cognizant

Executive white paper series on enterprise physical AI autonomy

Grid intelligence without coordination

Why AI-driven decisions are failing at the system level in energy and utilities

Table of contents

Authors

Vijay Narayan

EVP, Global Head for Physical AI and Head of Manufacturing, Logistics, Energy & Utilities

https://www.linkedin.com/in/vijaynarayan1/

Dr. Samih Fadli

CTO for Physical AI and Chief AI Officer of Manufacturing, Logistics, Energy & Utilities

https://www.linkedin.com/in/samihfadli/

Executive summary

Energy and utilities enterprises are leaning on AI to manage grid complexity, improve forecasting and enhance operations across assets, field execution and customer engagement. While these systems deliver measurable improvements within individual domains, their impact remains fragmented at the enterprise level. Utilities are becoming more intelligent in isolated functions yet struggle to coordinate decisions across them. This coordination gap reflects a structural shift—affecting reliability, cost efficiency, safety and customer experience. Lack of shared decision context to reconcile trade-offs in real time leads to fragmented reasoning during high-stakes events such as wildfire mitigation, storm response and capital planning. The expansion of agentic AI intensifies this challenge by accelerating decision-making across workflows.

While automation improves speed and responsiveness, it can also amplify misalignment when systems act independently, creating an enterprise that is faster but not more coherent. This introduces risks, increases regulatory scrutiny, and limits ability to capture and reuse institutional knowledge.

Addressing this issue requires a shift from deploying isolated AI use cases to establishing a coordinated decision architecture. Ultimately, the future of grid intelligence will depend on the ability of these systems to act as one.

The coordination gap in enterprise AI for utilities

Energy and utilities enterprises are managing a level of operational complexity that is fundamentally different from the environment that many legacy decision systems were designed to support. As generation portfolios change, distribution networks become more dynamic. Electrification alters demand patterns, even as distributed energy resources are introducing new grid behaviors that utilities do not directly control. At the same time, customer expectations for reliability and transparency are increasing while field operations, asset planning, market participation and resilience programs increasingly interact as continuous systems rather than sequential processes. To manage this complexity, utilities are deploying AI across forecasting, work management, asset optimization, outage response, vegetation management, customer engagement and compliance workflows.

The benefits are real. Load forecasts are more precise. Distribution management systems can resolve switching operations across hundreds of thousands of nodes. Outage management systems correlate customer calls, AMI signals and field crew dispatch into coherent restoration plans within minutes of an event. Asset risk models are improving. Crew dispatch can be optimized. Customer communications can be personalized. Planning tools can evaluate more scenarios in less time. Yet these improvements often remain confined to individual domains. A system can become smarter at forecasting demand without understanding field crew constraints. An asset model can rank risk without reconciling customer impact. A work management system can optimize crew schedules without fully incorporating system-wide reliability priorities. The utility becomes more intelligent in pieces but not necessarily more coordinated as an enterprise.

This is the emerging coordination gap in enterprise agentic AI for energy and utilities. It is not simply a technology issue. It is a system-level decision problem. As the sector shifts from human-centered coordination supported by tools to AI-enabled coordination across many platforms, the absence of a shared enterprise decision context can produce faster local decisions while leaving system-level reasoning fragmented. The utility has deployed intelligence everywhere, but it cannot yet think as one system, at a moment when thinking as one system is the competitive and regulatory frontier.

The utility enterprise is becoming a network of decision engines

The modern utility no longer operates through a single decision center. It operates through a network of platforms and workflows that each make or influence decisions. Distribution management systems evaluate real-time constraints and execute switching operations. Outage management systems sequence restoration and coordinate with field crews. Distributed energy resource management platforms reconcile solar, battery, electric vehicle and demand response assets across service territories that were, a decade ago, managed as passive distribution networks. Asset management platforms prioritize replacement and maintenance against lifecycle economics. Work management systems assign crews and track field execution. Customer information systems manage communications, service commitments and billing complexity. Planning tools analyze capital allocation. Market and energy management systems optimize economic performance. Geographic information systems carry the topology of the grid at a fidelity that field engineers once could only approximate from paper maps. Increasingly, AI-enabled agents assist or automate parts of these workflows.

Each system has a legitimate mandate. Grid operations focus on stability and reliability, asset teams on risk reduction and lifecycle economics, customer teams on communication and service experience, and field operations on safety, crew availability and execution practicality. Finance teams reinforce capital discipline while regulatory teams ensure traceability and compliance. The problem is that high-stakes utility decisions require these mandates to be reconciled in context and current architecture rarely provides a persistent mechanism for doing so.

When decision intelligence is distributed across platforms, the enterprise can accumulate many accurate partial views. The challenge is that utility operations are deeply interdependent. A decision to defer maintenance on a transformer may affect outage risk, crew allocation, capital planning, customer reliability metrics and future storm response. A decision to prioritize one restoration sequence may improve customer minutes interrupted in one area while delaying service to critical load elsewhere. A decision to activate demand response may reduce immediate strain while changing customer behavior and market exposure. These are not isolated decisions. They are system decisions made in an environment where the cost of incoherence is measured in reliability, safety and public trust rather than in margin alone.

Scenario one

Wildfire season and the limits of fragmented decision intelligence

Consider a utility operating in a high fire-threat environment as conditions elevate toward red-flag status. Weather intelligence systems issue forecasts that the vegetation management system has already informed through its risk assessments. The distribution management system is evaluating which circuits are candidates for elevated protection. The asset management platform carries the inspection history that determines which segments warrant particular scrutiny. The public safety power shutoff (PSPS) decision-support tool is assembling the customer impact footprint. Each system is reasoning about the same fire risk through a different lens and each is producing a different recommendation.

The operations executive responsible for the shutoff decision convenes a call. The call takes three hours. In those three hours, conditions shift, the forecast updates, the vegetation risk reclassifies and half of the analysis the executive is hearing is already stale. The decision is made, and it is defensible, but no system in the utility can reconstruct the reasoning that underlay it at the moment it was made because the reasoning existed only in the minds of the people on the call. When the regulator asks later why the de-energization footprint was drawn where it was drawn, the utility reconstructs the answer by correlating timestamps across systems that did not know they were participating in the same decision.

Scenario one

This pattern is not a failure of any individual system. The forecasting tool is accurate, the vegetation risk model is sophisticated and the distribution management system is performing as designed. The failure is that no mechanism exists to reason across these systems at the speed and with the governance that the event demands. The enterprise lacks a shared reasoning framework for determining how ignition risk, customer impact, asset vulnerability and operational feasibility should be balanced in the moment and to retain the reasoning long enough to answer the regulatory questions that inevitably follow.

Scenario two

Storm response and the limits of fragmented agentic workflows

Storm response reveals the same problem under greater pressure. Weather models identify risk areas. Outage management systems organize restoration plans. Work management systems assign crews. Distribution management systems execute switching operations that isolate faults. Customer information systems communicate with affected customers. Materials and logistics platforms coordinate inventory. Executive teams need a consistent picture of readiness, exposure and response options. Each workflow may be AI-assisted, yet the enterprise still needs a coordinated interpretation of what the event means and what actions should be prioritized.

Fragmented architecture can generate many useful alerts without producing a unified response posture. Within the first hour of a major event, all four operational systems are reasoning from slightly different assumptions about the actual state of the grid because the latency between field events and system updates varies across the four platforms. A crew clears a fault that the outage management system still shows as active. The distribution management system reconfigures around an isolation point that the work management system has already scheduled a crew to inspect. The customer information system issues a restoration notification for a block that is still physically offline because the AMI data that would contradict the notification has not yet propagated through the reconciliation pipeline. Every system is doing exactly what it was designed to do. The composite behavior is a utility that communicates confidently with customers while its internal picture of the grid is several minutes behind reality.

Scenario two

The consequence is not merely operational inefficiency. In a utility environment, decision fragmentation can affect public trust, service reliability, safety and regulatory confidence. Every minute of reconciliation lag is a customer complaint, a trust erosion and a regulatory metric that moves in the wrong direction. The organization may know a great deal about the storm, the grid, its assets and its customers, yet still struggle to convert that knowledge into a coherent enterprise response. The limiting factor is not observability. It is coordinated decision intelligence.

Scenario three

Capital planning and work execution diverge

A quieter but equally consequential pattern appears in capital and asset decisions. Asset analytics identify a set of equipment with elevated risk. Capital planning models prioritize replacement based on long-term reliability value and budget constraints. Work management systems evaluate crew availability and outage windows. Customer systems identify communities with heightened sensitivity to service interruptions. Grid operations teams assess near-term operating constraints. Each system has a defensible view of priority, but the enterprise must decide which work should happen first, why and under what constraints.

In a fragmented environment, capital plans and executional realities often diverge. A project may rank highly in the planning model but be difficult to execute because outage windows are limited or field access is constrained. Another asset may have lower long-term risk but greater near-term operational exposure due to weather, load or customer criticality. A work order may be efficient from a crew-routing perspective but misaligned with broader resilience priorities. The enterprise needs to reason across planning, operations, field feasibility and customer impact as one decision context.

Scenario three

Without that shared context, the utility can overinvest in planning precision while under-performing in execution coherence. Leaders may approve capital programs that look optimized on paper but degrade when translated into field reality. Alternatively, field execution may optimize immediate constraints and slowly drift away from strategic asset priorities. Both outcomes reflect the same problem: Decision intelligence exists, but it is not unified across the lifecycle of the work.

The consequences compound as autonomy increases

As utilities embed AI across more decisions, the impact of fragmentation does not remain contained. It compounds over time, affecting resilience, speed of response, governance and the enterprise's ability to learn.

Resilience risk

Energy and utilities systems are coupled networks. A decision that improves one metric can create stress elsewhere. When decision systems optimize in isolation, small misalignments can accumulate into larger operational exposure. The enterprise may respond quickly to individual signals while missing the system-level implication of those signals in combination. Over the course of a single wildfire season or storm cycle, this accumulates into incremental exposure that compounds across risk, litigation and regulatory posture in ways that no individual post-event review fully captures.

Decision latency during high-consequence events

AI-enabled tools may accelerate analysis, but fragmented recommendations still require human reconciliation. Operators and leaders must determine which system to trust, how to weigh conflicting recommendations and when to escalate. This creates a paradox: The enterprise has faster analytics, but the decision process can remain slow because the architecture does not resolve conflict. In markets where regulators increasingly weight performance-based ratemaking frameworks, the gap between coherent and merely adequate response is a direct determinant of allowed return.

Governance and audit complexity

Utilities must be able to explain high-impact decisions, especially when they affect reliability, customer service, capital allocation or safety. Frameworks such as NERC CIP and IEC 62443 mandate sovereign architecture for critical infrastructure control systems—and state commissions in California and elsewhere—are asking utilities to demonstrate what they knew at the precise moment critical decisions were made. When the reasoning behind decisions is distributed across multiple platforms with different assumptions, workflows and records, the enterprise struggles to reconstruct a coherent decision narrative. Reconstruction is possible, but it is not evidence of coherent governance. Utilities that can demonstrate contemporaneous coherent governance across their intelligent systems occupy a meaningfully stronger regulatory position than utilities that cannot.

Institutional learning loss

A utility should learn from every outage, every miss, every demand event and every major field execution. It should understand not only what happened, but how decisions interacted across operations, assets, customers, crews and planning. In fragmented environments, the learning remains localized. The outage system learns about restoration. The asset system learns about equipment. The customer system learns about communications. The enterprise may not learn the full event as an integrated decision episode. The reasoning that weighed evidence during a consequential event lives in the minds of senior dispatchers, grid planners and emergency response coordinators who have worked through a generation of storms, fires, heat domes and unusual events. Those people are entering a generational retirement in an industry that is simultaneously facing its most operationally demanding decade.

A utility that cannot institutionalize the reasoning behind its most consequential decisions is a utility whose operational capability is bound by whatever its current workforce happens to remember.

Why more dashboards will not create system intelligence

Many utilities have responded to complexity by improving visibility. They build control rooms with richer displays, create executive dashboards, connect operational data sources and increase reporting cadence. These efforts are valuable, but visibility is not the same as system intelligence. A dashboard can show competing facts side by side. It does not automatically reconcile them into a governed decision.

Data centralization has similar limits. A unified data environment may improve analytics, but a utility still needs a shared decision context. What is the priority under this condition? Which constraint governs? What trade-off is acceptable? Which recommendation should be escalated? How should the outcome be remembered? These questions are not answered by data availability alone. They require an enterprise approach to coordinated reasoning. Utilities that have consolidated operational data into unified analytical platforms report that the analytical coherence of their data has improved substantially, while the decisional coherence of their enterprise has not. Analytics describe what happened on the grid. They do not arbitrate what should happen next.

The deeper requirement is to establish coherence across decision systems without eliminating the specialized platforms that already support the utility. Grid operations, asset management, customer systems, market systems and field operations each need their own capabilities. The opportunity is not to collapse them into one monolithic application. The opportunity is to allow them to contribute to a common enterprise understanding of conditions, decisions and outcomes.

Why agentic AI changes the utility operating model

In energy and utilities, agentic AI is not merely a digital productivity enhancement. It changes how decisions move across the enterprise. AI-enabled workflows can prioritize work orders, draft customer communications, recommend restoration sequences, assess asset risk, summarize field conditions and coordinate planning inputs. These actions sit at the intersection of reliability, safety, cost, customer trust and regulatory expectations. As a result, agentic AI introduces a new operating model challenge: The enterprise must ensure that decisions generated in one workflow remain coherent with decisions generated elsewhere.

Historically, utilities relied on procedural governance, experienced operators and function-specific control points to manage complexity. Those mechanisms remain important, but agentic AI can compress the time between insight and action. A system can identify asset risk, generate a work recommendation, initiate a planning update, inform a customer workflow and escalate a compliance implication faster than traditional coordination processes can absorb. The enterprise then faces a difficult question: Is speed producing better system decisions, or simply accelerating fragmented execution?

This issue becomes more important as utilities digitize more of the operating model. The same event can touch control room operations, field crews, customer service, asset engineering, finance and executive communications. If AI-enabled workflows reason from different assumptions, the utility may present one position to customers, another to operations and a third to capital planning. The issue is not simply inconsistent messaging. It is inconsistent institutional understanding, which in a regulated critical infrastructure environment becomes a regulatory and reputational exposure rather than a mere operational one.

The executive failure mode: Insight without authority

Utilities increasingly have access to insight, but insight does not automatically create authority. A forecasting model can identify risk and an asset model can rank vulnerability while a customer model can predict dissatisfaction, and a field system can identify execution constraints. The enterprise still needs a governed mechanism for determining which insight takes precedence under specific conditions. Without that mechanism, authority defaults to hierarchy, meeting cadence or whichever system is closest to the operational moment.

This creates management variability. In one event, operations may dominate. In another, customer impact may drive the decision. In another, capital discipline may prevail. Flexibility is necessary, but unstructured variability is costly. The enterprise should be able to explain why a decision was made, what trade-offs were considered, what evidence mattered and how the outcome will inform future decisions. Fragmented agentic systems make that explanation harder as more decision logic moves into software workflows.

The same ambiguity carries into regulatory posture. Every major utility operates under frameworks that require the utility to reconstruct—often months or years after the fact—the reasoning behind operational decisions. Why was this circuit energized before that one? Why was this asset prioritized for replacement over that one? Why was this customer's service restored in this order? Why did the shutoff footprint extend to this substation and not that one? Utility regulatory teams describe the reconstruction work as the single largest source of unbudgeted effort in the regulatory function, and they describe it as increasingly adversarial, because regulators are beginning to recognize that the reconstruction itself is evidence that coherent governance across systems did not exist when it mattered.

For executives, this means that the AI agenda cannot be limited to use-case deployment. It must include decision architecture. The enterprise needs a clear view of where AI-generated recommendations enter business processes, how conflicts are resolved, how policies are applied and how lessons are retained. The future utility will not only be more digital. It will need to be more institutionally coherent.

What maturity begins to look like

A mature utility would not expect every system to become a universal decision engine. Specialized platforms will remain essential. Distribution management should remain strong at grid state. Outage management should remain strong at restoration. Asset management should remain strong at lifecycle economics. Customer systems should remain strong at experience and communication. The maturity shift is to ensure that these systems do not reason in isolation when their decisions interact. The utility needs a way to preserve specialization while creating coherence across domains.

In practical terms, this means that consequential decisions should be interpreted against a shared enterprise context that reflects the utility's safety, reliability and regulatory obligations. A PSPS footprint decision should not only be evaluated against weather and vegetation risk, but against customer vulnerability, field feasibility, regulatory precedent and the reasoning from similar prior decisions.

A capital program should not only be evaluated against risk models, but against execution realities, operational constraints and the lessons from prior program cycles. The utility should be able to understand the decision episode as one event, even when multiple systems contribute to it, and retain that understanding long enough for the next similar event to benefit from it. Critically, this capability must preserve the human override that every utility requires and that every regulator will continue to require for the foreseeable future, while applying the institution's rules consistently in the moments when human override is not invoked.

The way forward

Looking forward, the evolution of grid intelligence will depend less on deploying additional AI into individual domains and more on unifying the intelligence that already exists. Utilities will need a persistent and governed way to coordinate decisions across operations, assets, customers, field execution, planning and finance. This capability must respect the specialized logic of each domain while allowing the enterprise to reason across domains when the stakes require it.

Organizations that develop this capability will be better positioned to manage volatility, improve resilience and accelerate response without sacrificing governance. They will be able to understand how a decision in one part of the utility affects the broader system. They will learn from events as integrated episodes rather than as isolated platform records. They will meet regulatory expectations with contemporaneous evidence of reasoning, rather than with post-hoc reconstruction. And they will carry the operational experience of their most senior workforce forward as institutional capability, rather than as individual memory that retires with the people who hold it.

The future utility will not be defined only by how many intelligent systems it has deployed. It will be defined by whether those systems think as one. In an industry where the cost of incoherent decisions is measured in lives, reliability and regulatory standing, that is not a question that can be deferred. It is the defining question of operational capability for the remainder of this decade.

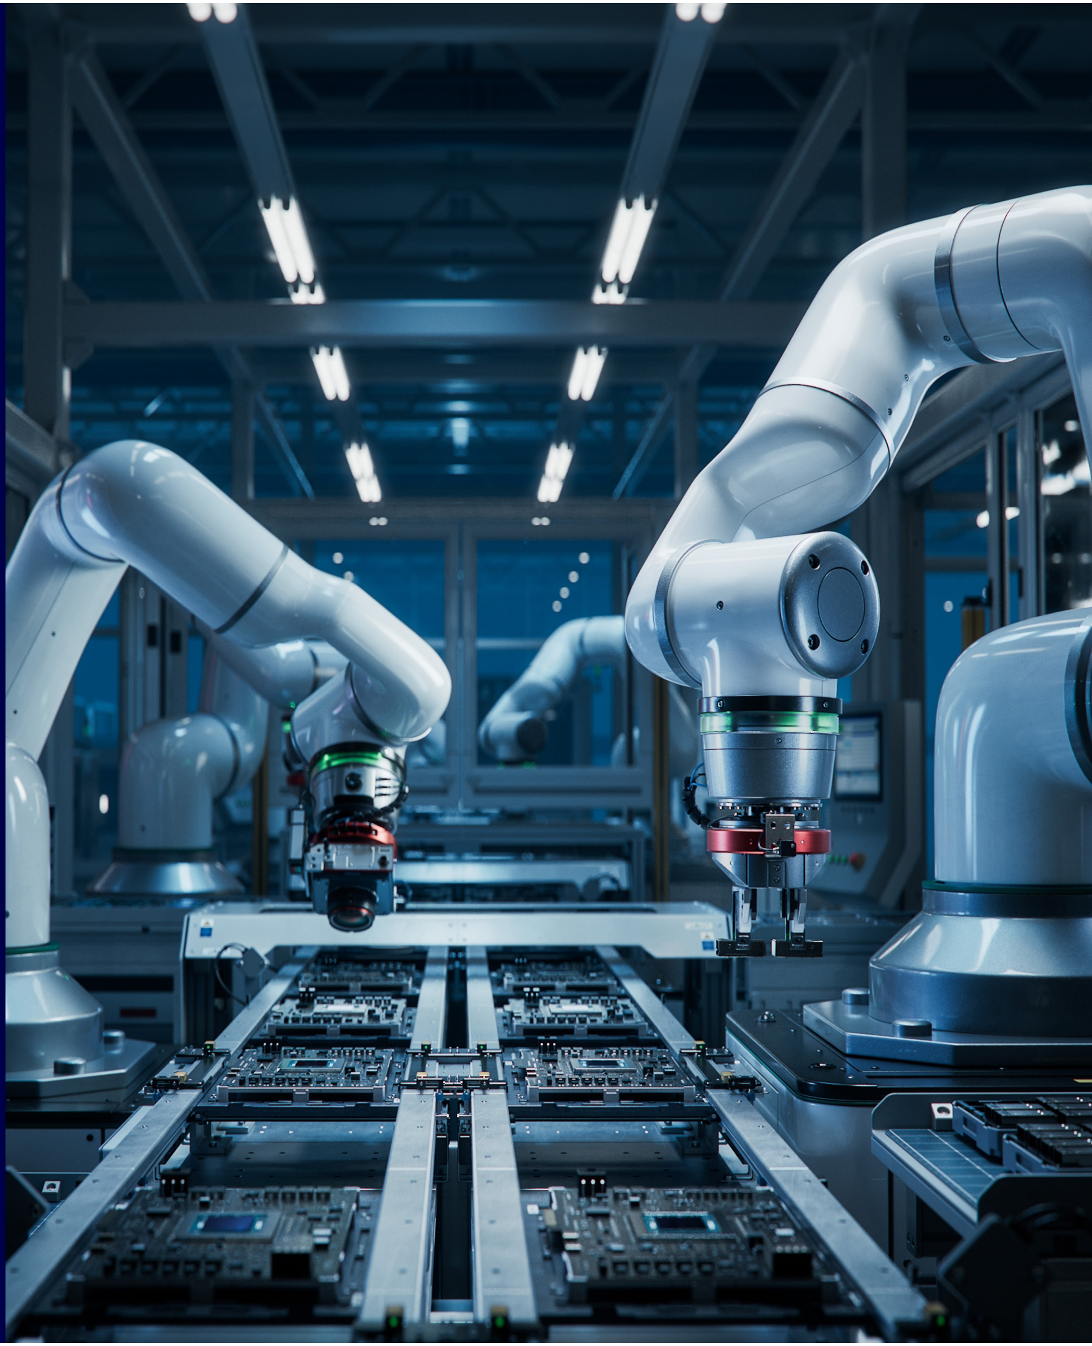

cognizant

Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

World Headquarters

300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Tel: +1 201 801 0233

European Headquarters

280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office

Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters

1 Fusionopolis Link, Level 5
NEXUS@One-North, North Tower
Singapore 138542
Tel: +65 6812 4000

WF4823550