI Microservices, NVIDIA Blueprints



## NVIDIA Application Frameworks



## NVIDIA Platform Software



NVIDIA HPC

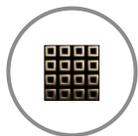


NVIDIA AI



NVIDIA Omniverse

## NVIDIA Accelerated Computing Infrastructure



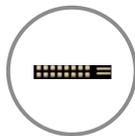
GPU



CPU



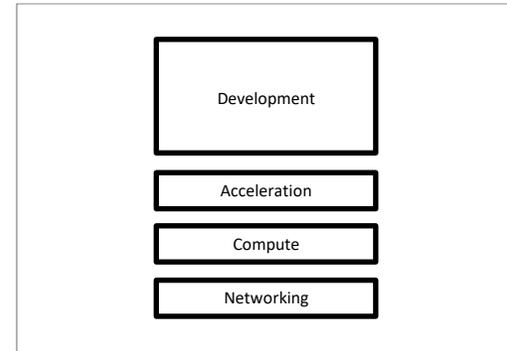
DPU



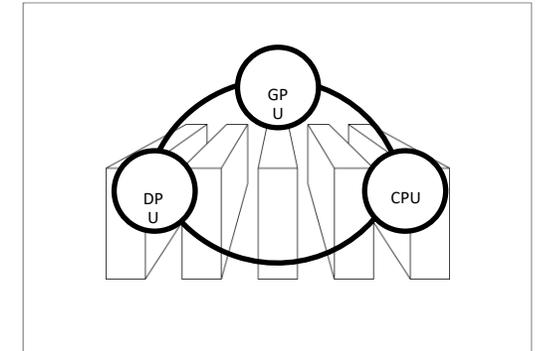
Networking

## 4 Core Platform Tenets

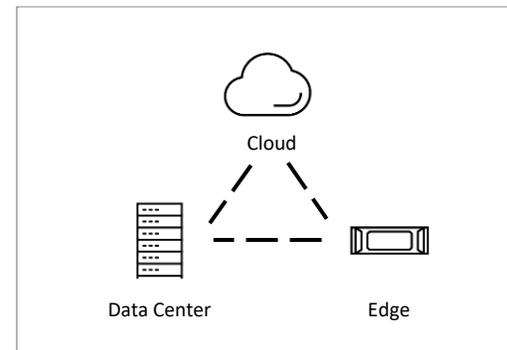
### Full Stack



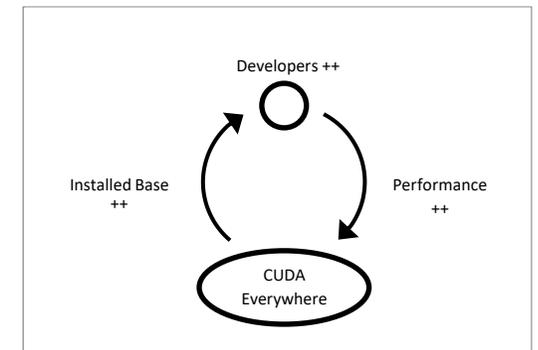
### Data Center Scale



### Cloud to Data Center to Edge



### One Architecture



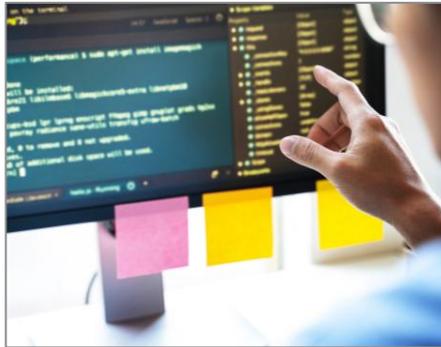
# AI Center of Excellence (CoE)

HR for AI onboarding



## Identified Use Cases

Business Need, Internal Alignment



## Quality Data

Understand Data Sources, Format



## AI Skills & Expertise

IT, Data Scientists, MLOps



## AI Infrastructure

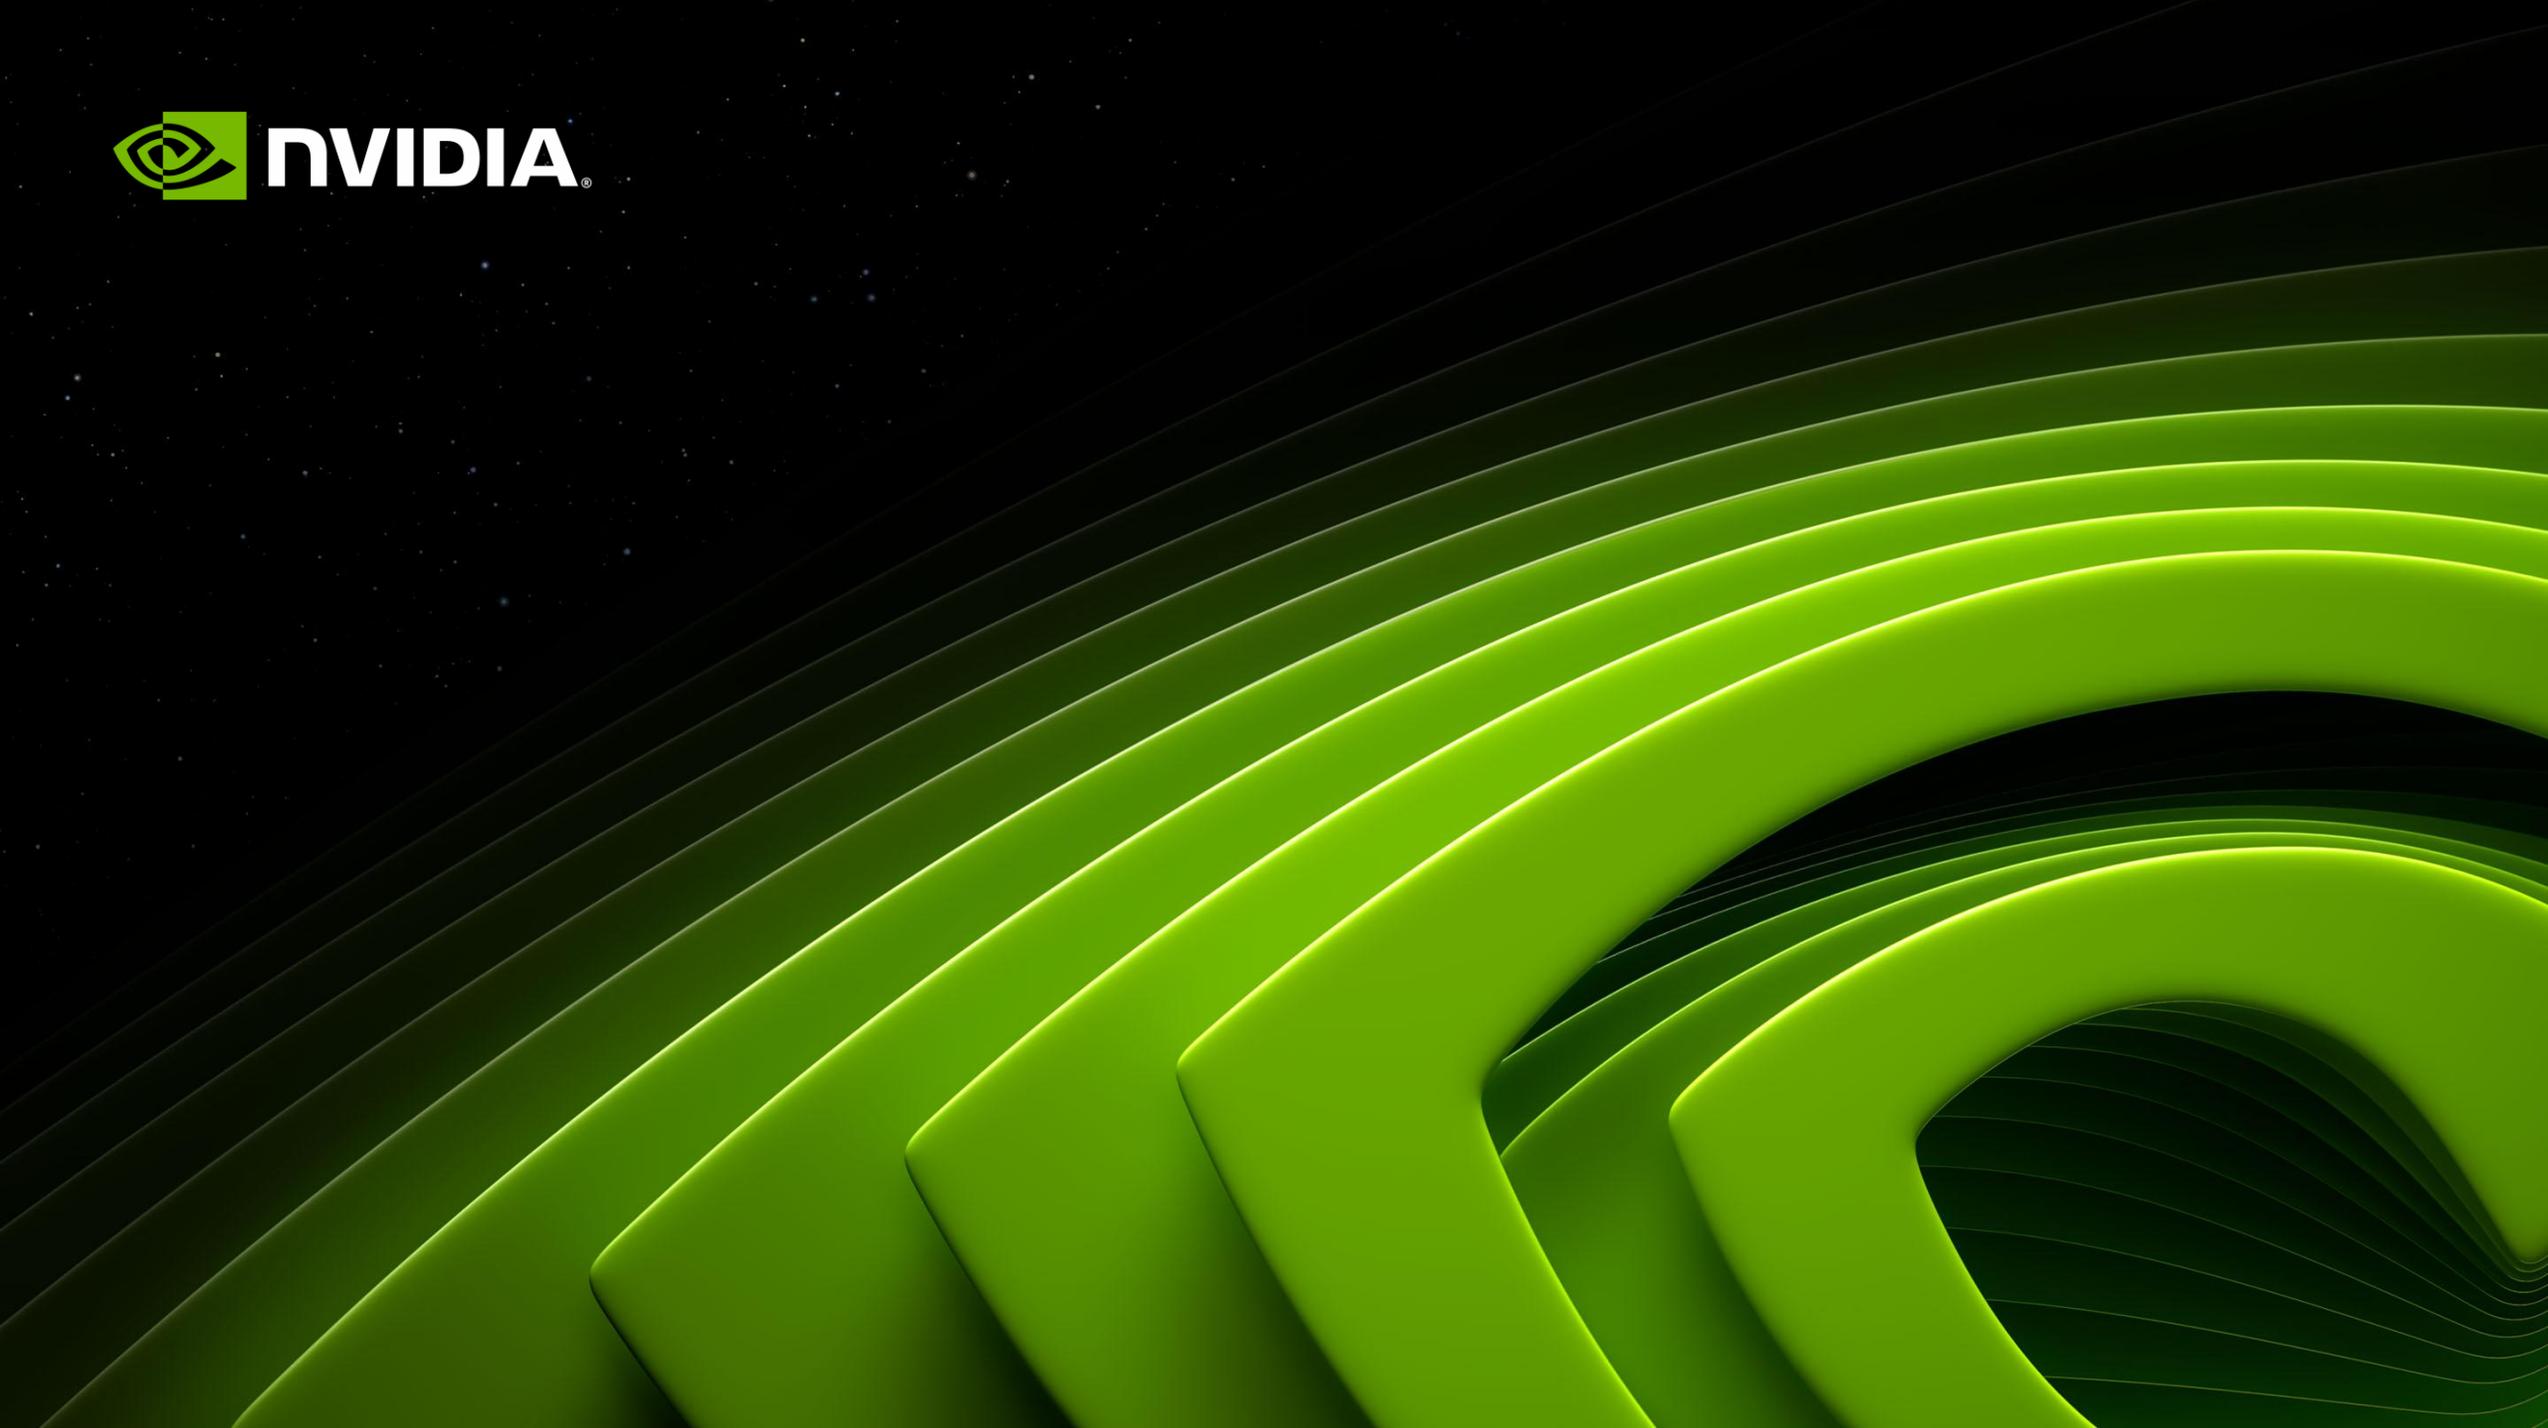
Consider Agility & Energy Efficiency



## AI Governance

Responsible AI, Data Privacy

Successful AI Deployments Start with an AI Strategy





DISTRIBUTECH 2026

**From Generative to Agentic:**

**Reimagine Grid Control Room of the Future**

**Xing Wang, PhD**

Global Leader Grid Modernization  
AWS Energy and Utilities



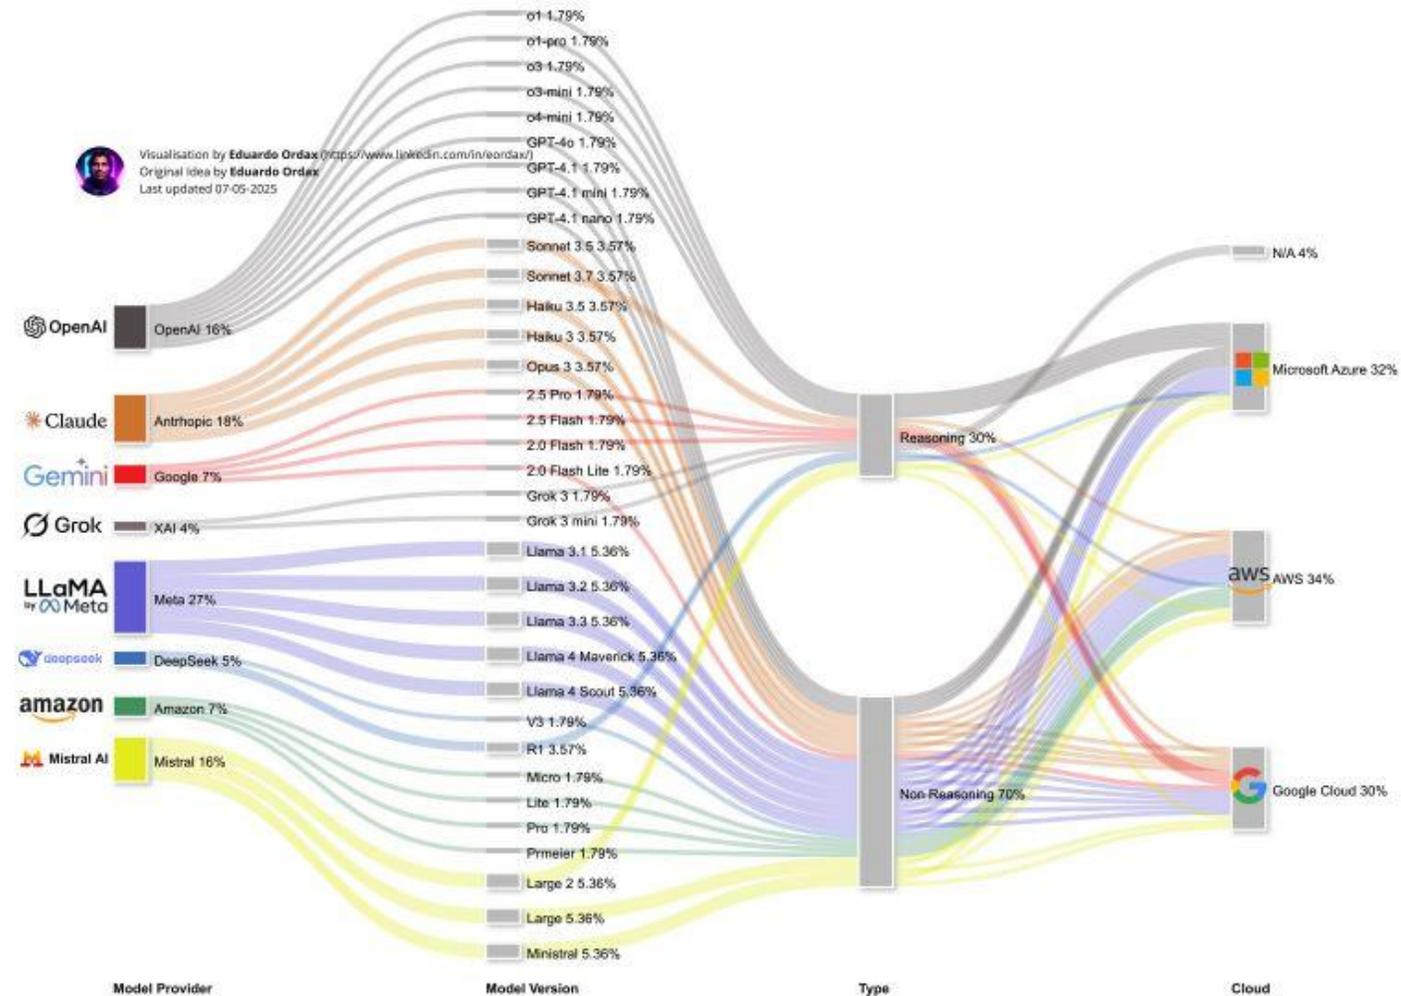
# The 2025 → 2026 Landscape of LLMs

- ◆ No more clear front-runner: We've gone from "everyone chasing one leader" to a fairly even playing field. For most use cases, model differences are small and often not that relevant.

- ◆ Model choice is the new normal: Customers now expect the ability to test, compare, and switch between models with ease. This shift is driving interest in evaluation frameworks and model routing tools.

- ◆ Reasoning-first models are rising: Many providers are clearly moving toward models optimized for reasoning—fueling the surge of Agentic AI architectures.

- ◆ Proprietary still leads, but just barely: Open-source and open-weight models are quickly closing the gap.



# Agentic AI will help realize the true promise of LLMs

## Gen AI



Primary purpose: Content generation

Dependency: Prompts and input data

Agency level: None (requires more human prompting)

Ability to execute workflows: None / low

## Agentic AI

Primary purpose: Goal achievement

Dependency: Access to tools, data and agents

Agency level: Higher (less to no human oversight)

Ability to execute workflows: Medium / high

# AWS's Investment in Neuro-Symbolic AI

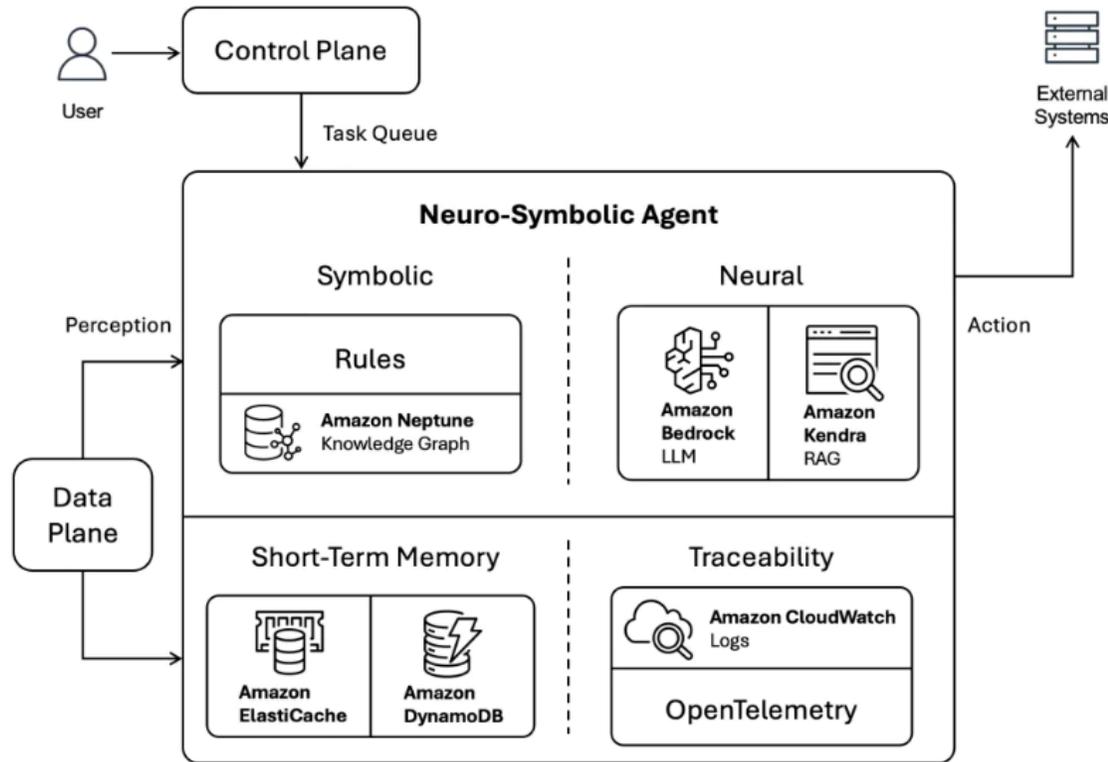


Figure 3: Neuro-symbolic agent on AWS



**Matt Garman** · Following  
CEO at Amazon Web Services (AWS)  
1mo ·

Scientists across [Amazon](#) are revolutionizing how we verify AI systems with automated reasoning. Our teams are making AI more trustworthy, applying the same mathematical logic we've used for years to verify security at [Amazon Web Services \(AWS\)](#).

It's called "neurosymbolic AI", and you can see this in action with robots like the one below navigating Amazon warehouses more efficiently, Rufus providing shoppers with more accurate answers, and AWS teams using this approach for everything from verifying AI-generated code correctness before shipping new Kiro features publicly, to Graviton chip optimization.

Last week's launch of Automated Reasoning checks in Amazon Bedrock Guardrails—catching AI errors with up to 99% accuracy—is another proof point in how we're making AI safer for everyone. Learn more in the [The Wall Street Journal](#) about this game-changing verification technology. <https://lnkd.in/g9H-hXnS>

- Understand language using neural networks.
- Ground understanding in structured knowledge bases.
- Execute logic using explicit rules and policy frameworks.
- Explain decisions through symbolic traces while adapting to new inputs through learned representations.



# AWS AI portfolio

## APPLICATIONS AND AGENTS

- Kiro
- Amazon Quick Suite
- AWS Transform
- Amazon Connect
- AWS Marketplace

## FRONTIER AGENTS

- Kiro autonomous agent **New**
- AWS DevOps Agent **New**
- AWS Security Agent **New**

## AI & AGENT DEVELOPMENT SOFTWARE & SERVICES

### AGENT BUILDERS

#### Vertically Integrated

- Amazon Nova Act **New**

#### Flexible/OSS

- Strands Agents

### AMAZON BEDROCK

#### Models

- Amazon Nova
- 3P Models

#### Capabilities

- Optimization
- Guardrails
- Customization
- Knowledge Bases

#### AgentCore

- Runtime
- Gateway
- Memory
- 1P Tools
- Identity
- Observability
- Policy **New**
- Evaluations **New**

## INFRASTRUCTURE

### AMAZON SAGEMAKER

#### Custom Model Building

- Model building
- Model training
- Deployment
- MLOps
- Hyperpod
- Nova Forge **New**

#### Data Foundation for AI

- Data Processing
- Governance
- Storage & Databases
- Vector Search

### AI COMPUTE

- Trainium
- Inferentia
- GPUs

INTERFACES & PROTOCOLS (MCP/A2A)

SECURITY AND POLICIES



# AWS Agents4Energy (In Progress)



**Amazon Bedrock**  
Access to FMs to AI infrastru



stability.ai



Reasoning-first models are rising: Many providers are clearly moving toward models optimized for reasoning—fueling the surge of Agentic AI architectures.



# Gartner Predicts AI Adoption in 40% of Power and Utilities Control Rooms by 2027

STAMFORD, Conn., January 15, 2025

## 94% of Power and Utility CIOs Plan to Increase Their AI Investments in 2025, With an Average Spending Increase of 38.3%

By 2027, 40% of power and utilities will deploy AI-driven operators in control rooms, reducing human error risks, but increasing cyber-physical system security vulnerabilities, according to Gartner, Inc.

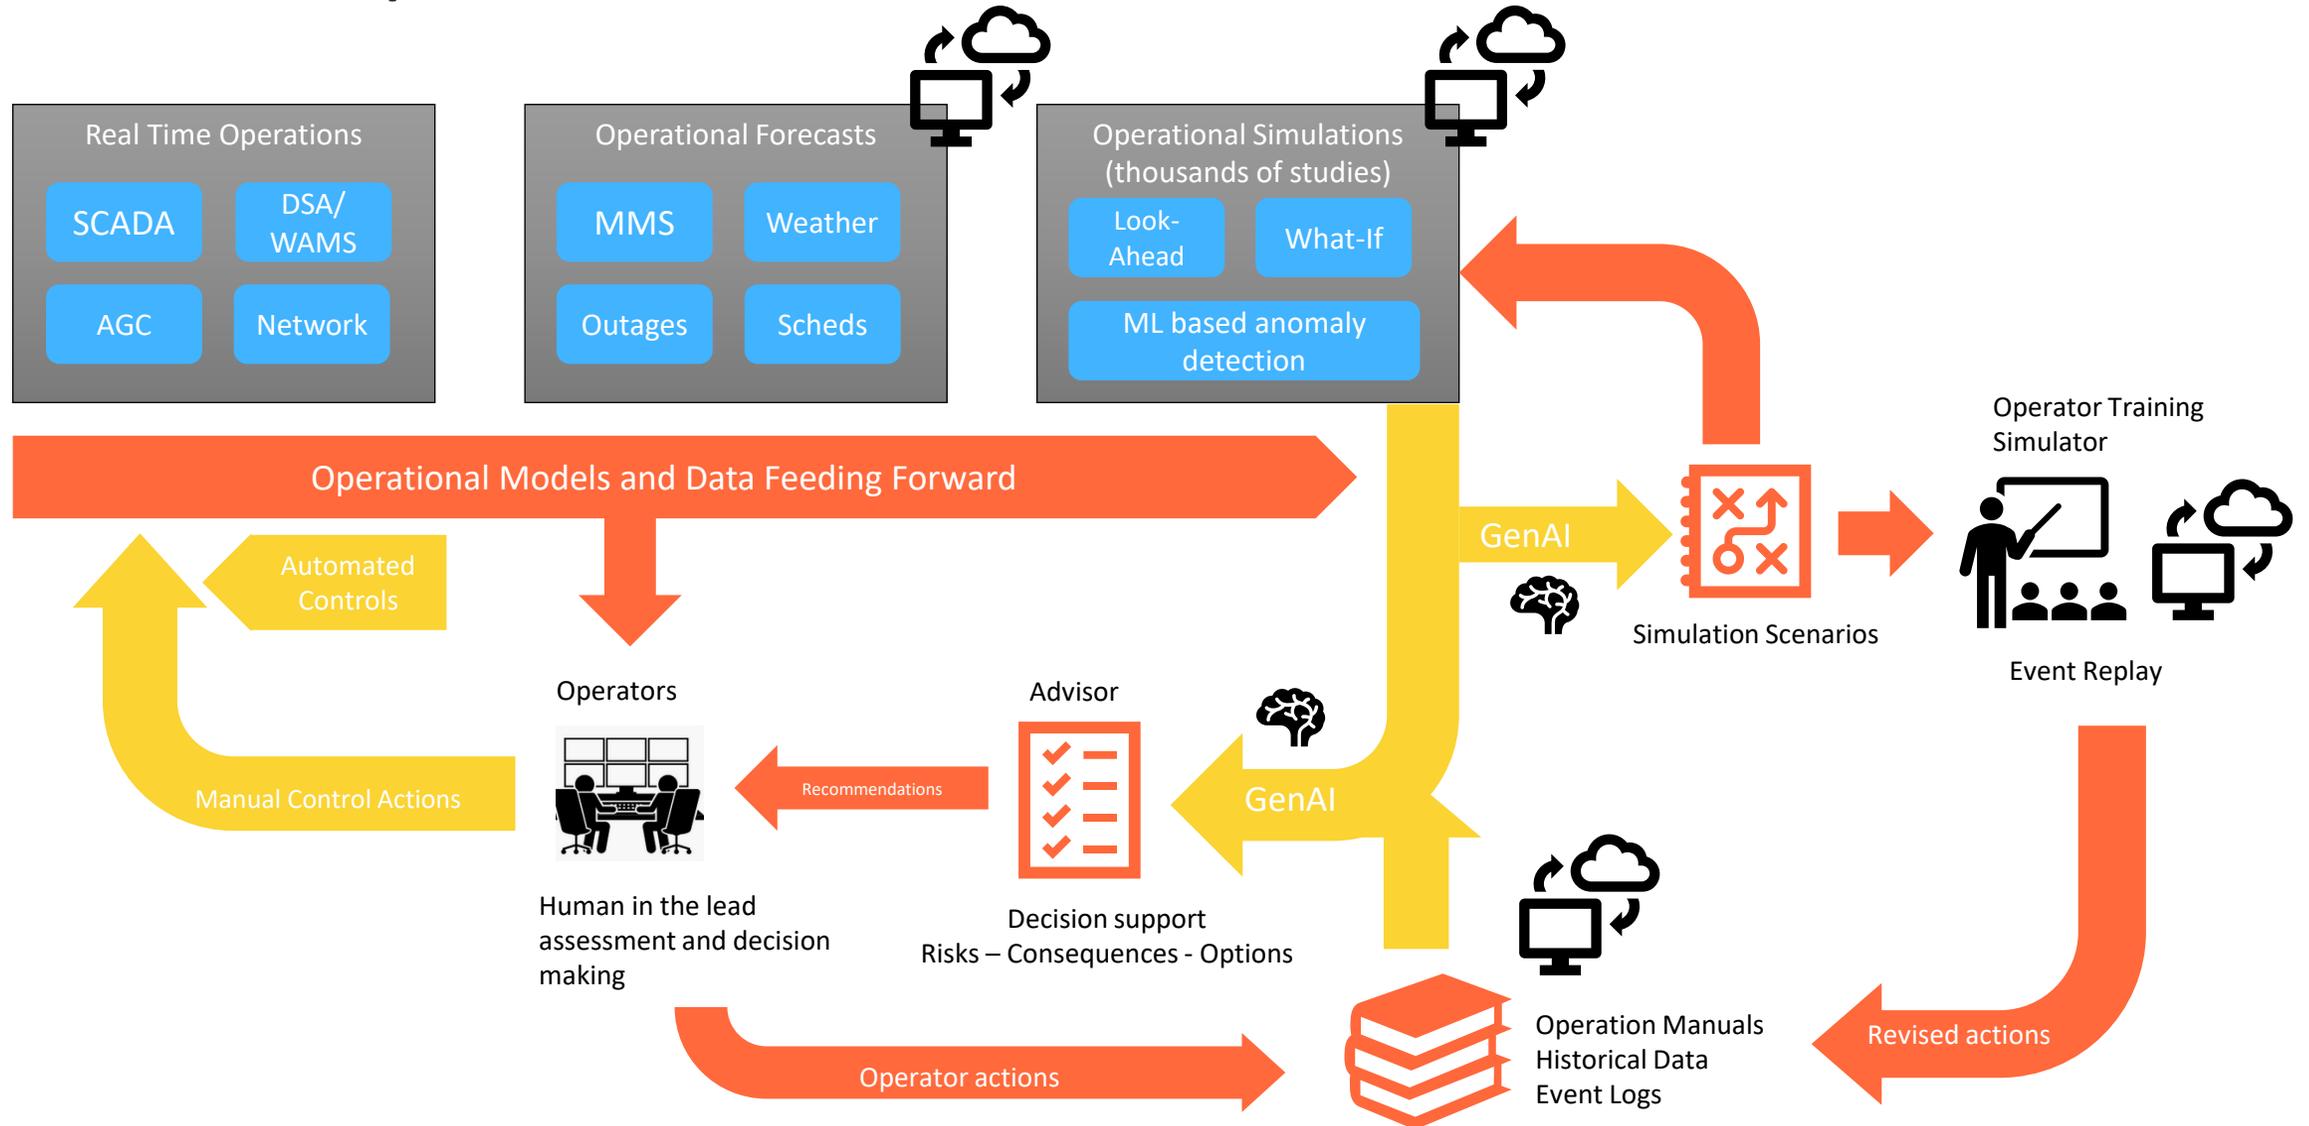
The 2025 [Gartner CIO and Technology Executive Survey](#) indicated that 94% of power and utility chief information officers (CIOs) plan to increase their AI investments in 2025, with an average spending increase of 38.3%.

“AI technology is poised to transform the power and utilities sector,” said [Jo-Ann Clynych](#), Sr Director Analyst at Gartner. “Human [decision-making](#) is critical, but it is also a significant factor in industrial accidents. AI-driven operations offer a compelling solution, performing tasks with repeatability, precision, and without bias when effectively governed.”

<https://www.gartner.com/en/newsroom/press-releases/2025-01-15-gartner-predicts-ai-adoption-in-40-percent-of-power-and-utilities-control-rooms-by-2027>



# GenAI Co-Operator for Grid Control Room

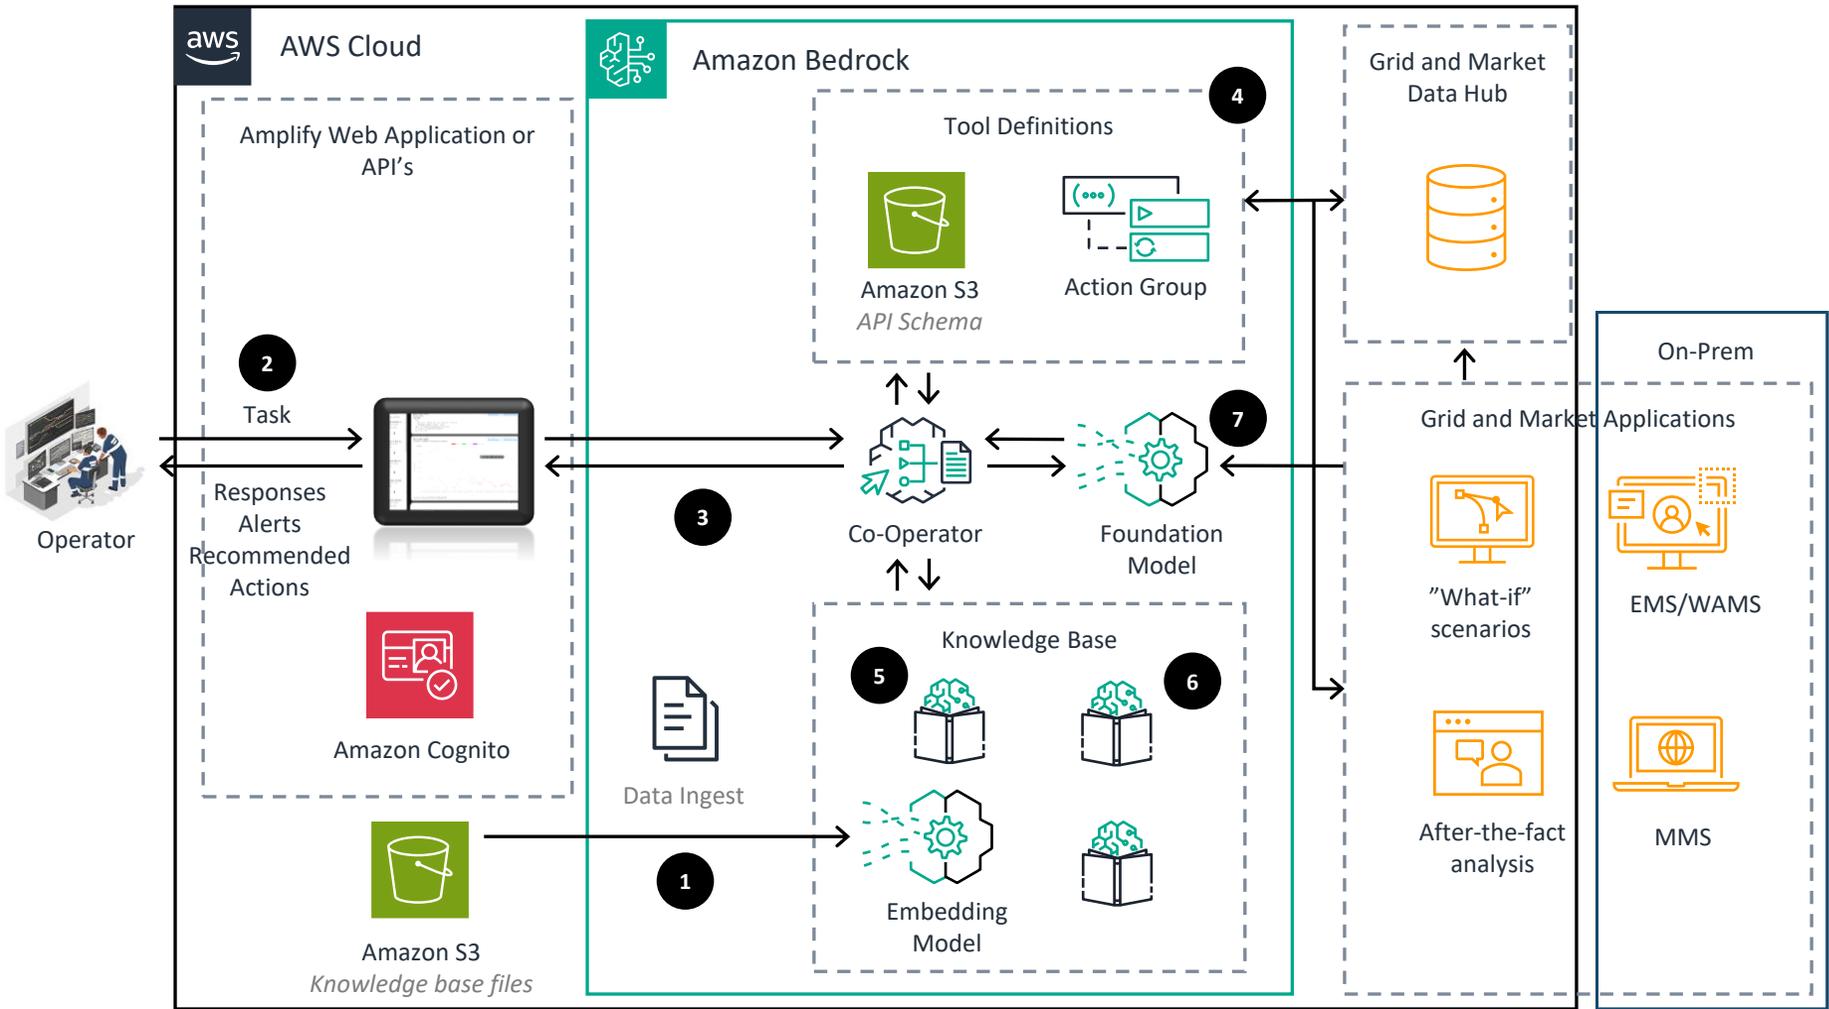


# Multiple Agents Interacting with Operators

- **Data integration agent** integrates data layers from SCADA, PMU, IoT, GIS, Weather, etc. to create a comprehensive grid model and situation awareness;
- **Simulation agent** conducts large number of “what-if” grid simulations and identify system anomalies, leveraging state-of-the-art grid analytical tools powered by Cloud HPC;
- **Intelligent alarm agent** processes alarms from various systems and identify the root cause.
- **Dispatch agent** produces dispatch recommendations based on the simulation agent’s results and trustworthy knowledge base (grid operational protocols, grid common information model (CIM), and historical operator event logs, etc.);
- **Feasibility check agent** performs comprehensive system impact test to ensure the feasibility and safety of recommended operator actions;
- **After-fact “perfect dispatch” agent** simulates the best possible dispatch actions assuming operators had “perfect” prediction of the system. “Perfect dispatch” agent provides improvement feedback for future dispatch recommendations.
- **Supervisor agent** coordinates, validates, and approves the entire real-time system dispatch process.

# Guidance for a Grid Co-Operator on AWS

This architecture diagram demonstrates how grid operators can build AI agents to provide decision support to operators in grid control room . Through an integrated chat interface or triggered events, customers will get truly intelligent system alerts and recommended remedial actions.



1

Documents such as Operation manuals can be ingested from S3, SharePoint, etc. into an Amazon Bedrock Knowledge Base (KB) that contains vectorized embeddings serverlessly stored on Amazon Aurora or Amazon OpenSearch.

2

Users interact directly with the agents through an Amplify web application secured by Amazon Cognito.

3

Amazon Bedrock Agents plan and execute complex workflows to solve natural language provided requests. Chain-of-thought reasoning provides human-like workflow processing.

4

Action Groups enable agents to take action against existing grid/market data and applications.

5

Agents can contribute to their own knowledge base to continue learning from new interactions..

6

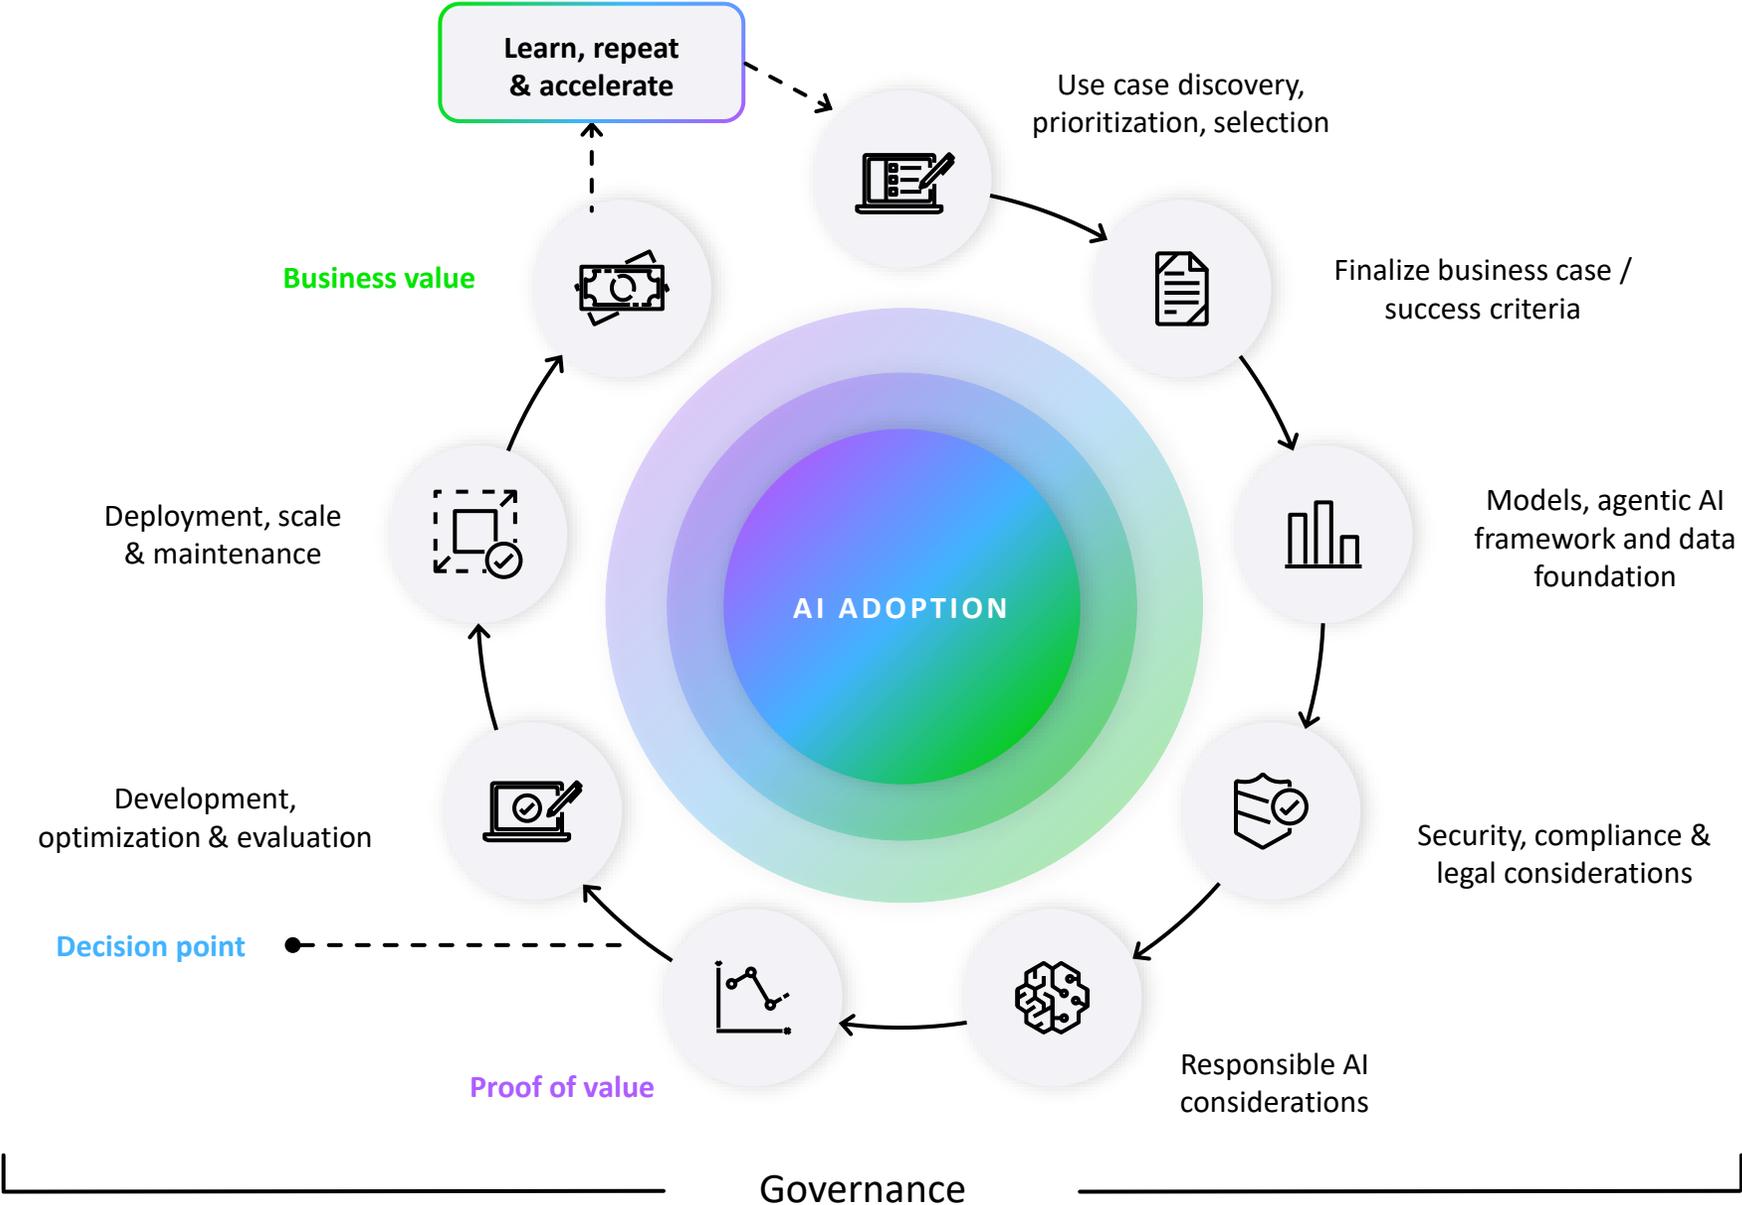
Periodic sync operation will vectorize new data into the knowledge base to ensure agents have up-to-date information.

7

Make the LLM "physics informed"



# Start now



# Thank you!



# Next Generation Asset Management



# The Case for Change: Traditional Inspections Fall Short

## The Inherent Flaws of Manual Methods

Traditional power grid inspections are fundamentally limited and pose significant risks to a modern reliability program.



**Subjectivity and Inconsistency:** Manual methods lack scientific supervision and objective evaluation metrics. The quality of an inspection often depends on an individual inspector's sense of responsibility and experience, which can vary dramatically.



**Poor Data Management:** Inspection data is typically recorded on paper, making collection, organization, and retrieval incredibly difficult and inefficient. This hinders the development of data-driven reliability methods.

# The Path Forward: A Modern Reliability Program Must Be...



**Data-Driven:** Leveraging digital data for trend analysis and predictive insights.



**Consistent:** Applying the same rigorous standards to every inspection.



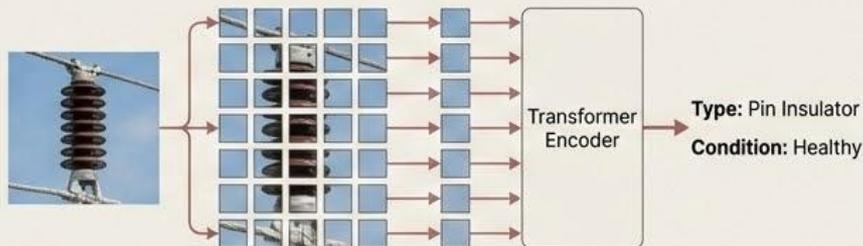
**Scalable:** Efficiently covering a vast and growing infrastructure network.

# Vision Transformers (ViT) for Asset Inspection and Classification

The Classifier: Seeing the Bigger Picture in Every Image

ViT models represent a paradigm shift, analyzing images as a sequence of patches to understand global context and complex relationships, not just pixels.

## How It Works



### Application in Asset Inspection:

Exceptionally good at image classification. Accurately identifies asset type and classifies its overall condition.



### Key Advantage:

Achieves high performance, often with less training data than traditional CNNs, making them highly adaptable.

# Masked Autoencoders (MAE) for Anomaly Detection

## Learning What's Normal to Find What's Not

One of the biggest hurdles in AI is the need for massive, hand-labeled datasets. Masked Autoencoders (MAE) offer a powerful solution through self-supervised learning.



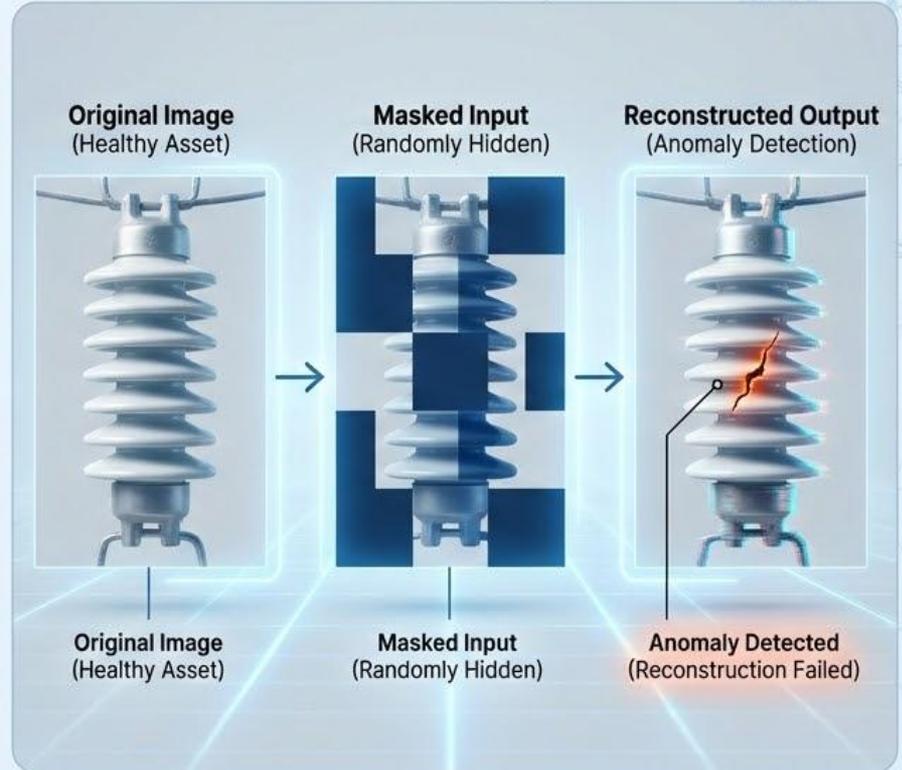
**How It Works:** The MAE model takes an image, randomly hides a large portion of it, and then tries to reconstruct the original. This teaches the model a deep understanding of what normal, healthy assets look like.



**Application in Asset Inspection:** MAEs are perfect for pre-training models on vast archives of unlabeled footage. The pre-trained model becomes an excellent anomaly detector, flagging any defect that it struggles to reconstruct.



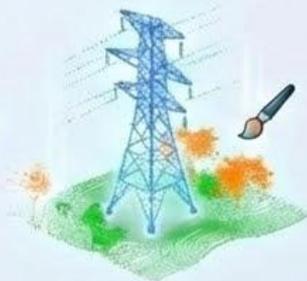
**Key Advantage:** This drastically reduces the time and cost of manual data labeling and leverages the full value of all collected visual data, not just the labeled portions.



# 3D Semantic Segmentation for Spatial Precision

## Building an Intelligent 3D Map of the Grid

While 2D images are useful, they lack depth. For true spatial understanding, we turn to **3D Semantic Segmentation**, which provides an unparalleled level of detail from LiDAR data.



**95%**   
**Accuracy**

 Transformer-based  
segmentation model



### How It Works:

Semantic Segmentation assigns a classification label to every single point in a 3D point cloud, effectively 'painting' a detailed 3D digital twin of the environment.



### Application in Asset Inspection:

This is indispensable for tasks requiring precise spatial measurement. The most critical use case is vegetation management, calculating the exact distance between power lines and foliage.



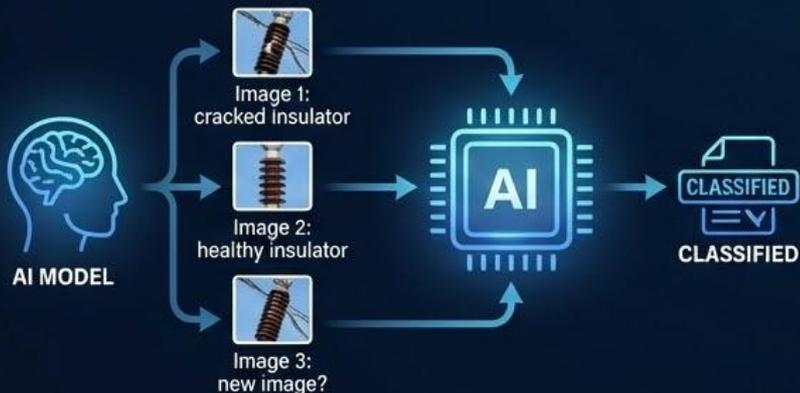
### Key Advantage:

Transformer-based segmentation models achieve over 95% accuracy in identifying grid components, providing a highly reliable and automated way to manage spatial risks.

# Advanced Techniques for Complex Analysis

Beyond individual models, the next wave of AI brings sophisticated capabilities for deeper, contextual insights.

## Few-Shot Prompting of Frontier Models (e.g., Gemini)



**The Concept:** Leverage massive, pre-trained models by providing just a few examples ("shots") in a natural language prompt to perform a new visual inspection task instantly.

**Example Prompt:** "You are an expert grid inspector. The following is a cracked insulator: [Image 1]. The following is a healthy insulator: [Image 2]. Now, classify this new image: [Image 3]."

## The Power of Multimodality



**Holistic Analysis:** A multimodal model can analyze an image of a "hot spot," read the corresponding maintenance log (text), and process its real-time temperature data. By fusing this information, it provides a far more comprehensive assessment.

# Scoring for Reliability: Why the F2 Score is Critical

## Precision vs. Recall: A Critical Trade-off



### Precision

When the model flags a defect, how often is it right?  
(Minimizes false alarms)



### Recall

Of all the real defects out there, how many did the model find?  
(Minimizes missed problems)

## The F2 Score: Prioritizing Safety and Prevention



The F2 Score weights **Recall** twice as heavily as Precision. It is designed for situations where the cost of missing a real problem is far greater than the cost of a false alarm.



### False Negative (The Catastrophe)

A critical crack is missed, leading to failure. This is what we must avoid at all costs.

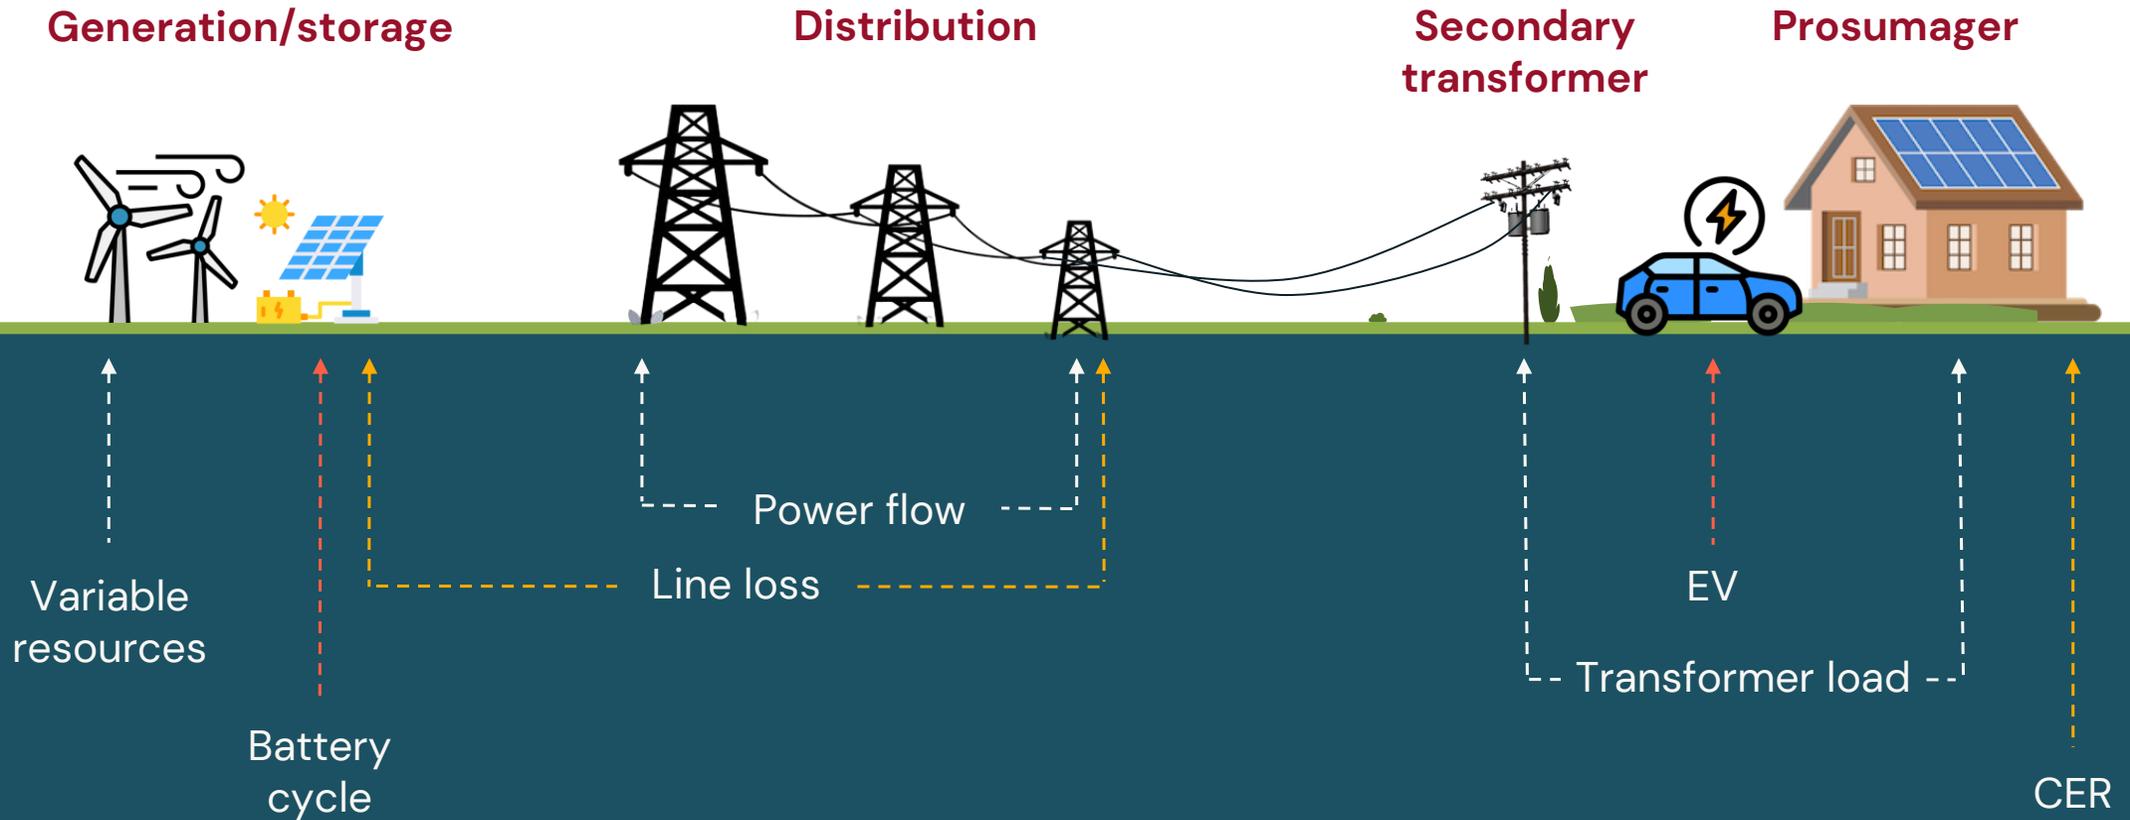


### False Positive (The Inconvenience)

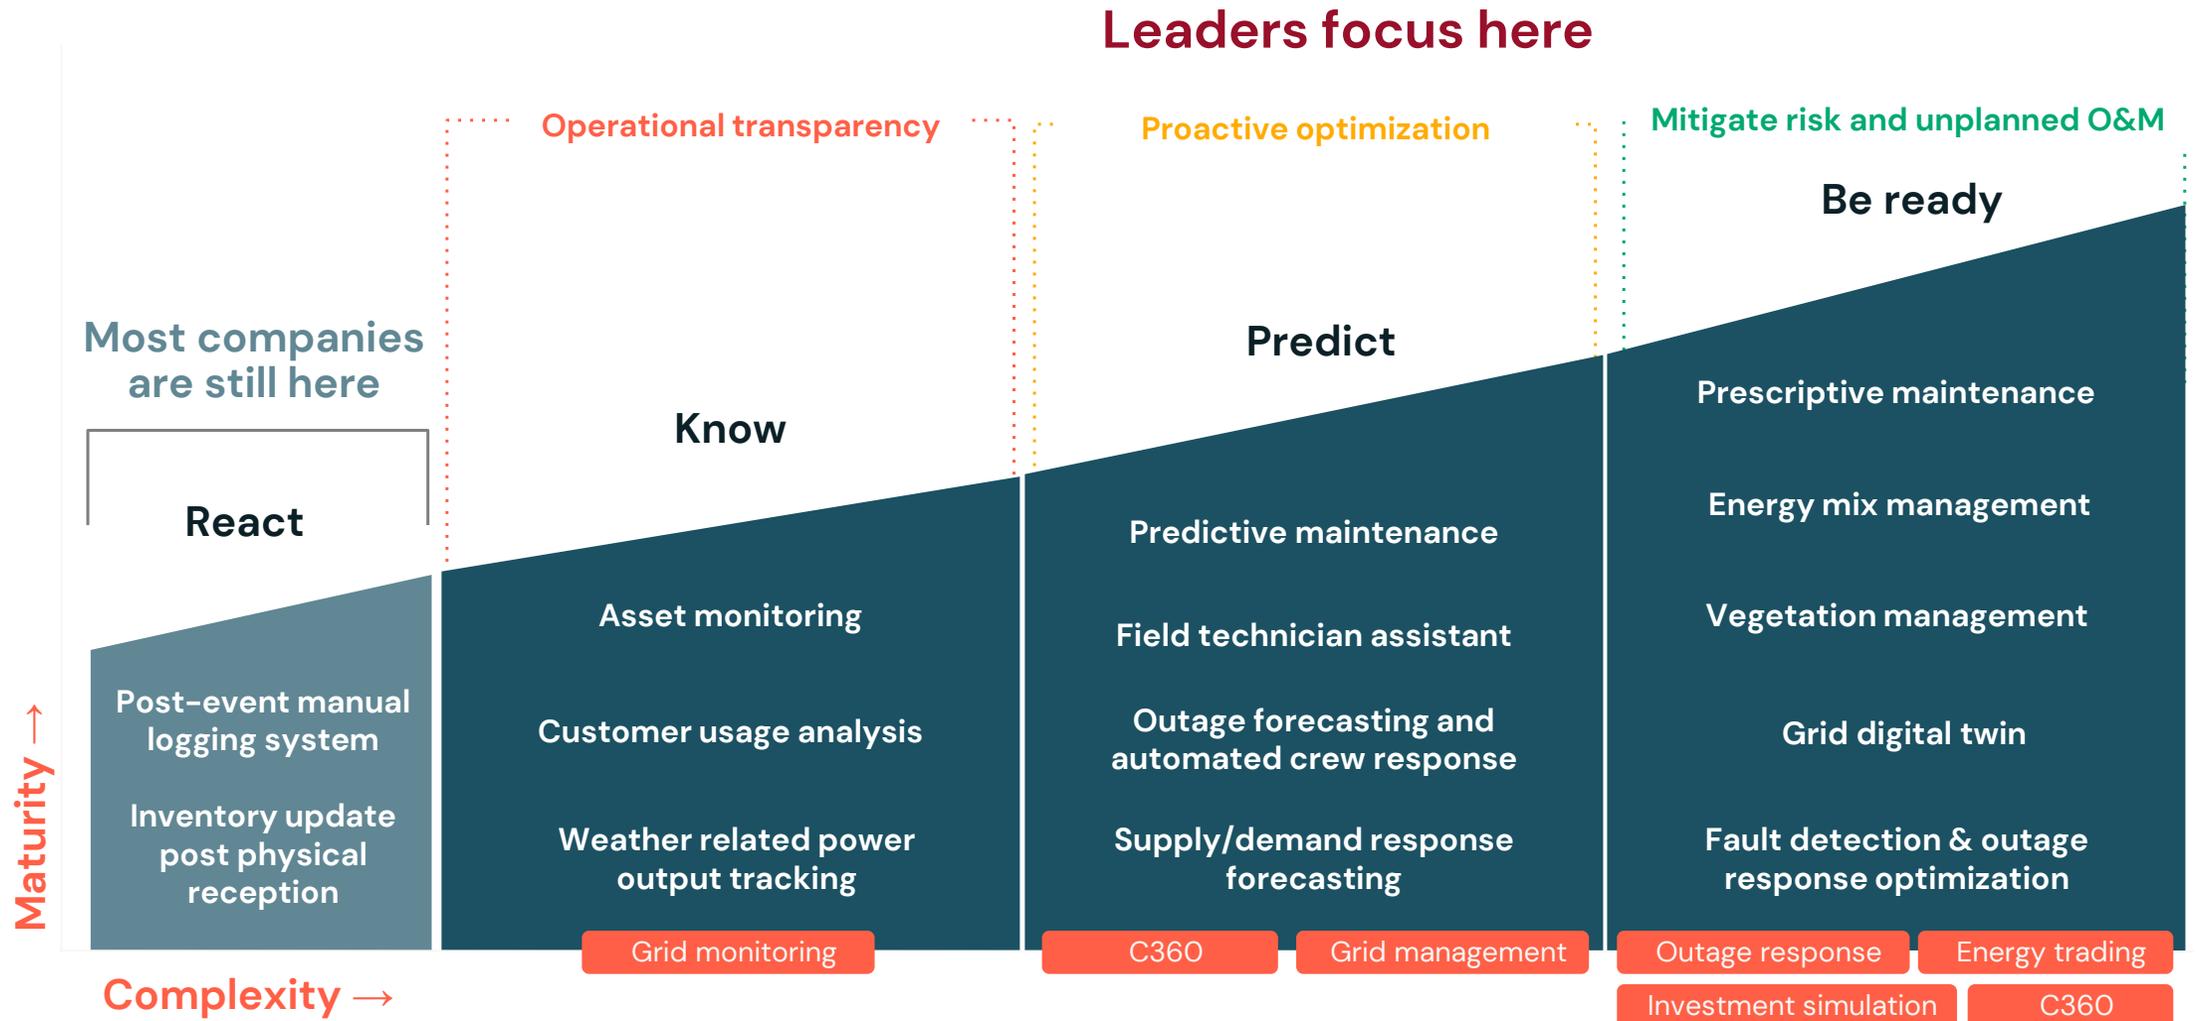
A healthy component is flagged, resulting in an unnecessary manual inspection.

By optimizing for the F2 Score, we train our AI to prioritize finding every potential threat, aligning with the primary goal of any reliability program: safety and the prevention of failure.

The electricity grid's complexity has exploded in recent years.  
Consumers are using more electricity and doing so in less predictable ways.

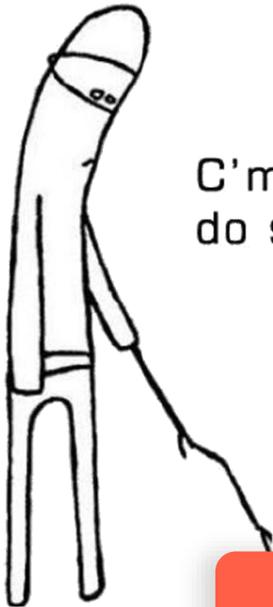


# Meeting utilities where they are on their AI journey



# Requirements

What's needed to get there



C'mon,  
do something...

**Warehousing**

What happened yesterday



# Requirements

What's needed to get there

## Predictive models

What's going to happen

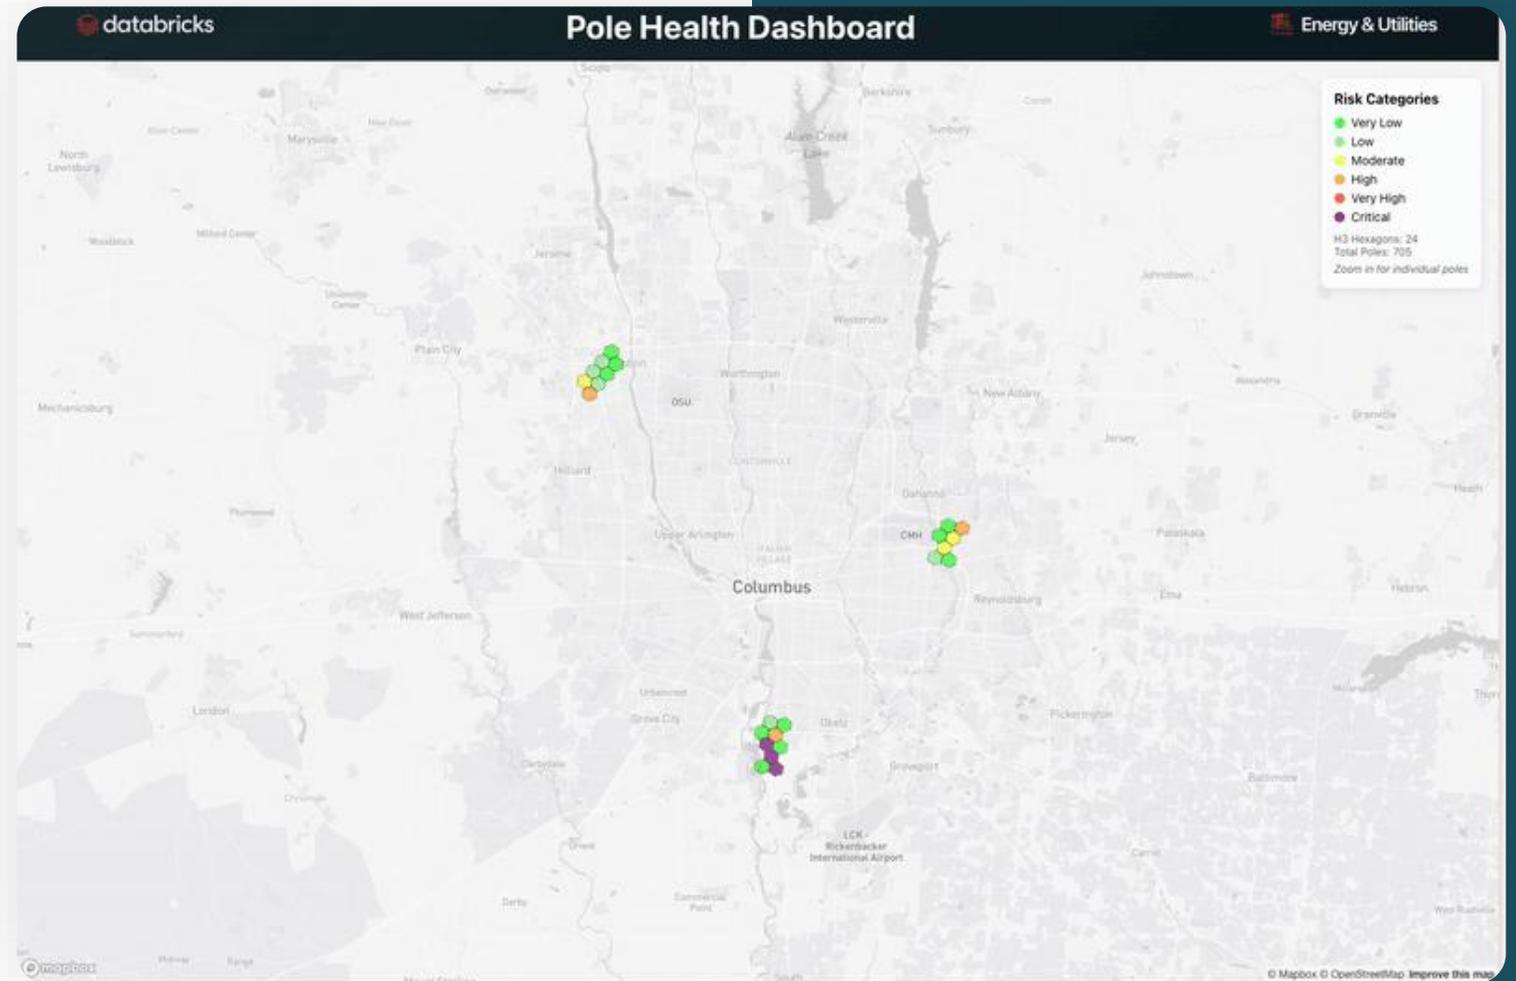
## Warehousing

What happened yesterday



# Requirements

Predictive models



# Requirements

Fast data

**Fast data**

Sensor data

**Predictive models**

What's going to happen

**Warehousing**

What happened yesterday



# What does this mean for the network?

1

## Local network congestion

Significant reverse power flow

Raises local voltage

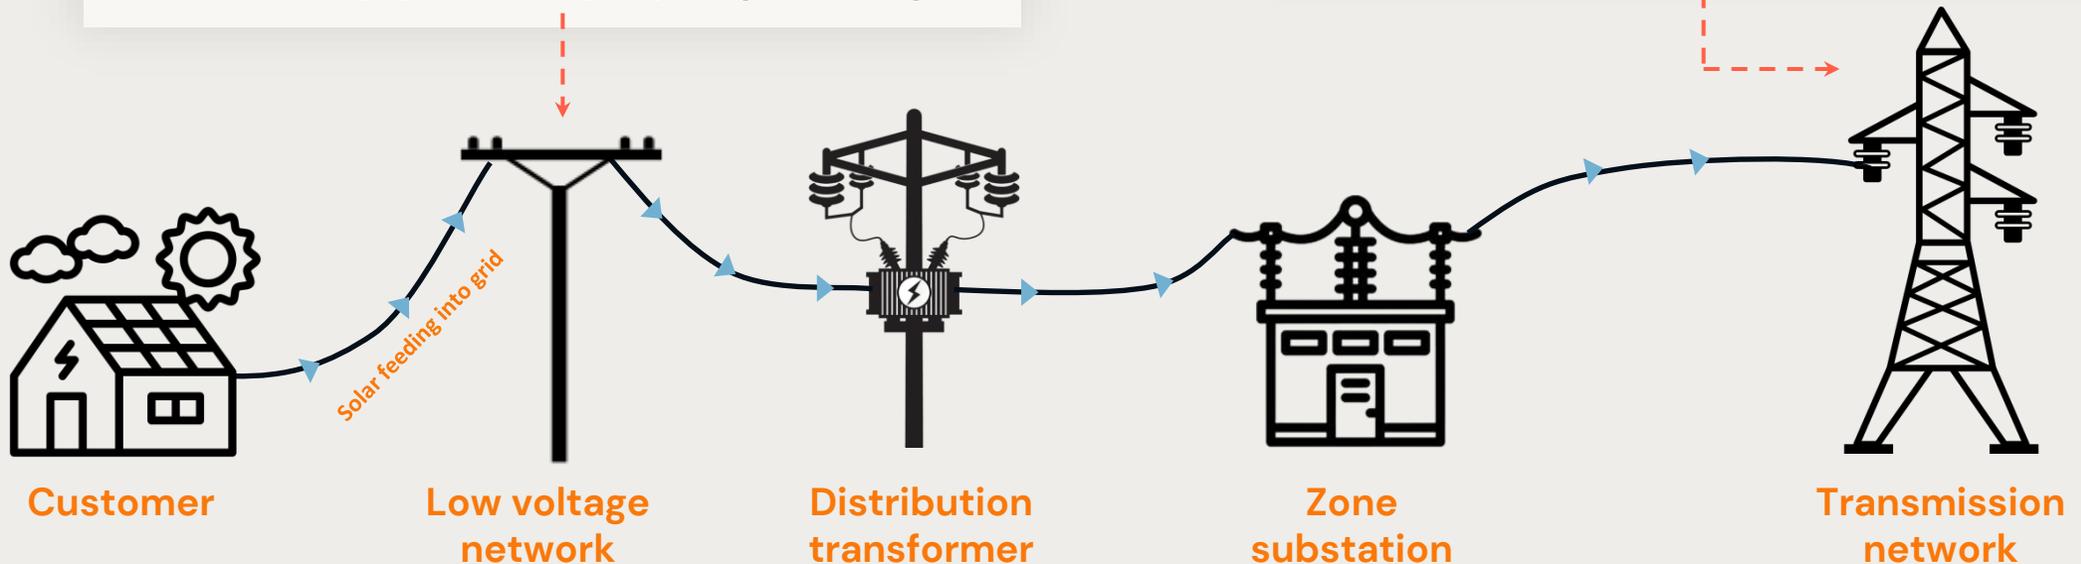
Customer equipment may trip or get damaged

2

## Maintaining system security

Entire state powered by solar

Reduces system inertia



# Flexible exports

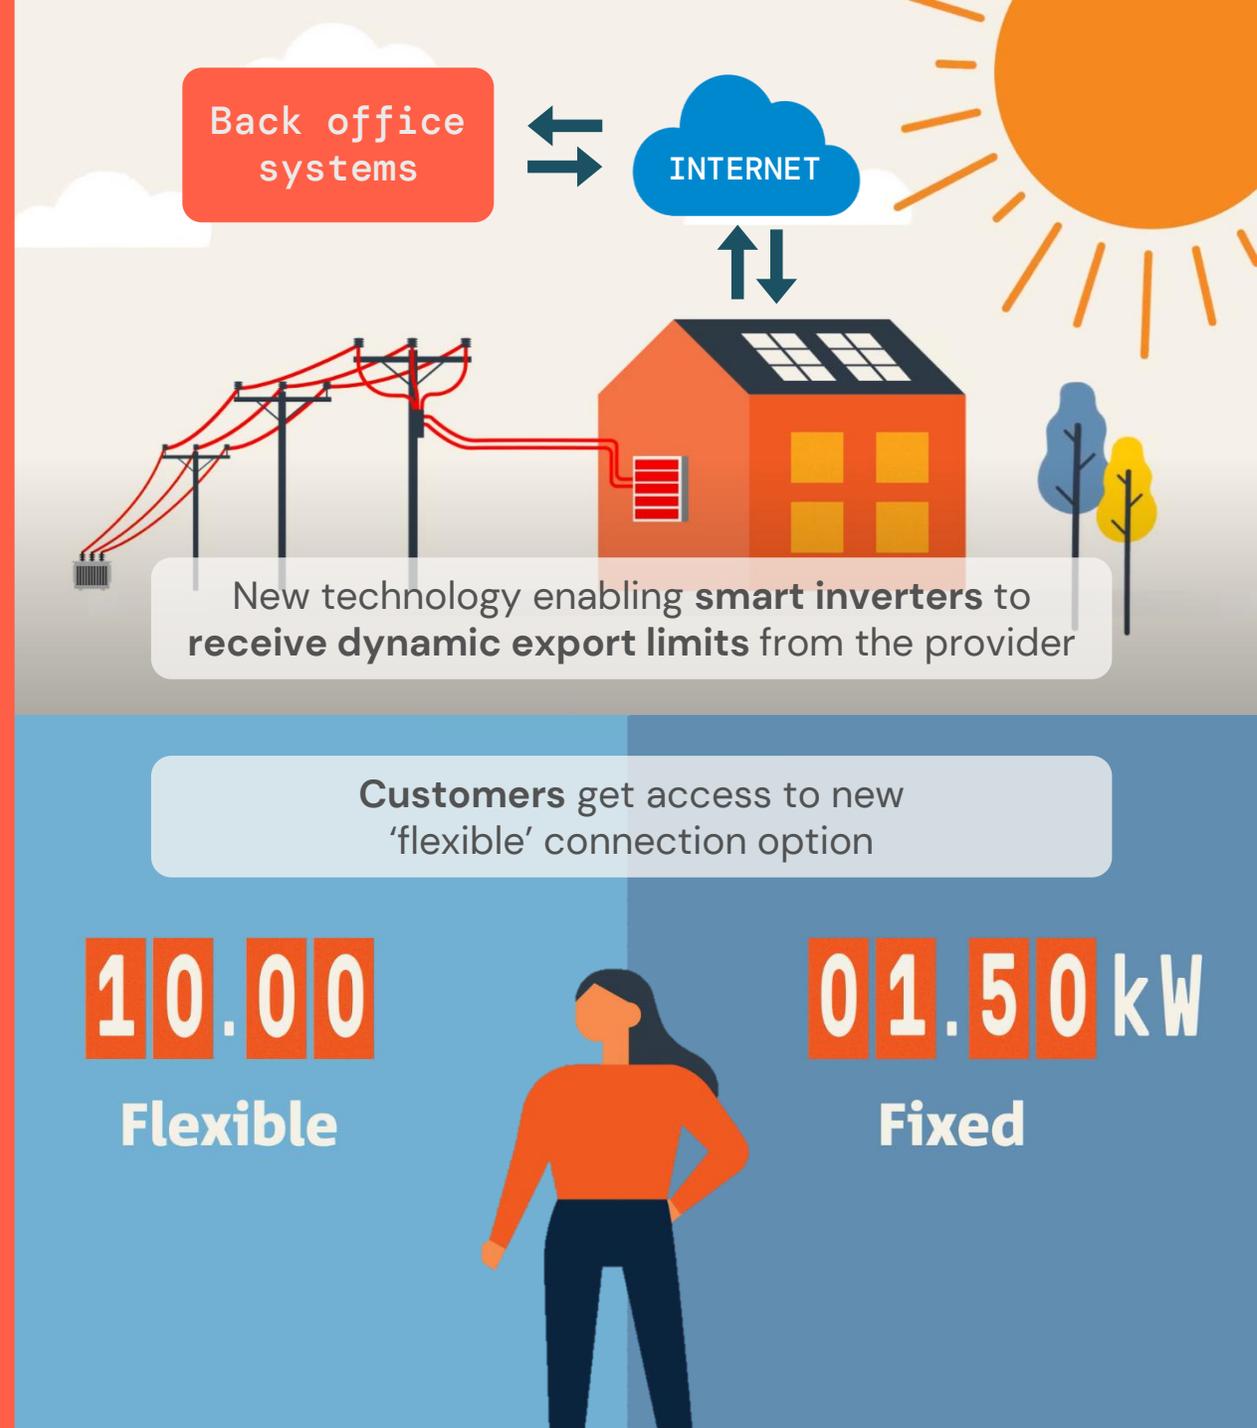
Our first flexibility offer!

Our traditional offer was fixed 5kW/phase

We're now offering a choice between

- Fixed export 1.5kW/phase
- Flexible Exports 1.5–10kW/phase

Flexible export limits are at 10kW for 95%+ of the time



# Requirements

What's needed to get there

## Multimodal

Images/PDF/Geospatial

## Fast data

Sensor data

## Predictive models

What's going to happen

## Warehousing

What happened yesterday



# Requirements

Multimodal: geospatial



**1.4M**

Acres burned per year



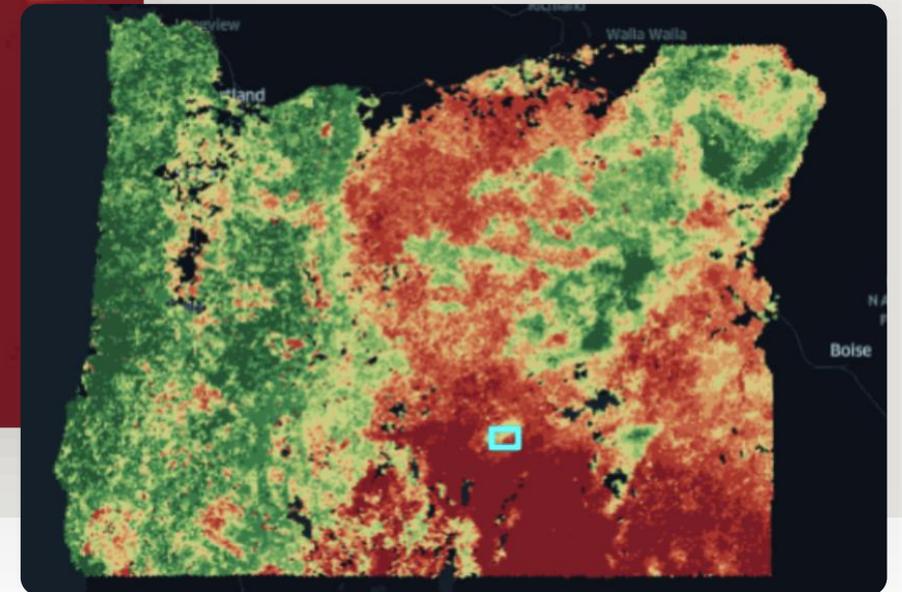
**\$1.2B**

in property damage per year



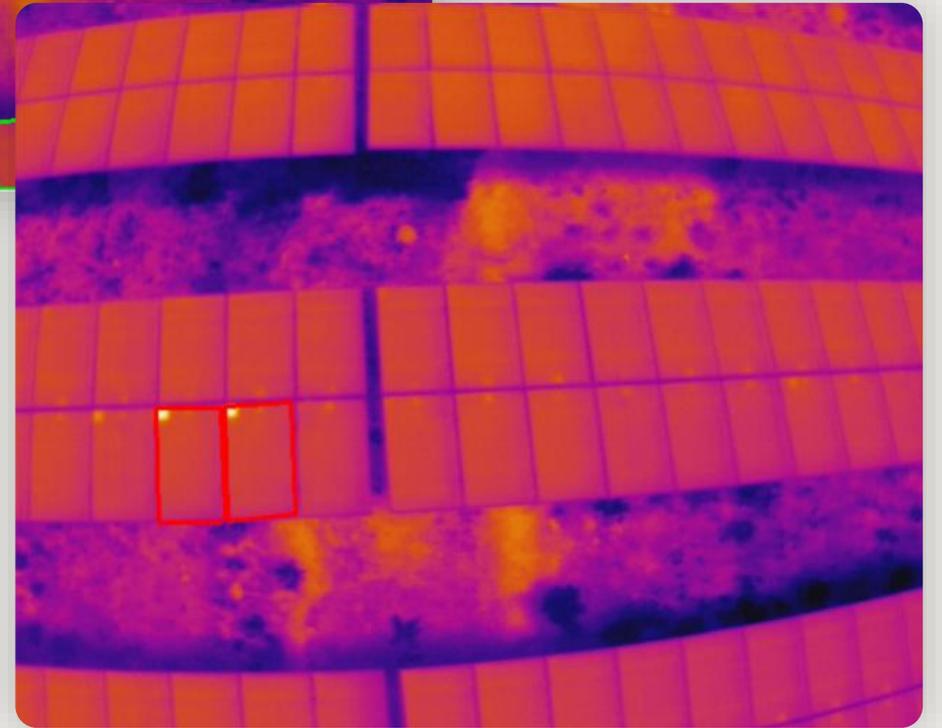
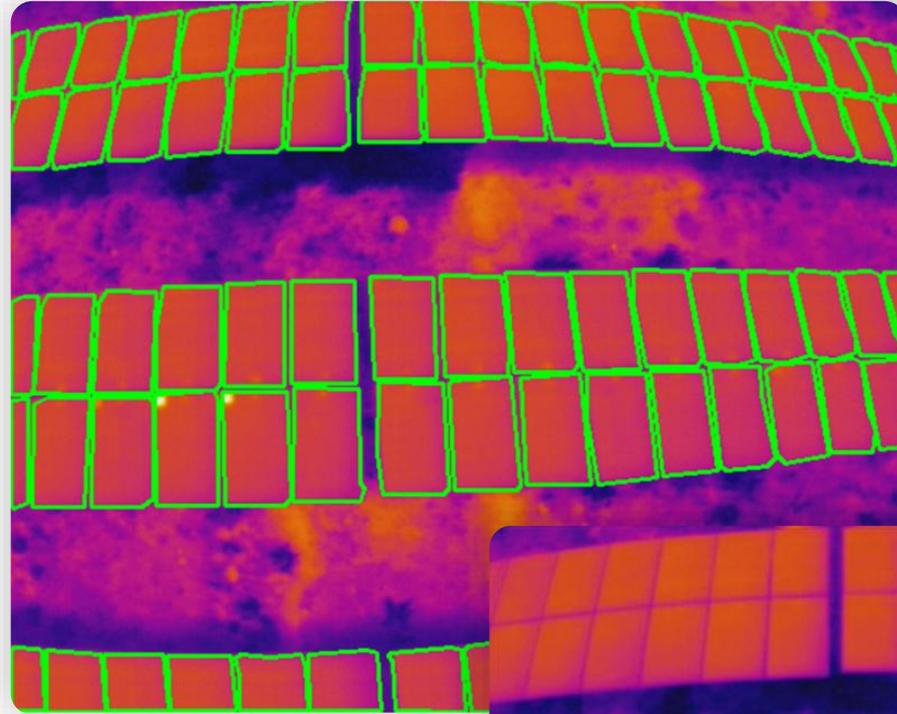
# Requirements

Multimodal: geospatial



# Requirements

Multimodal: images



# Requirements

What's needed to get there

**Generative**  
Agentic models

**Multimodal**  
Images/PDF/Geospatial

**Fast data**  
Sensor data

**Predictive models**  
What's going to happen

**Warehousing**  
What happened yesterday



# Requirements

Multimodal: PDF + Agentic

```
{
  "system_condition_explain" : "Mouse chewed wires"
  "panel_age" : "15"
  "minimum_photo_requirements" : [
    0 : "Electrical box with panel off"
  ]
  "wiring_type" : "Copper"
  "year_built" : "2010"
  "hazards_present" : [
    0 : "Loose wiring"
  ]
  "date_inspected" : "2025-08-04"
  "year_last_inspected" : "2010"
  "insured_name" : "John Doe"
  "system_condition" : "Unsatisfactory"
  "brand" : "EX"
  "address_inspected" : "535 W Hollywood"
  "policy_number" : "99831-2"
}
```

4-Point Inspection Form

Insured/Applicant Name: John Doe Application / Policy #: 998312  
Address Inspected: 535 W Hollywood  
Actual Year Built: 2010 Date Inspected: 2025-08-04

Minimum Photo Requirements:  
 Dwelling: Each side  Roof: Each slope  Plumbing: Water heater, under cabinet plumbing/drains, exposed valves  
 Main electrical service panel with interior door label  
 Electrical box with panel off  
 All hazards or deficiencies noted in this report  
A Florida-licensed inspector must complete, sign and date this form.

Be advised that Underwriting will rely on the information in this sample form, or a similar form, that is obtained from the Florida licensed professional of your choice. This information only is used to determine insurability and is not a warranty or assurance of the suitability, fitness or longevity of any of the systems inspected.

**Electrical System**  
Separate documentation of any aluminum wiring remediation must be provided and certified by a licensed electrician.

<b>Main Panel</b> Type: <input checked="" type="checkbox"/> Circuit breaker <input type="checkbox"/> Fuse Total Amps: <u>120</u> Is amperage sufficient for current usage? <input type="checkbox"/> Yes <input type="checkbox"/> No (explain)	<b>Second Panel</b> Type: <input type="checkbox"/> Circuit breaker <input type="checkbox"/> Fuse Total Amps: _____ Is amperage sufficient for current usage? <input type="checkbox"/> Yes <input type="checkbox"/> No (explain)
--	--

Indicate presence of any of the following:  
 Cloth wiring  
 Active knob and tube  
 Branch circuit aluminum wiring (if present, describe the usage of all aluminum wiring):  
\* If single strand (aluminum branch) wiring, provide details of all remediation. Separate documentation of all work must be provided.  
 Connections repaired via COPALUM crimp  
 Connections repaired via AlumConn

<b>Hazards Present</b> <input type="checkbox"/> Blowing fuses <input type="checkbox"/> Tripping breakers <input type="checkbox"/> Empty sockets <input checked="" type="checkbox"/> Loose wiring <input type="checkbox"/> Improper grounding <input type="checkbox"/> Corrosion <input type="checkbox"/> Over fusing	<input type="checkbox"/> Double taps <input type="checkbox"/> Exposed wiring <input type="checkbox"/> Unsafe wiring <input type="checkbox"/> Improper breaker size <input type="checkbox"/> Scorching <input type="checkbox"/> Other (explain)
---	---

General condition of the electrical system:  Satisfactory  Unsatisfactory (explain)  
Mouse Chewed Wires

Supplemental information

<b>Main Panel</b> Panel age: <u>15</u> Year last updated: <u>2010</u> Brand/Model: <u>EX</u>	<b>Second Panel</b> Panel age: _____ Year last updated: _____ Brand/Model: _____	<b>Wiring Type</b> <input checked="" type="checkbox"/> Copper <input type="checkbox"/> NM, BX or Conduit
---	---	--

Sample Form Insp4pt 01 18



# Requirements

What's needed to get there

**Apps**

Business facing

**Generative**

Agentic models

**Multimodal**

Images/PDF/Geospatial

**Fast data**

Sensor data

**Predictive models**

What's going to happen

**Warehousing**

What happened yesterday

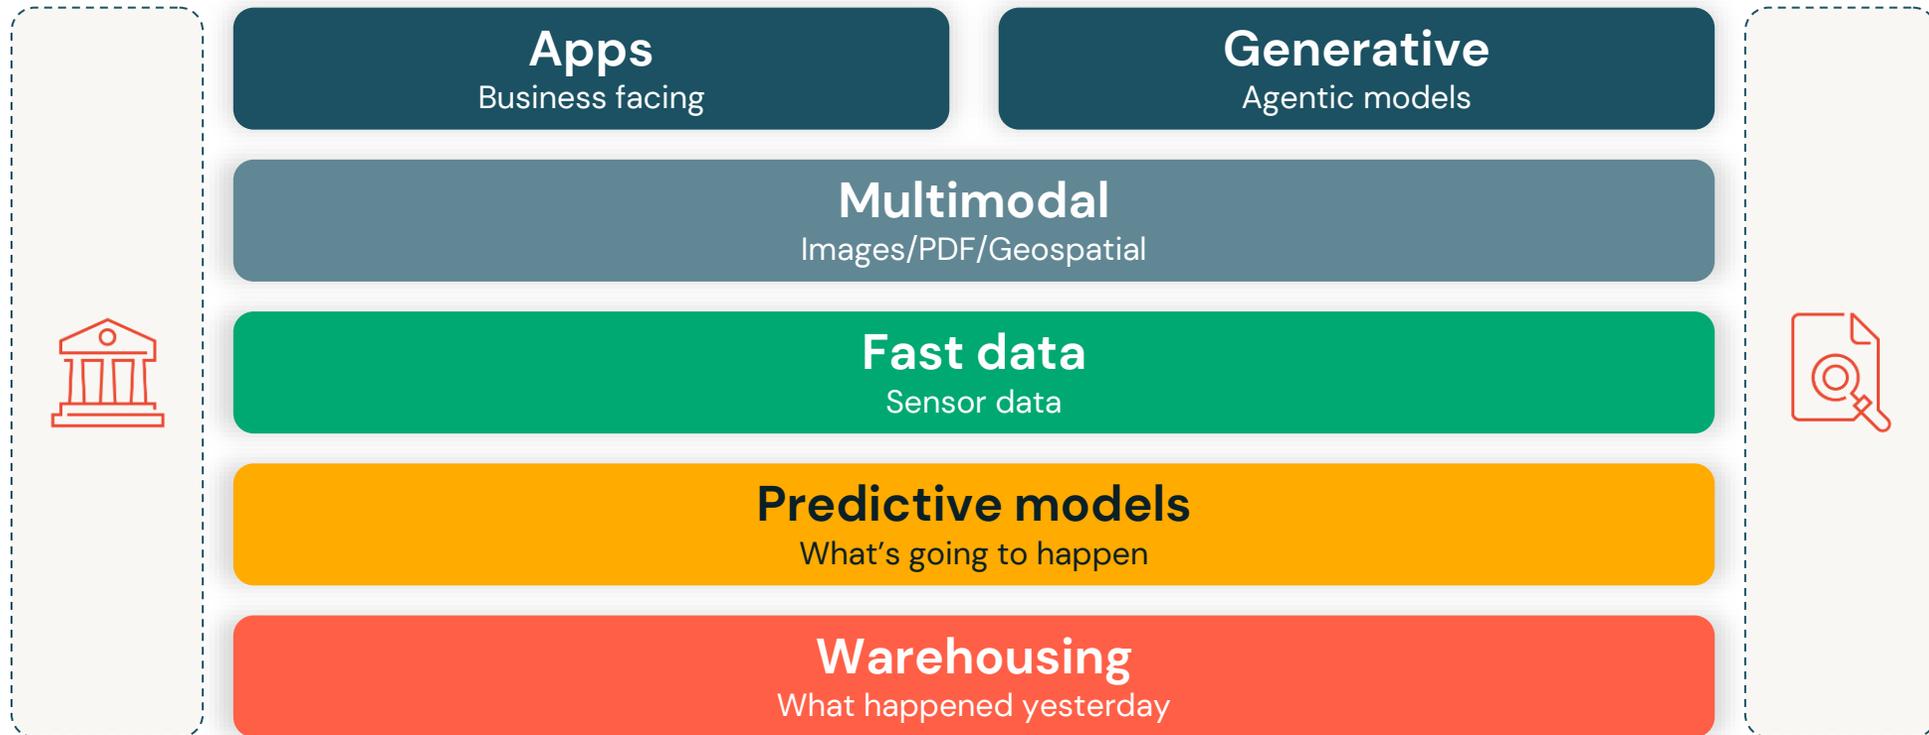


**DEMO break**



# Requirements

What's needed to get there



# Resources



# Appendix



# Segment Anything

Upload an image, then draw a bounding box OR enter a text prompt (or both)



Click to upload or drag and drop

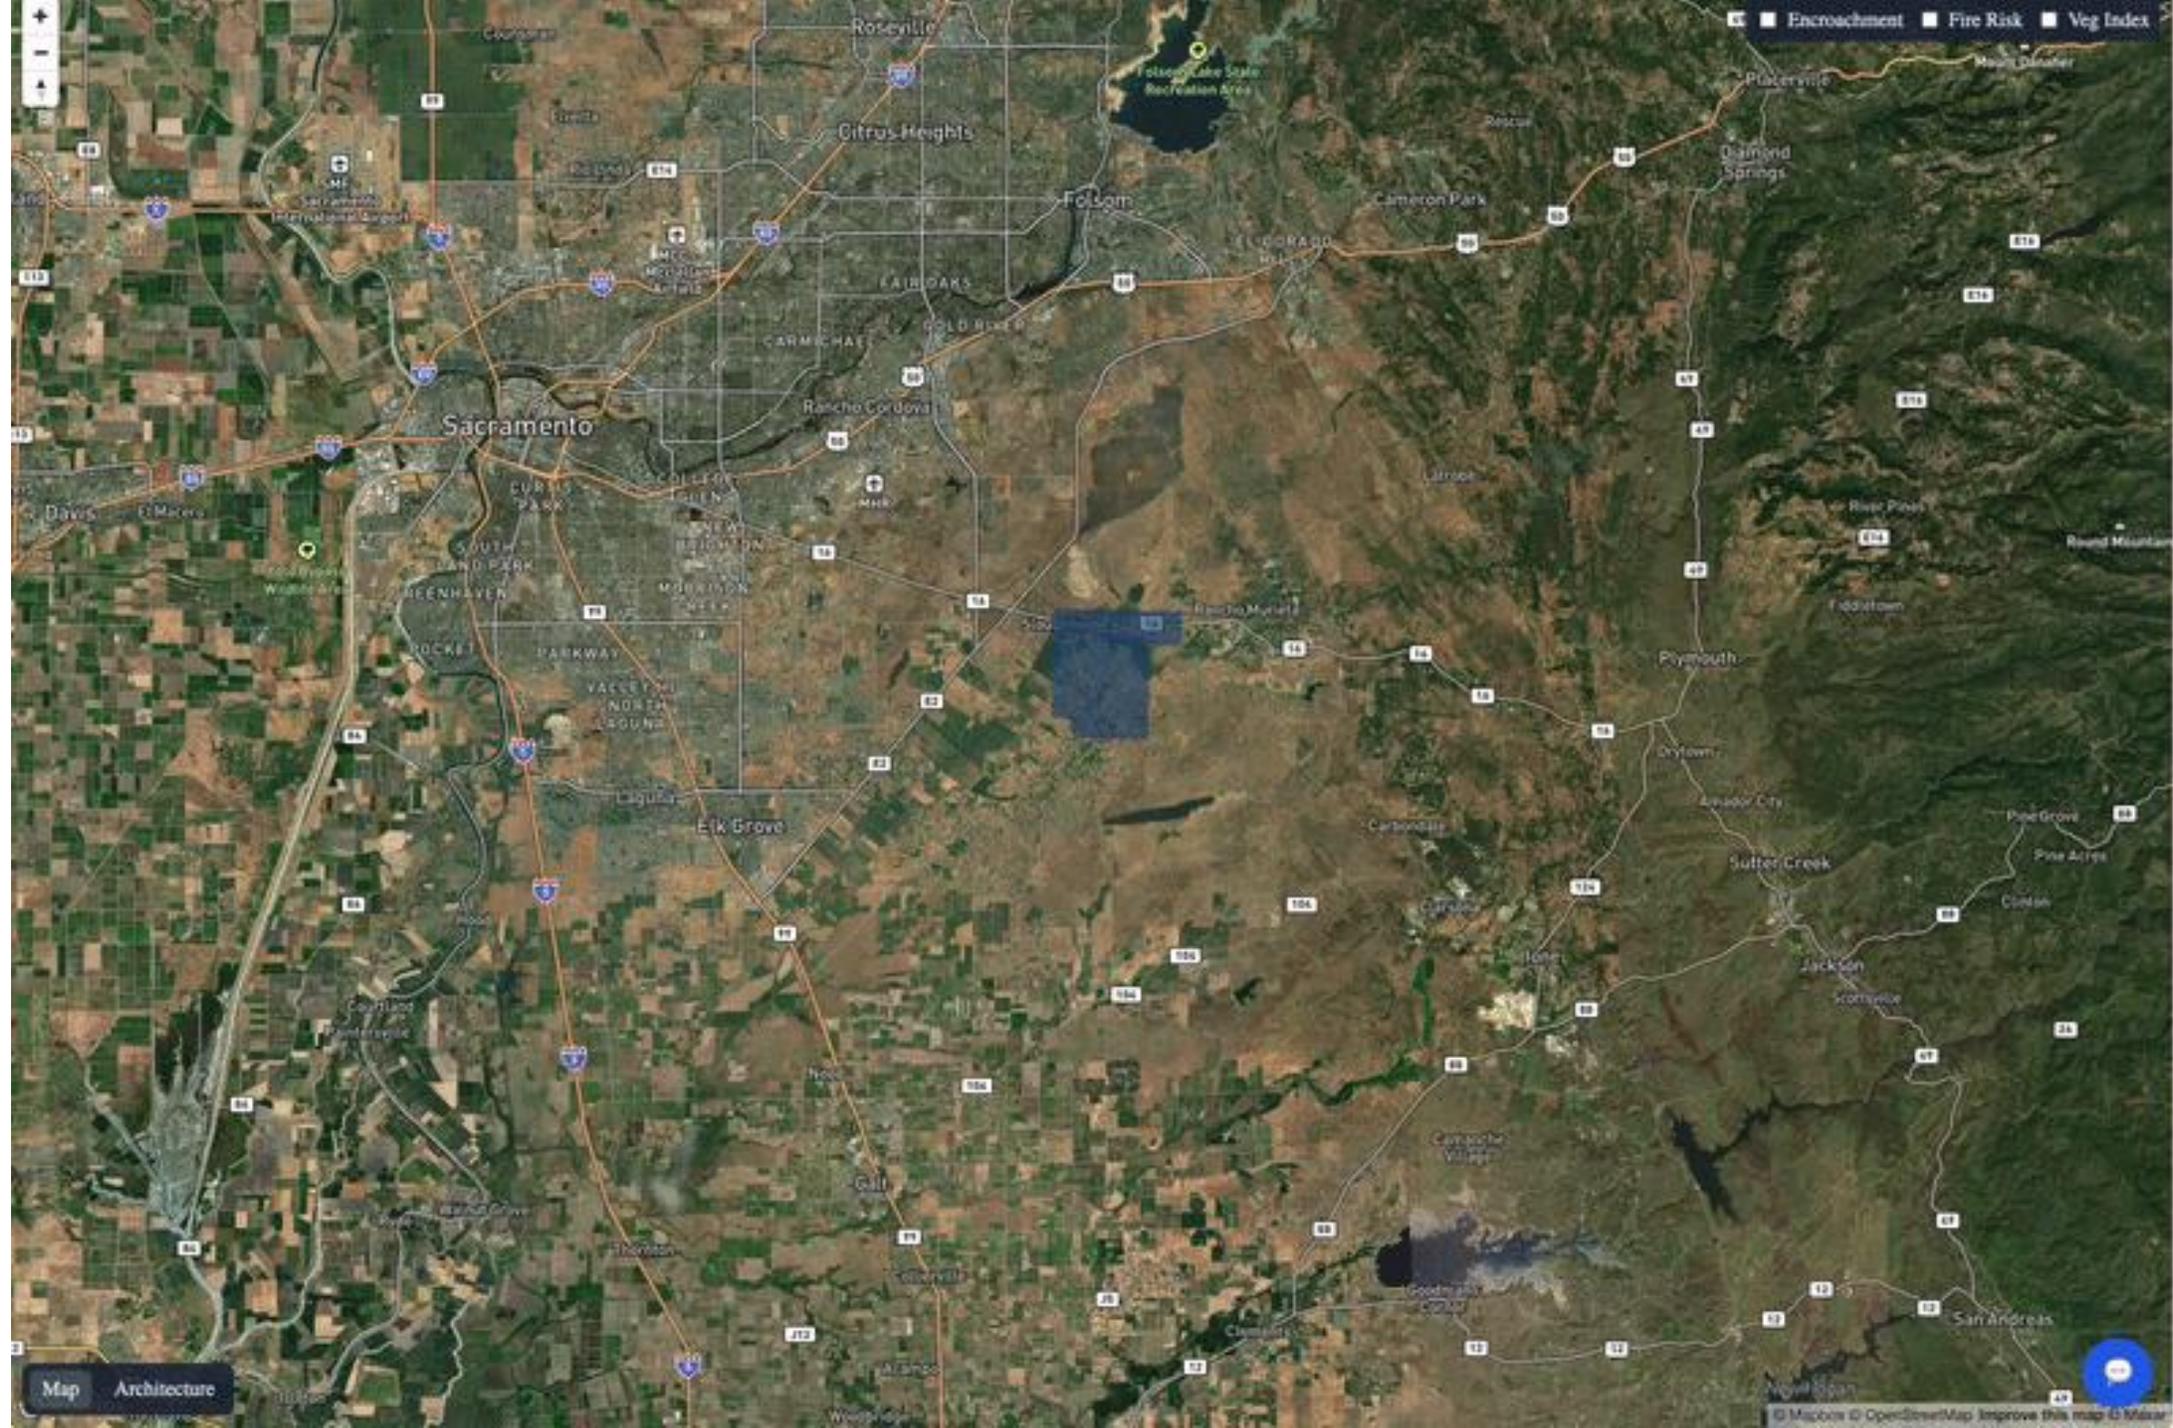
PNG, JPG, GIF up to 10MB

 Use Webcam

 Load Sample

Built with  
**APX** apx







*Exceptional service in the national interest*

# Towards advance artificial intelligence (AI) tools in the energy sector: Data Acquisition for AI Research and Inference (DAQAIRI)

Arturo Bretas, *PhD*  
Distinguished Member of Technical Staff  
Sandia National Laboratories

Monday, February 23, 2026



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525

SAND2025-15008PE

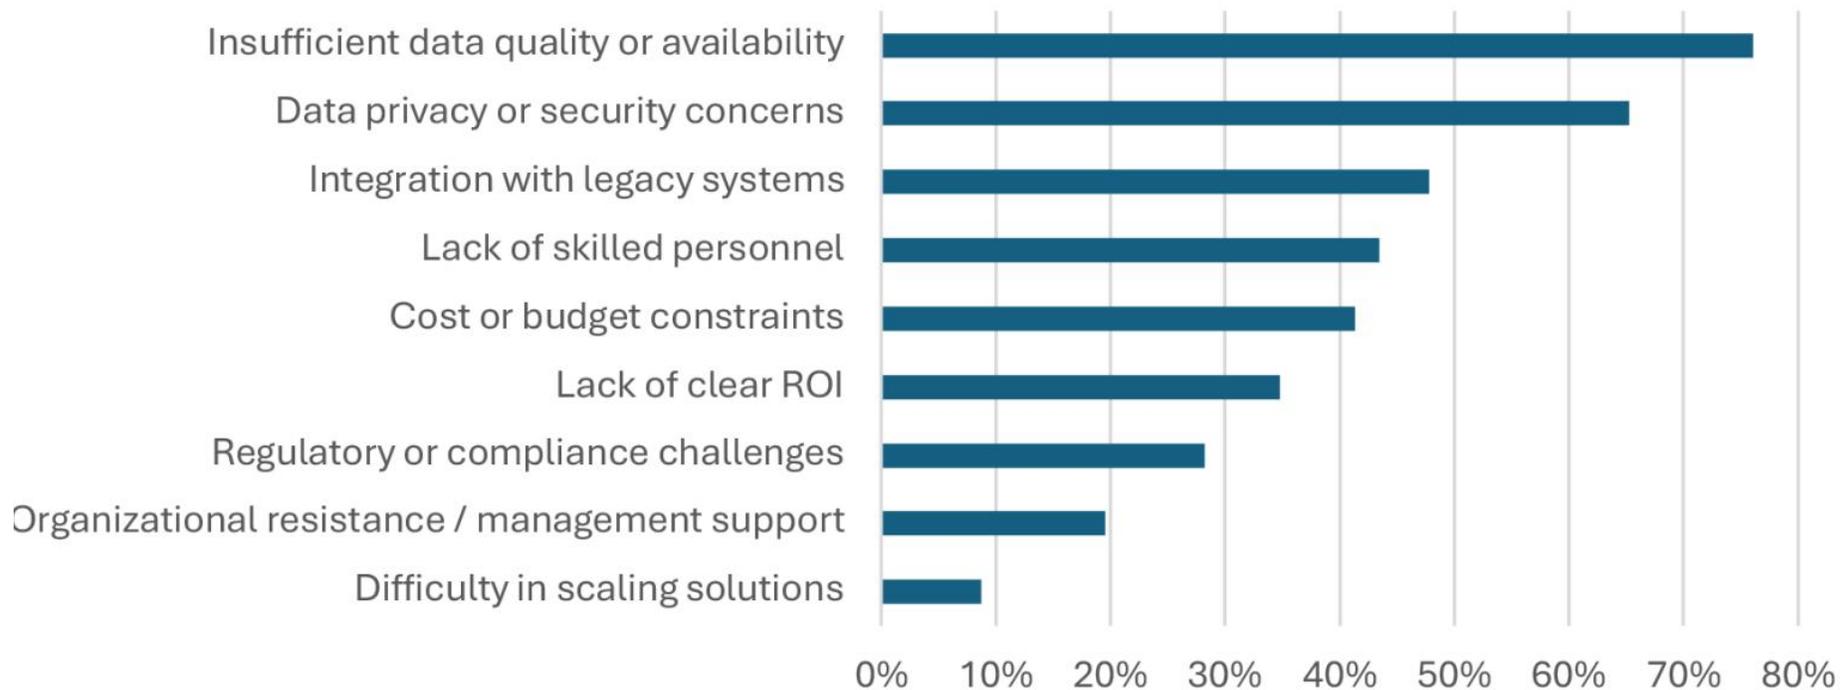
# AGENDA

- Introduction
- DAQAIRI's Platform
- AI Models

# Background

# Overview: High quality data is needed

What are the primary challenges or barriers to AI adoption in your organization? (select all that apply)



**Figure Note:** The survey included responses from 52 organizations including 28 utilities, 22 energy technology and AI providers and 2 research organizations where each organization could select all applicable options.

# Overview: Why DAQAIRI?

Energy datasets already exist

DOE's Open Energy Data Initiative (OEDI)

Grid Event Signature Library (OE)

North American Synchrophasor Initiative (OE/EPRI)

Open Power AI (EPRI)

IEEE DataPort (IEEE)

American Science Cloud (SC)

Transformation AI Models Consortium (SC)

# Overview: Why DAQAIRI?

Existing data is often:

Synthetic

From test systems, not production systems

Single data type/source

Limited time/space coverage (i.e. not full year of data)

Designed to answer only one analysis question, not necessarily cybersecurity

# Overview: Why Daqairi?

DAQAIRI will have OT datasets that are:

From real, production systems

Comprehensive: wide temporal/spatial coverage

High quality (metadata, structured, labels)

Multi-Modal

Environmental/Physical Sensor Data

AMI

PMU

SCADA

Communication data

IED Data



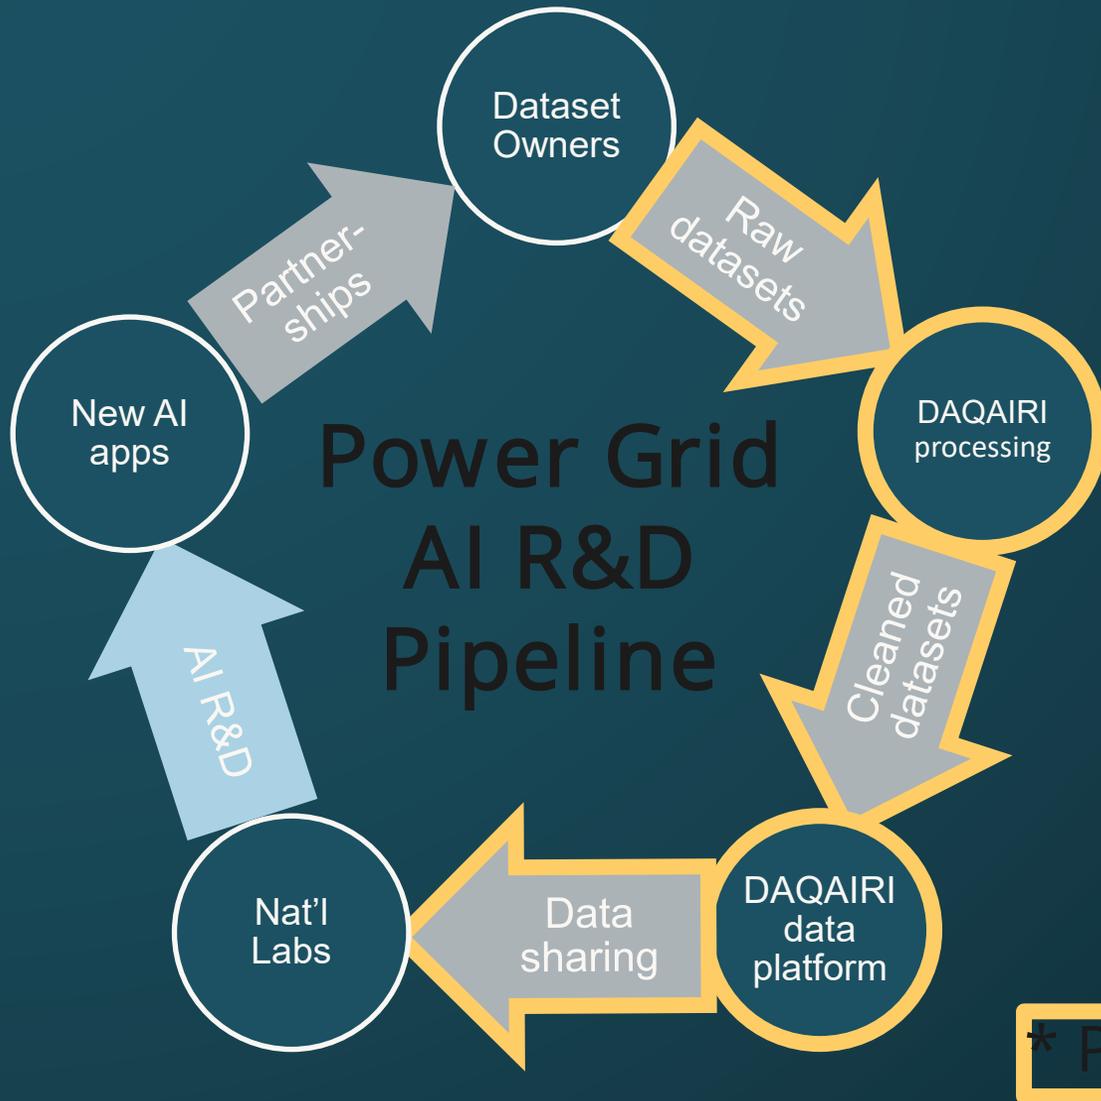
Modern AI R&D needs lots of high quality data

# DAQAIRI: what is it?



The design of DAQAIRI's platform consists of 4 distinct software modules.

# Overview



DAQAIRI will be a platform for dataset owners to share OT datasets with national labs  
National Labs will share the developed AI tools with dataset owners  
Goal: datasets for AI R&D for grid security/reliability

# Overview: Potential AI Use cases



Workflow automation

Energy cost prediction models

Cybersecurity

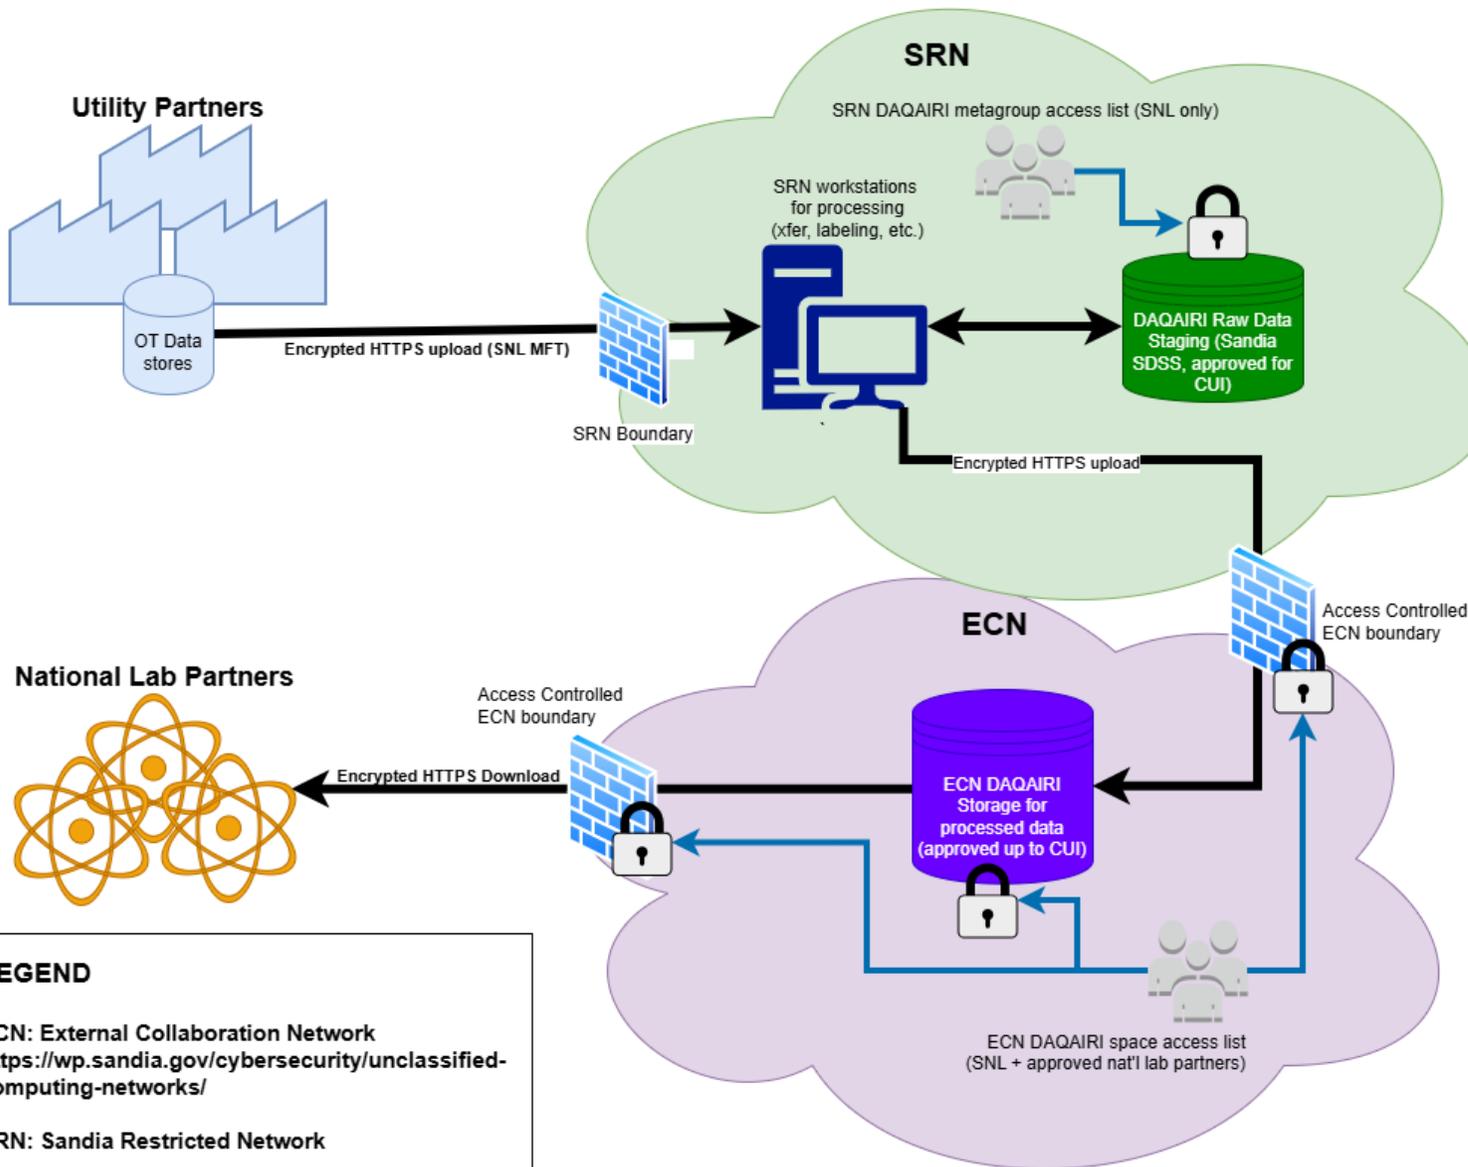
Resilience modeling

Model and simulation validation

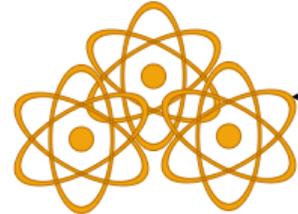
Fault location

Load prediction

# DAQAIRI Data Flow



National Lab Partners



## LEGEND

ECN: External Collaboration Network  
<https://wp.sandia.gov/cybersecurity/unclassified-computing-networks/>

SRN: Sandia Restricted Network

➔ : Dataset Movement

➡ : Access Control

- Architecture leverages existing SNL infrastructure
- No direct cost to DAQAIRI, covered by SNL overhead
- Approved for CUI/proprietary data
- Raw data: SNL access only
- Processed/cleaned data: SNL + APPROVED partner access only
- May need expansion for large datasets (>100's of GB)



# Data Management Plan

User/Role based Access controls for all data

Access controlled by DAQAIRI team

Encryption for all data transfers (HTTPS)

Encryption at rest available if required by data owner

All data will be sanitized for sensitive information before sharing

Partner sharing will be approved by all data owners before partner access

Data stores have 30 day recovery backup

DAQAIRI will maintain at least 1 long term data backup





# Cybersecurity for the electrical grid program breadth

*Research including, exploration of fundamental energy security concepts, applied analytical tools, field-based systems integrations, and codifying best practices as precursors to new codes and standards. We partner with diverse, multidisciplinary internal and external organizations bringing to bear Sandia's considerable capabilities in addressing national grid security.*

*Significant internal investments, decades of experience, and deep cybersecurity expertise supporting DOE applied energy offices, notably CESER*

## RESEARCH TOPICS



*Industrial Control Systems*



*Network Intrusion Detection (NID) & Mitigation*



*Comms. & Cybersec. for Energy Distribution Networks & Edge Devices*



*System/Device Cyber Vulnerability Analysis*



*Cybersecurity Standards*





# Presentation OVERVIEW



## PROJECT HIGHLIGHTS

- SolarSnitch
- sDERC
- DER HERO
- ECS OT Devices
- CyTRICS
- BESS CS
- Goaltender
- VEDIZAR





# SolarSnitch | Overview

*SolarSnitch is a cybersecurity solution that provides both cyber and physics-based (or physical) detection and mitigation capabilities.*

## CURRENT STATE

*Most Cybersecurity solutions are designed for traditional cyber IT networks and only process cyber data*



- Deployed within utility SCADA networks and centralized implementation at a control or data center



- Costly and extensive configuration for individual systems



- Does not monitor power system behavior or correlate cyber events that are causing power system events

## SOLARSNITCH

*A low-cost, bump-in-the-wire cybersecurity defense uniquely suitable for grid-edge systems*

- Distributed implementation and peer-to-peer communication deliver a scalable solution
- Flexible implementation as software containers or bump-in-the-wire on a low-cost single-board computer
- Monitors both cyber and physical data for increased detection capability and situational awareness at the grid-edge
- Correlates cyber and physical events and deploy proactive mitigation

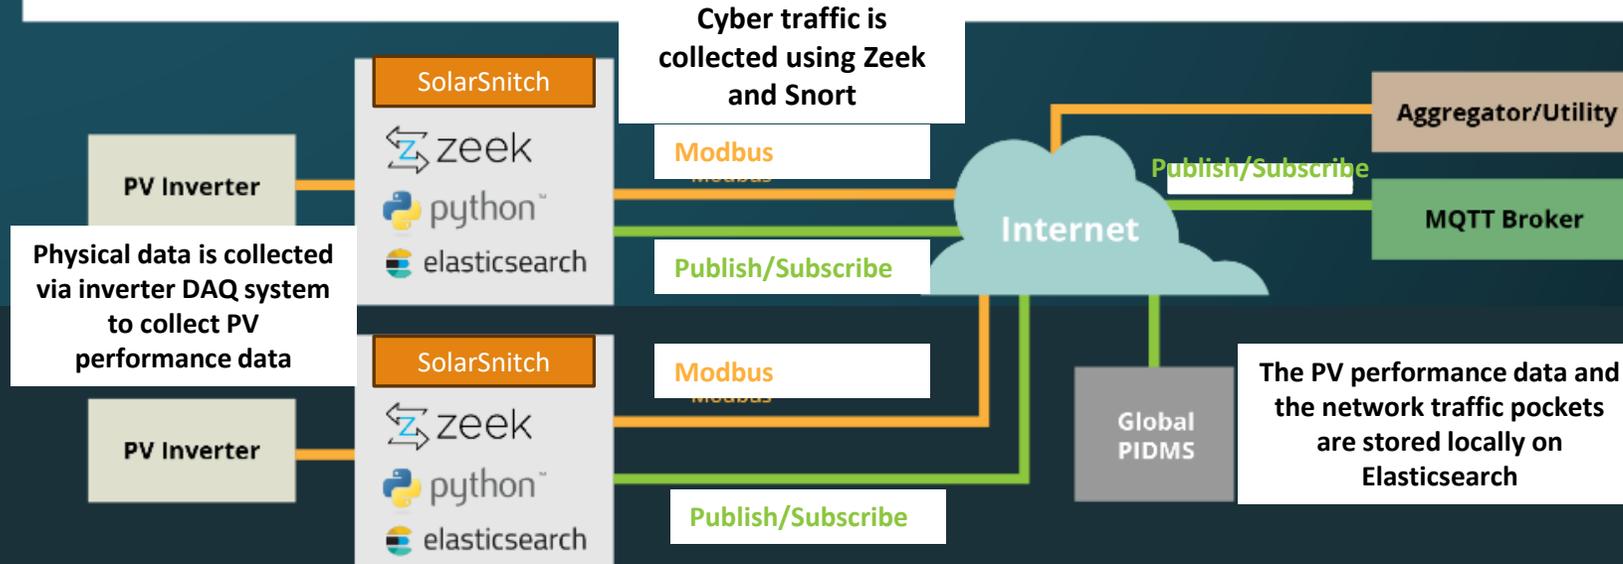




# SolarSnitch | Approach: Cyber-Physical Analysis

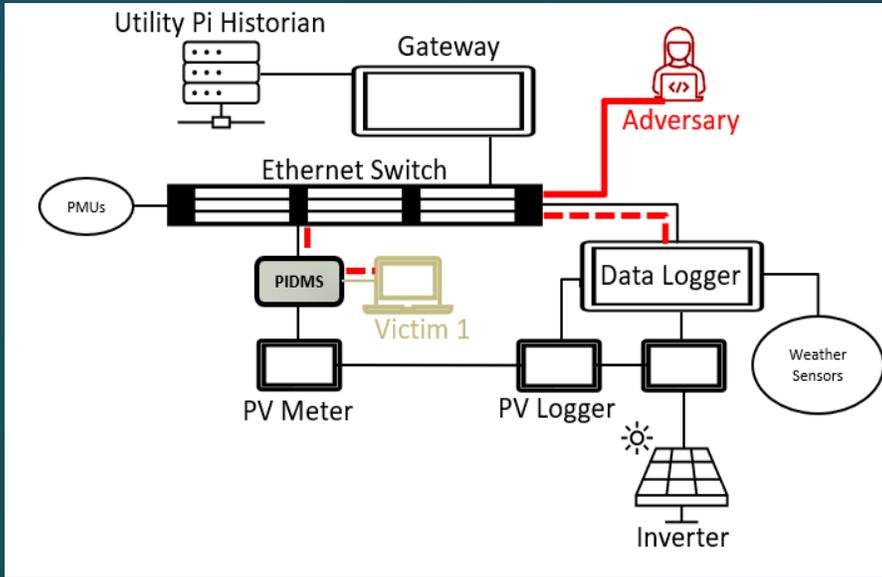
SolarSnitch

Previous Name

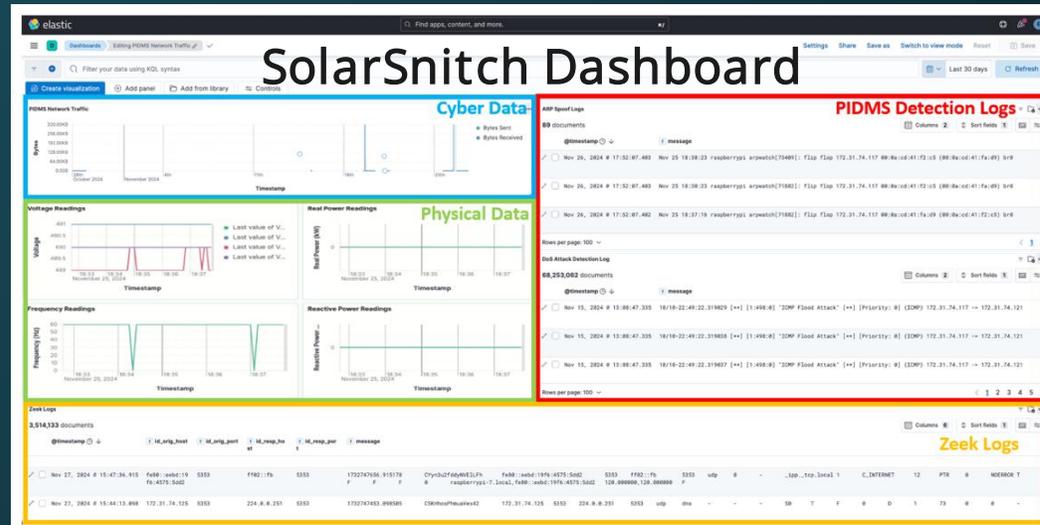
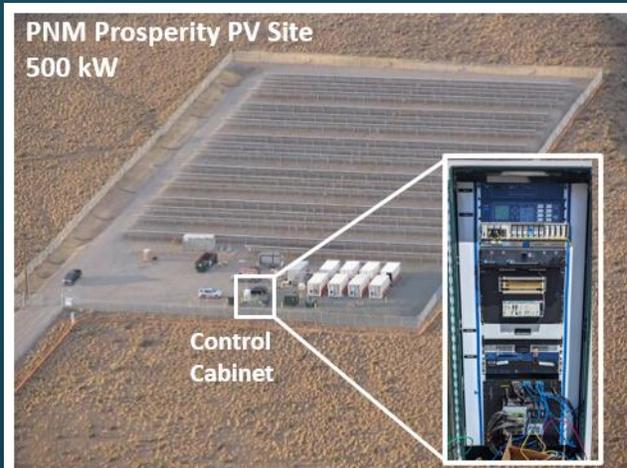


*The SolarSnitch is designed to collect cyber-physical data, process the data using machine-learning algorithms, detect abnormal events, and deploy mitigations*

# SolarSnitch | Field Testing Results



- Testing SolarSnitch at partner PNM's Prosperity PV Site
- Tested against ARP Spoof and Machine-in-the-Middle Attacks successfully
- Next steps: focus on resilience against data integrity issues and improving machine learning robustness





# DER HERO | Overview

## A Hybrid and Explainable DER Network Intrusion Detection System

### CURRENT STATE

*In FY24, we successfully completed the project titled, "DER NIDS: Application Layer Maps to Extract Real-Time Physical and Control Data".*

A protocol-agnostic DER NIDS based on the digital twin technology to:

- Enable the cyber-physical analysis of DER network traffic
- Detect stealthy adversary attacks (e.g., FDIA)

The digital twin utilizes a ruleset of physical rules to detect threats in the critical infrastructure/grid

### CURRENT CHALLENGES & WHAT OUR SOLUTION ADDS

17.238

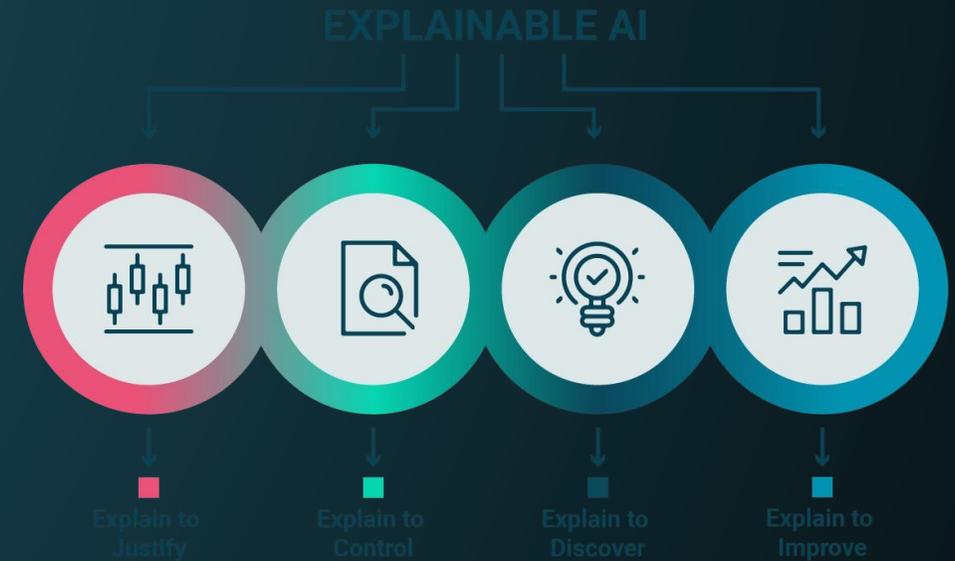
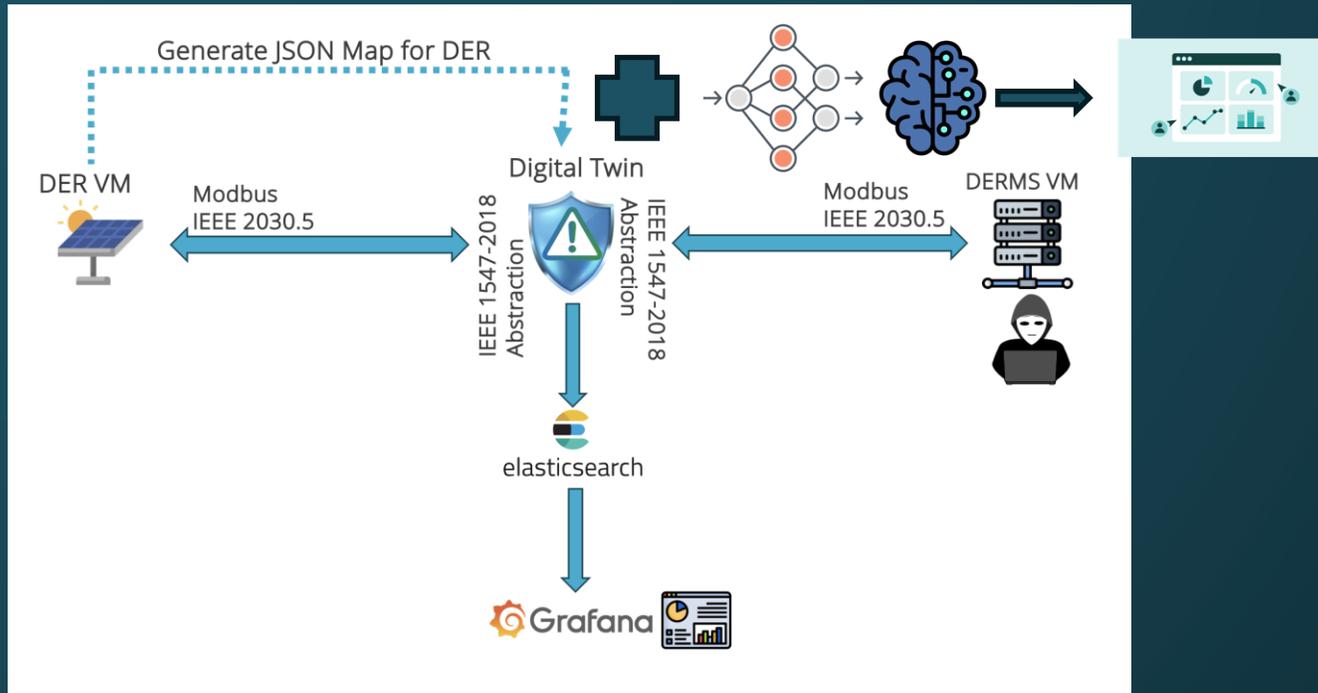
**Challenge 1:** We cannot capture all the potential malicious actions with a ruleset – we would need thousands of rules

**Addition:** Enrich the deterministic ruleset of the digital twin with a ML pipeline for high-fidelity data processing

**Challenge 2:** "How can the operators trust AI and ML for intrusion detection to provide cyber-physical security in the grid?"

**Addition:** Provide transparency to the operators regarding the ML decision-making via explainable AI

# DER HERO | Project plan



## Explainable AI-driven cyber-physical security with a digital twin

- Create an ML-driven digital twin that analyzes cyber-physical traffic (Modbus, IEEE 2030.5)
- Implement explainable AI wrappers around the ML models to understand why the ML models predicted something as "good"/"bad" and reduce the misclassifications (essential in a cyber-security use case)
- Transform the digital twin from a 'black box' to a 'white box' and protect the critical infrastructure
- Create xAI visualizations so the grid security operators trust ML-driven cyber-physical analysis

# DER HERO | Current Progress – Results

INSIDE Directory Policy News Search Sandia

Innovation Portal Search Help

### Technical Advance - TAQ586 (Pending DC Review)

#### GENERAL INFO

**Is this invention classified?**  
No

**Do you consider this invention to be work related?**  
Yes

**Title**  
DER HERO

**Summary**  
The proposed "DER HERO" is an Explainable Artificial Intelligent (AI)-based digital twin for real-time cyber-physical threat detection within the Distributed Energy Resources (DER) ecosystem. Specifically, DER HERO is a novel protocol-agnostic digital twin that ingests information from DER equipment via deep packet inspection and analyzes it real-time with end-to-end Machine Learning (ML) for stealthy cyber-physical attack detection, e.g., False Data Injection Attack. However, the grid operators do not trust black-box ML models in production systems if they do not understand how crucial decisions are made. The proposed digital twin is different from any developed DER/cyber-physical threat detection system because it utilizes explainable AI (XAI) wrappers around the ML pipeline to be transparent, trustworthy and reliable to use in production systems. DER HERO answers the question: "How can the operators trust AI for intrusion detection to provide cyber-physical security in the grid?"

**Point of Contact**  
Georgios Fragkos - 08812

---

#### ORIGINATORS

**Sandia Originators**

Name	POC	Org.	Status	Type	Phone	Email	Company
Georgios Fragkos	Yes	08812	Active	Employee	505-228-7587	gfragko@sandia.gov	
Timotius Oentung	No	08815	Active	Employee	505-249-1085	toentun@sandia.gov	
Nitin Sharan	No	08761	Active	Employee		nsharan@sandia.gov	

**External Originators**  
N/A

## Technical Advance for DER HERO

Visualization that shows the threats detected by the digital twin, the IPs and where the threats came from



## ML-driven cyber-physical analysis

# Goaltender | Cloud-Based Defense and Response Tools for the DER Ecosystem



*Defense-in-depth solutions for Distributed Energy Resource Management Systems (DERMS) with Next Generation Firewalls (NGFWs) and Security Orchestration, Automation, and Response (SOAR) solutions*

**CURRENT STATE:** Modern DERMS and Charging Station Management System (CSMS) systems protected at cloud boundary with firewall and IDS systems representing a perimeter defense approach; lack of introspection upon data before reaching legitimate backend systems.

## OUR SOLUTION

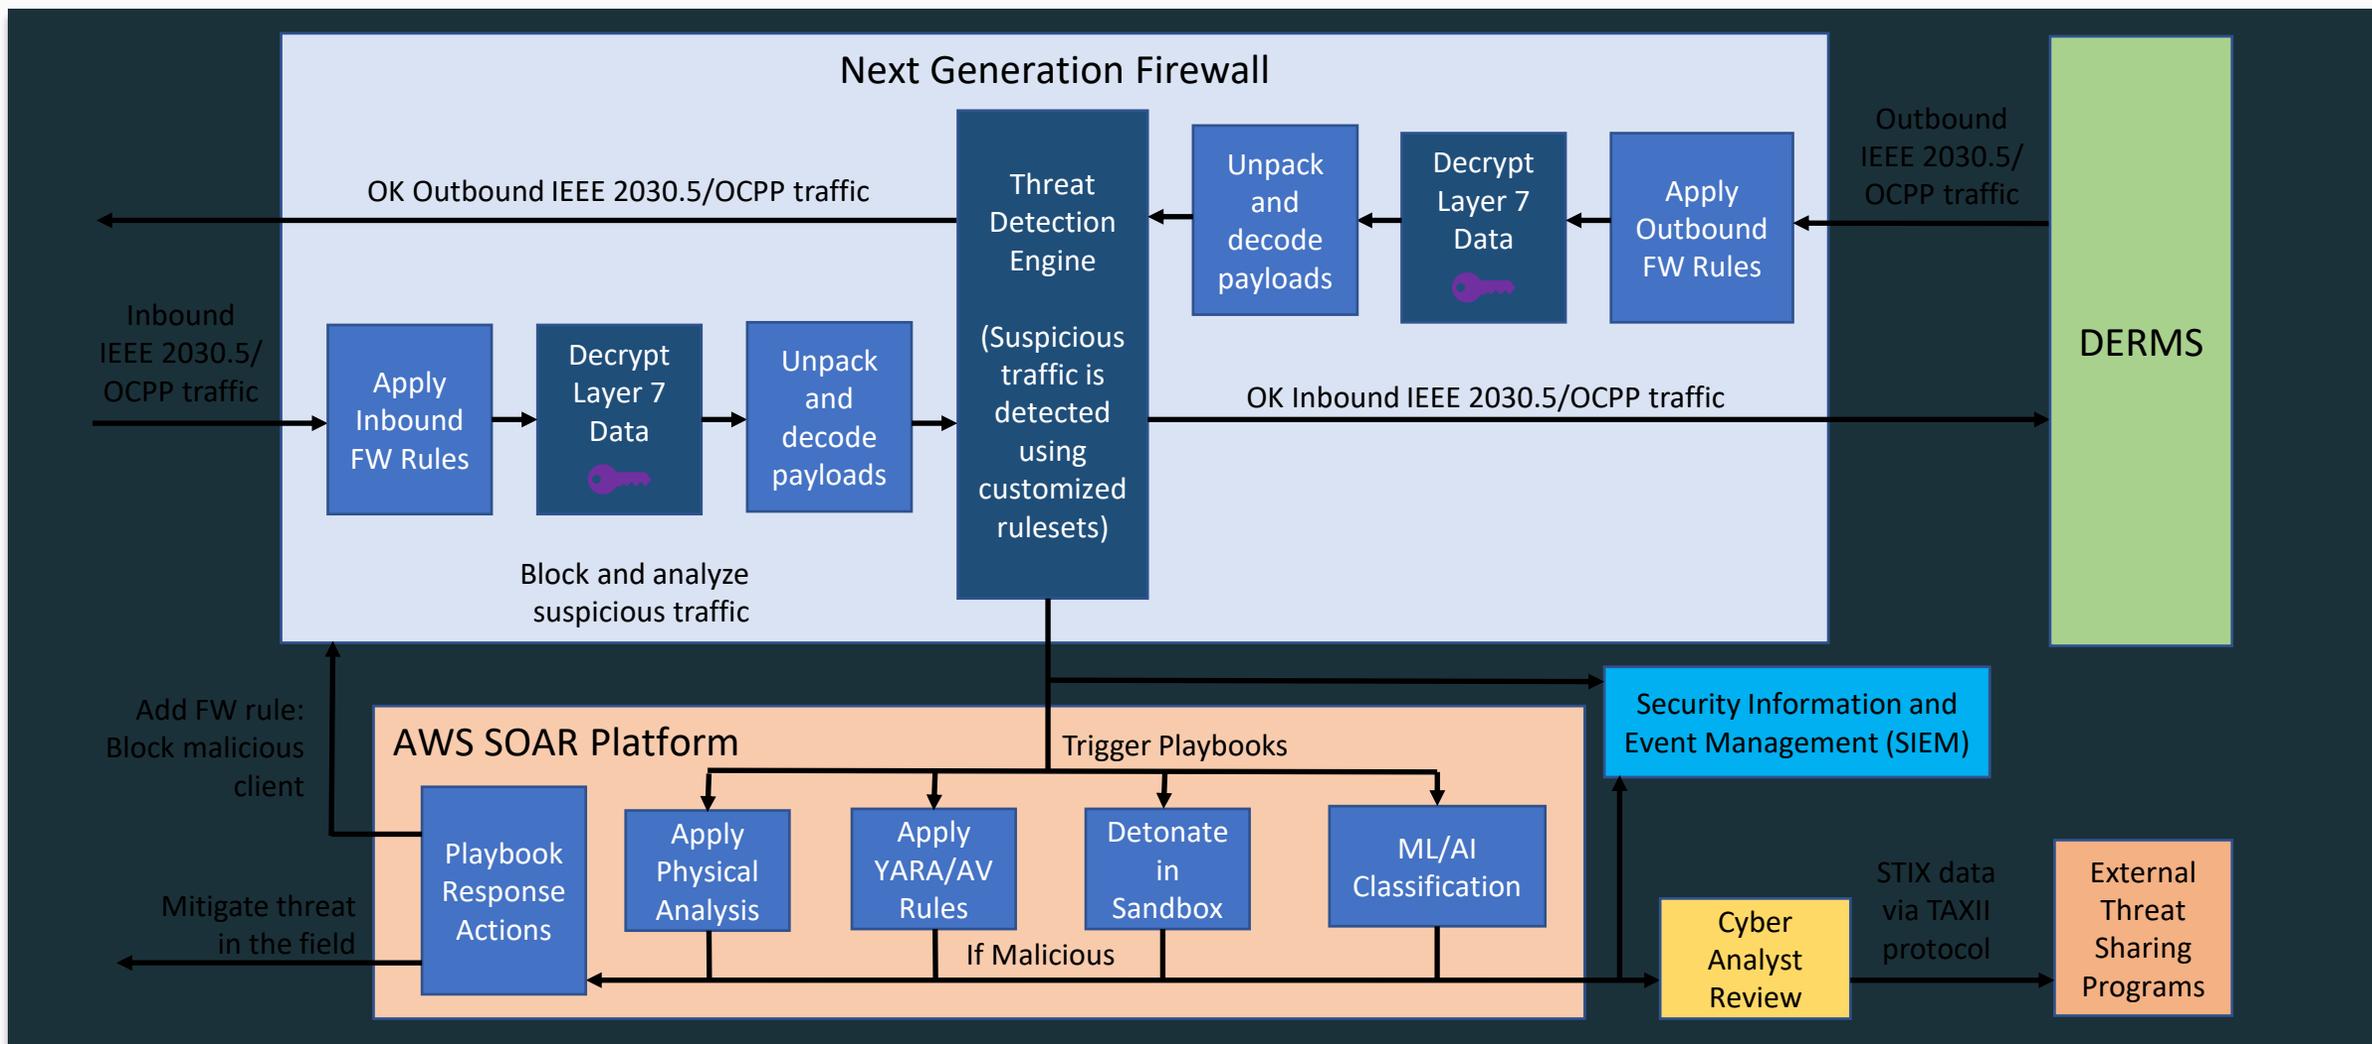
NGFW positions additional checks and security analysis before traffic reaches legitimate backend systems

SOAR technologies allows an **active approach** to mitigating misconfiguration, adversarial actions, and misbehavior of devices

An add-on approach allows utilities to adopt the system without requiring **large-scale modifications** to their DERMS/CSMS architecture

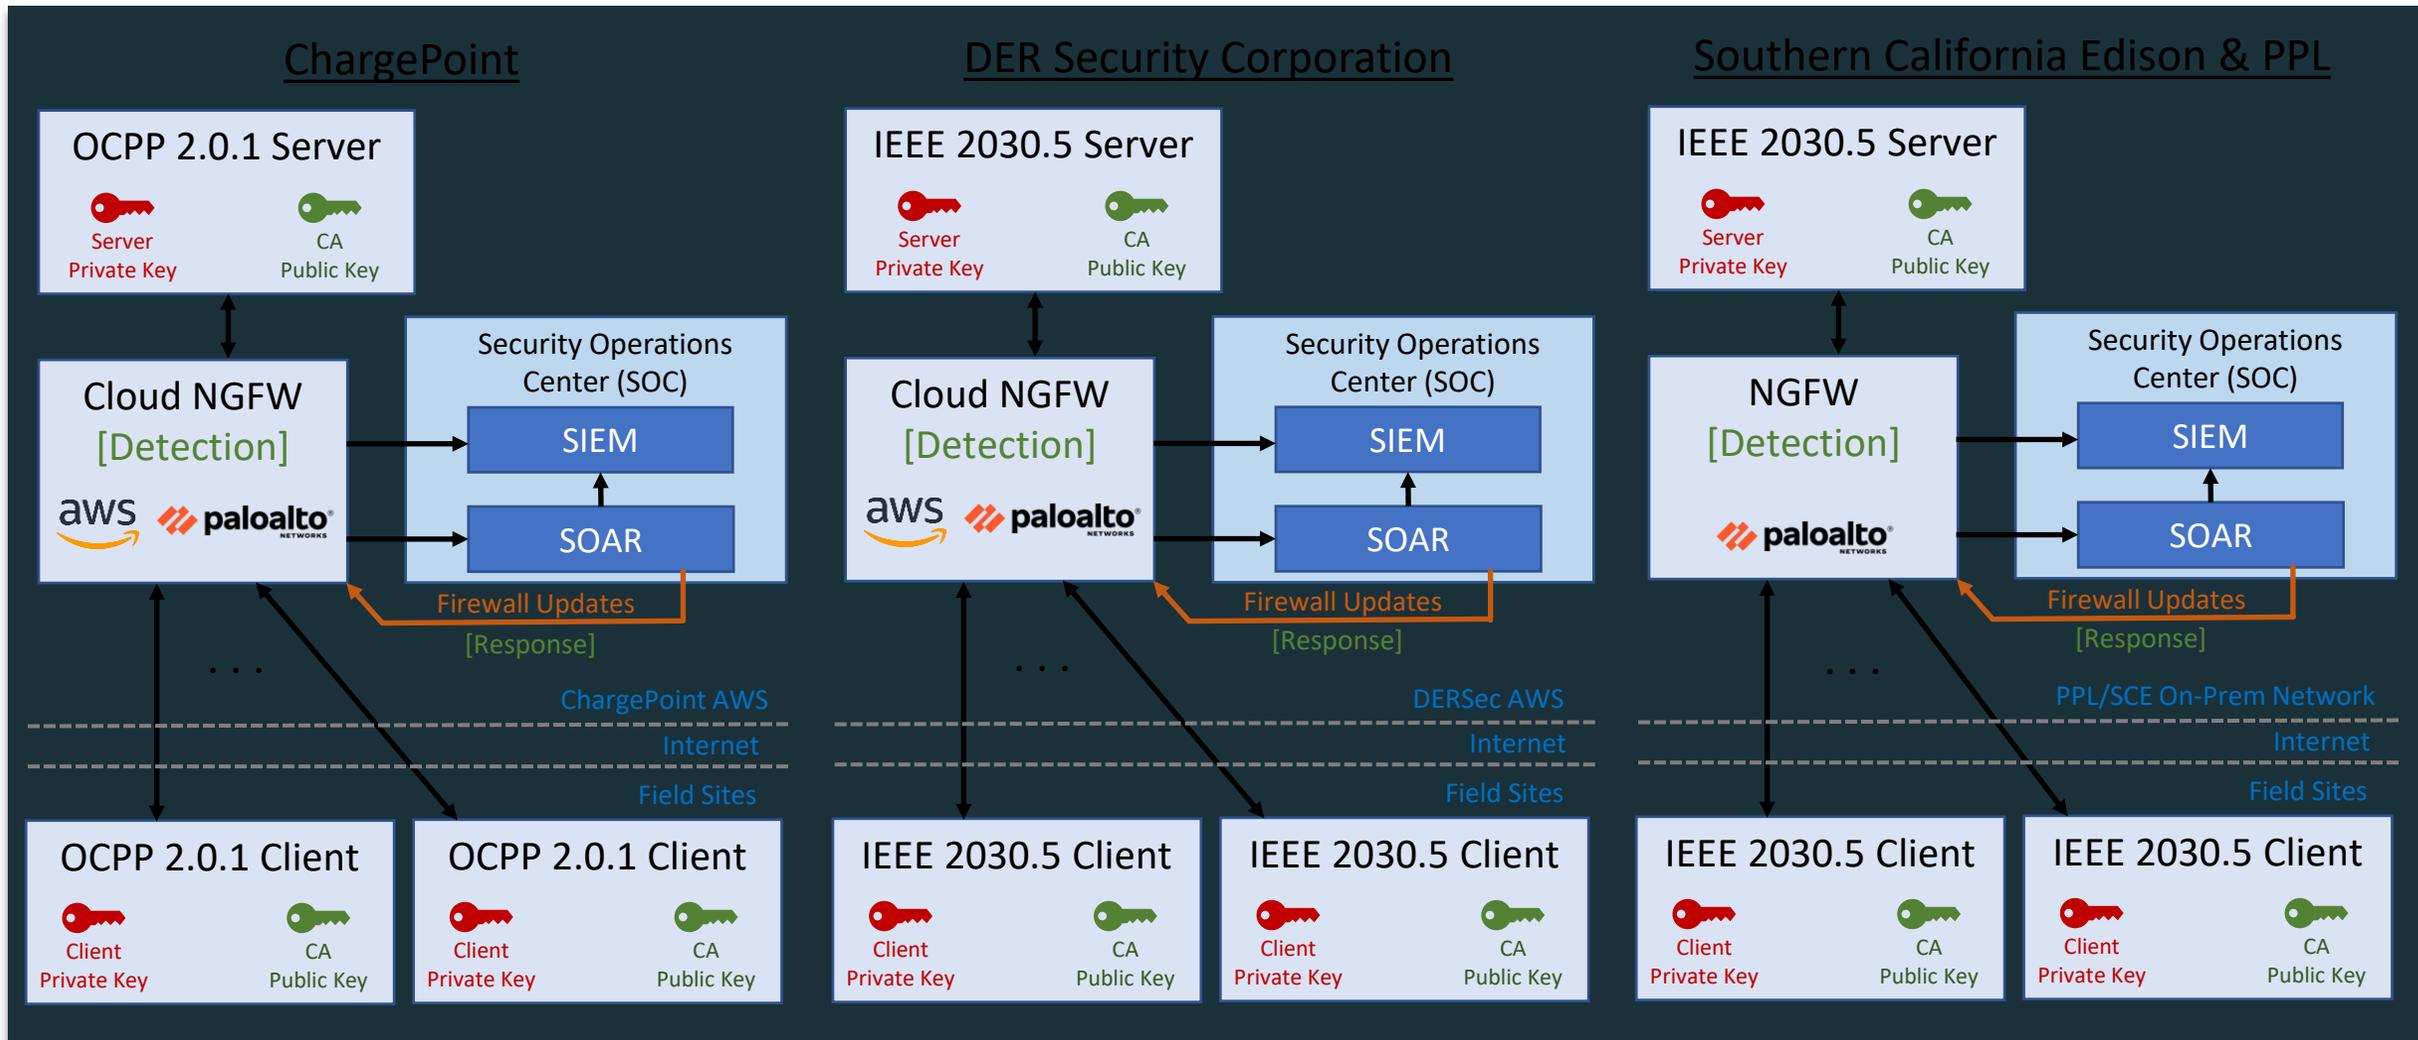


# Goaltender | Next Generation Firewall + AWS SOAR Platform





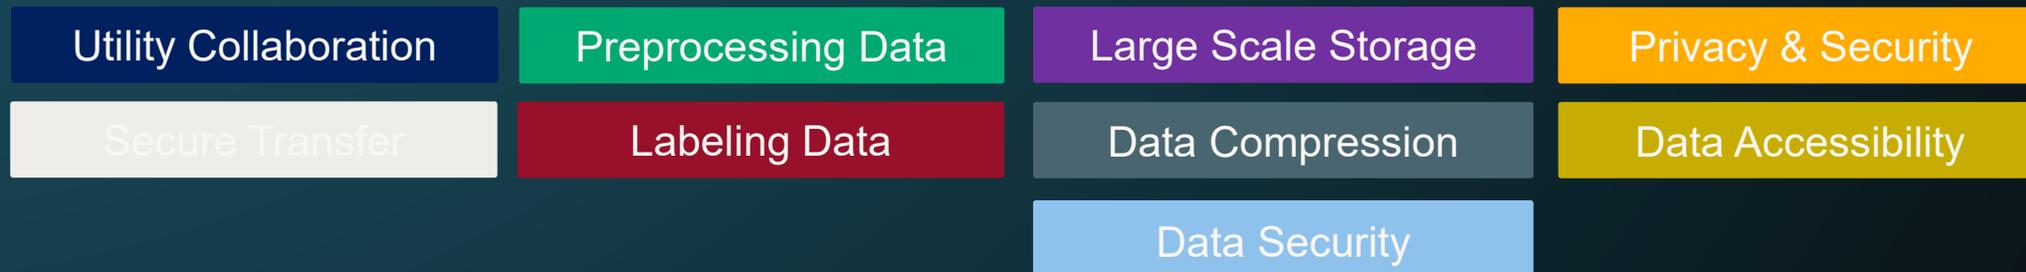
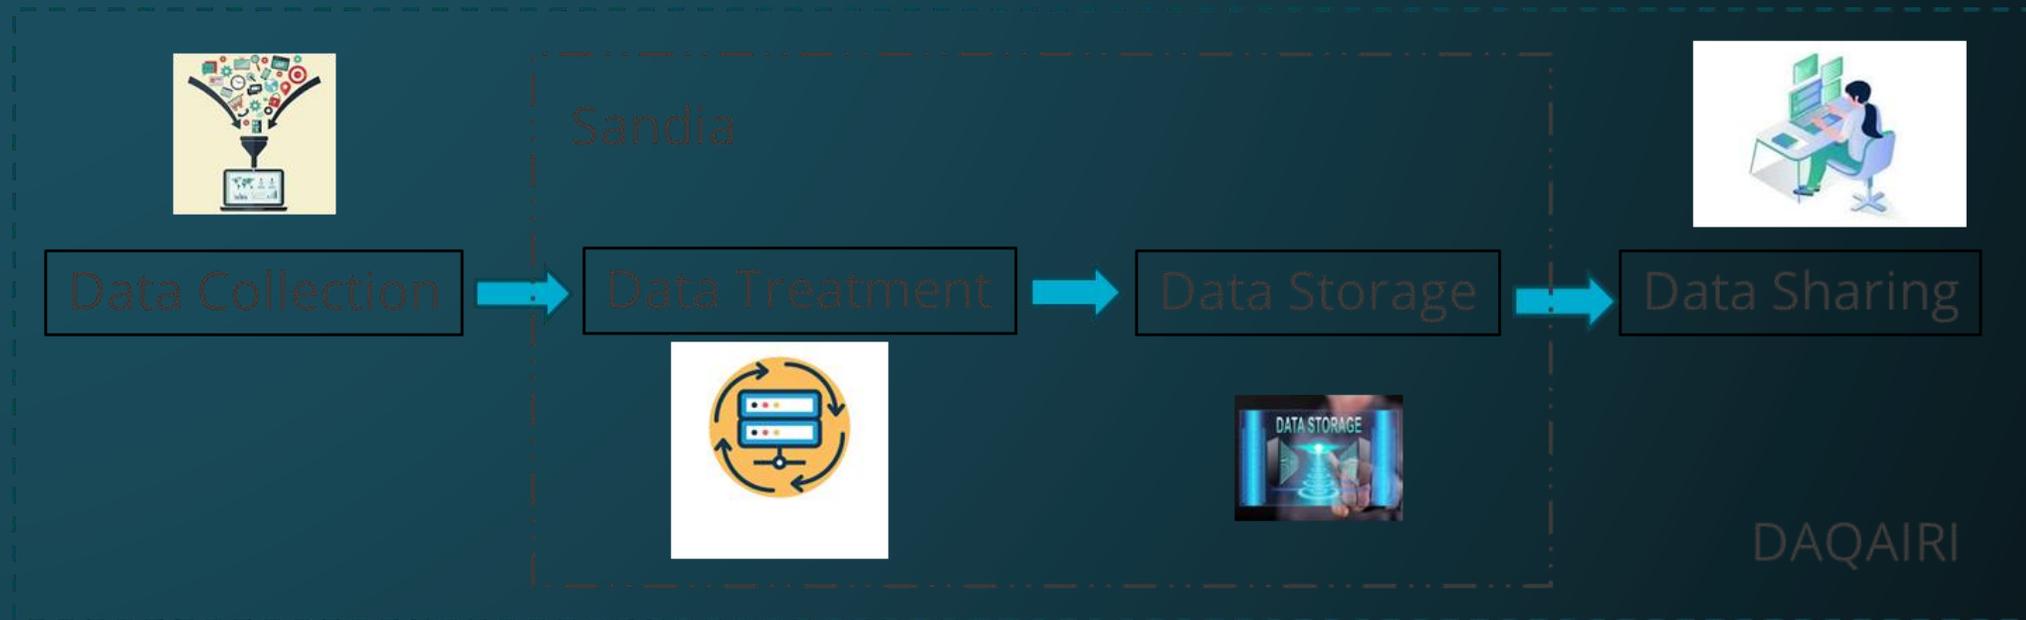
# Goaltender | Three Cases



# Main Takeaways

# RECAP of DAQAIRI's OVERALL Framework

DAQAIRI will provide a secure platform for utilities to share raw SCADA datasets with national labs, where the data will be treated, properly stored, and used for AI research and development.



# Conclusion: KEY Features of Daqairi's PLATFORM

- **Data Collection**
  - Engages with partners to identify and gather high-quality OT datasets essential for AI research.
- **Data Treatment**
  - Transforms raw utility data into clean, anonymized, and consistently formatted datasets through cleaning, privacy protection, and validation processes.
- **Data Storage**
  - Implements secure storage solutions for datasets, ensuring data integrity and long-term management while facilitating easy access for authorized users.
- **Data Sharing**
  - Establishes secure sharing protocols and agreements, enabling seamless access to datasets for national laboratories and promoting collaboration in AI research.



# DAQAIRI Platform's Goals

## Facilitating Collaboration

DAQAIRI's platform fosters strong partnerships between utility companies and national labs, promoting innovation and knowledge sharing in the energy sector.

## Comprehensive Data Management

DAQAIRI encompasses data collection, treatment, storage, and sharing, ensuring a streamlined process for researchers to access and utilize datasets effectively.

## Advancing AI Research and Development

DAQAIRI's infrastructure is designed to enhance AI research and model development within the DOE National Laboratories by providing access to high-quality utility datasets.