cognizant

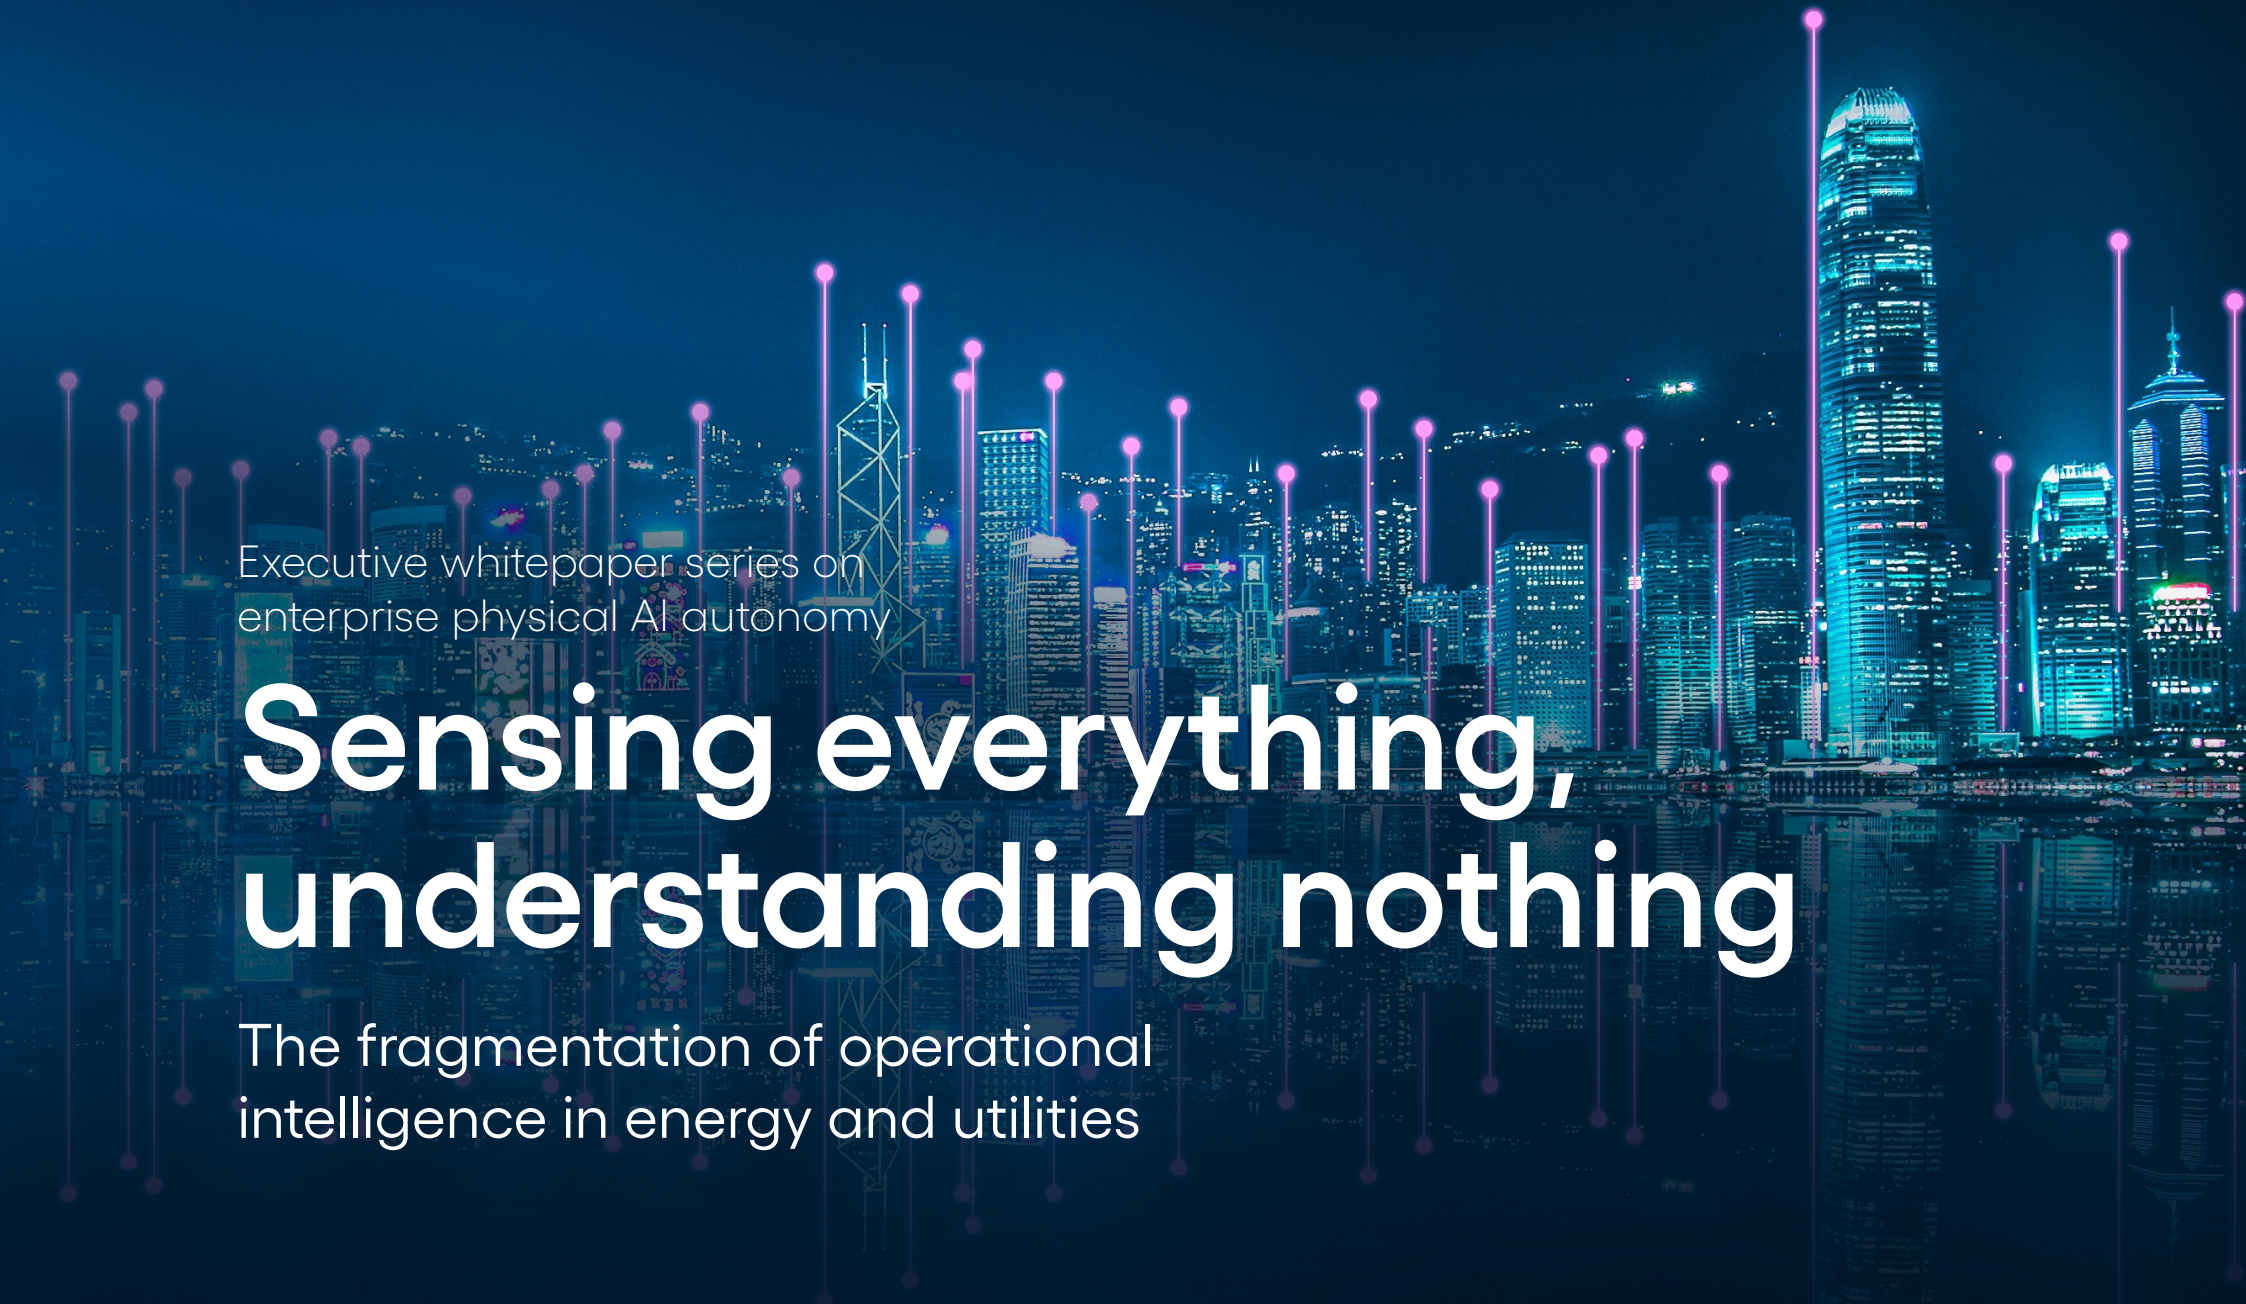

Executive whitepaper series on
enterprise physical AI autonomy

Sensing everything, understanding nothing

The fragmentation of operational intelligence 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 entering an era of near-continuous observability. Grid telemetry, smart meters, line sensors, substation monitoring, pipeline instrumentation, imaging and weather analytics are expanding visibility into physical conditions that were previously delayed, partial or invisible.

But observability is not understanding. The enterprise can sense voltage variation, asset stress, vegetation proximity, soil movement, pressure, flow and demand without forming a coherent interpretation of how signals interact. Intelligence often remains confined to separate platforms.

Four scenarios illustrate the gap:

1. Environmental stress events where weather, load, asset vulnerability, vegetation exposure and crew constraints must be synthesized under time pressure
2. Pipeline integrity decisions where inspection, fiber sensing, cathodic protection, operations data and terrain conditions update faster than prioritization can keep pace
3. Substation investigations that require weeks of reconstruction across multiple systems and do not become reusable intelligence
4. Divergence between control room signals and field realities that complicate safe, effective execution

The consequences are systemic: reduced anticipatory capability, weaker prioritization under resource constraints, slower and less consistent response and fragmented learning that increases regulatory and public exposure. Maturity requires moving from isolated sensing to integrated interpretation that connects signals, records evidence trails and turns events into institutional learning.

The observability paradox in energy and utilities

Energy and utilities enterprises are entering an era of near-continuous observability. Grid telemetry, smart meters, line sensors, substation monitoring, pipeline instrumentation, thermal imaging, LiDAR, weather analytics, field inspection data, and asset health systems are expanding the enterprise's view of the physical environment. Transmission corridors are observed at frequencies that capture vegetation encroachment, conductor sag, and structural anomalies long before they announce themselves in service. Substations carry continuous acoustic and partial discharge monitoring, dissolved gas analysis, and thermal imaging that collectively represent one of the most sophisticated asset observability postures in any industry. Pipeline networks are traced by in-line inspection tools, distributed fiber optic sensing, cathodic protection monitors, and aerial leak detection that together observe thousands of miles of buried infrastructure continuously. In many cases, the organization can observe conditions that were previously invisible, delayed, or available only through manual inspection. This has enormous value. It improves awareness, supports faster detection, and enables more proactive operations.

Yet observability is not the same as understanding. A utility can sense voltage variation, asset stress, vegetation proximity, soil movement, temperature, wind, pressure, flow, and customer demand without forming a coherent interpretation of how those signals interact. Physical AI systems can detect anomalies and forecast localized outcomes, but if their intelligence remains confined to separate platforms, the enterprise does not develop a system-level understanding of its infrastructure. The patterns that matter most in critical infrastructure, which are the cross-sensor, cross-asset, and cross-time patterns that emerge when multiple conditions correlate, are precisely the patterns that fragmented physical intelligence cannot see.

This is the fragmentation of operational intelligence in energy and utilities. The enterprise is sensing more of the physical world, but the meaning of those signals is distributed across systems that do not naturally reason together. The utility has deployed intelligence everywhere, yet it cannot perceive as one system. The result is a paradox that, in critical infrastructure, carries consequences beyond operational inefficiency: the organization may have unprecedented visibility into its infrastructure while still lacking the coordinated understanding required to anticipate, prioritize, and act at the system level.

The physical system is more observable, but more complex

The energy and utilities operating environments are becoming more dynamic. Electricity networks are shaped by distributed resources, two-way power flows, electrification, changing weather patterns, and aging infrastructure. Gas, water, and pipeline networks face similar pressures from asset age, environmental exposure, demand variability, and public expectations for reliability and safety.

Physical systems that were once managed through relatively stable operating assumptions now require continuous interpretation. Physical AI is being deployed because this complexity cannot be managed through traditional inspection and monitoring alone. Sensors detect equipment behavior. Imaging systems identify physical risk. Predictive models estimate asset degradation. Field data captures conditions across distributed networks. Environmental analytics assess external threats. These systems increase the enterprise's ability to sense the operating environment in real time or near real time.

The challenge is that each sensing domain operates with its own logic and context. Grid monitoring systems understand electrical behavior, asset health systems track equipment degradation, vegetation management systems assess proximity and growth, and weather models analyze environmental movement. Field inspection workflows capture localized conditions, while substation systems focus on signals such as dissolved gas and acoustic signatures, and pipeline systems monitor factors such as wall thickness and strain. The physical system, however, does not operate in these discrete categories. It functions as a single, interdependent environment where conditions continuously influence one another.

Operational intelligence must therefore move beyond isolated interpretation, recognizing patterns that emerge across systems and over time rather than within individual domains alone.

Scenario one

Environmental stress across a distributed grid

Consider a regional utility facing a period of severe environmental stress as conditions elevate toward red-flag status. Weather models indicate rising temperatures and shifting wind conditions. Demand models forecast higher load. Transmission assets operate under thermal stress. Distribution equipment in several areas shows elevated risk. Field data indicates vegetation exposure in specific corridors. Customer demand response programs may help reduce strain, but their impact varies by region and participation behavior. Satellite imagery provides thermal signatures at intervals that have reduced dramatically over the past five years. Camera networks observe fire-prone terrain continuously. Grid sensors register fault currents, voltage excursions, and momentary events with millisecond fidelity.

Each system sees part of the event. Weather analytics interpret environmental risk, grid telemetry reflects operating conditions, asset models assess equipment vulnerability, and vegetation systems evaluate physical exposure. Customer systems see demand flexibility, while field operations account for crew constraints and access limitations. All these perspectives matter. The operational risk emerges from their interaction, and the utility is observing the same landscape through independent systems that do not share a common operational picture.

In a fragmented environment, the utility may respond to these signals separately. Asset teams prioritize equipment, grid teams manage load, and vegetation teams assess exposure while field teams manage crews and customer teams activate programs. The organization can be busy, informed, and responsive while still lacking a unified understanding of the full system condition. When conditions align toward elevated ignition probability, the utility's operational leadership is making decisions based on the synthesis of information that no system in the utility actually synthesizes. The synthesis happens in human minds, under time pressure, with information arriving from sources that update at different cadences and describe the same phenomenon in different terms. When the decision is made, it is defensible. But when it needs to be justified to the regulator, the answer must be reconstructed, because the reasoning that weighed the evidence existed only at the moments it was applied. The most important question is not whether each signal is accurate. It is whether the enterprise can understand what the combined signals mean before the event escalates.

Scenario two

Pipeline integrity and the limits of localized intelligence

A similar pattern appears in pipeline and distributed infrastructure operations. A pipeline operator responsible for thousands of miles of buried infrastructure deploys in-line inspection tools that characterize wall thickness, corrosion, and geometric anomalies. Distributed fiber optic sensing observes strain, temperature, and acoustic events along the entire asset. Cathodic protection monitoring characterizes the electrochemical environment that drives corrosion progression. Pressure and flow instrumentation characterizes the operational conditions that interact with structural integrity. Aerial leak detection provides periodic observation of the surface environment above the pipeline. Maintenance records may show prior interventions in the same segment. Weather and terrain conditions may affect access and repair planning.

Each signal can be evaluated independently, but the operational decision depends on synthesis.

Is the pressure variation a transient condition or an early indicator of integrity risk? Does the flow anomaly correlate with weather, terrain movement, or equipment history?

Should the organization dispatch inspection crews, adjust operating parameters, isolate a segment, or continue monitoring?

These decisions involve safety, continuity, cost, field feasibility, and customer impact.

The integrity decisions that matter most, including where to prioritize inspection, where to execute integrity digs, where to re-rate, and where to replace, require the fusion of all observation modes into a unified representation of asset condition across the entire network. Pipeline integrity teams perform this fusion today through a combination of specialist judgment and custom analytical workflows that do not scale to the frequency at which the observations update. The result is a prioritization process that lags the observation cadence, which means the utility is making integrity decisions based on a picture of the asset that is weeks or months behind what its sensors actually know.

When physical intelligence is fragmented, the enterprise may either under-react because no single system crosses an escalation threshold or over-react because teams cannot confidently interpret the combined risk. Both outcomes are costly. Under-reaction increases exposure. Overreaction consumes resources and disrupts operations. The higher-value capability is to understand the physical environment as an integrated system so that action is proportionate, timely, and explainable to the regulators and communities that depend on the infrastructure.

Scenario three

The substation event that required three weeks to reconstruct

A significant event occurs at a substation.
The sequence of events record captures electrical behavior in high resolution. The dissolved gas analyzers register a signature that may or may not correlate with the electrical event, while the acoustic monitoring system captures a pattern that may or may not indicate incipient failure in a specific component. Thermal imaging from the most recent inspection campaign, conducted six weeks earlier, documents conditions that may or may not be relevant, and the maintenance history carries records of work performed that may or may not have contributed. The engineer responsible for investigating the event spends three weeks correlating across seven systems to produce an explanation that is defensible but partial, because the correlation that would produce a complete explanation requires a cross-sensor temporal analysis that no system in the utility's current architecture performs.

The investigation reaches a conclusion.
The conclusion is accepted. The next similar event, at a different substation, requires the same three-week reconstruction from scratch, because the prior reconstruction is documented as a report rather than captured as institutional intelligence that the next investigation can build on. The utility has observed everything it needed to observe. It simply cannot assemble what it observed into the sustained institutional understanding that would allow the enterprise to anticipate rather than reconstruct.

Scenario four

Field reality and control room intelligence diverge

A persistent challenge in energy and utilities is the gap between centralized visibility and field reality. Control rooms may have increasingly detailed operational data, while field crews encounter conditions that are more complex, constrained, or ambiguous than the models suggest. Access limitations, local hazards, equipment configuration, terrain, weather, and crew availability all shape what can actually be done. Physical AI systems may detect and prioritize risk, but the enterprise must still connect those signals to field execution in a coherent way.

This gap becomes consequential when the control environment and field environment produce different interpretations of urgency. A monitoring system may identify an anomaly that appears high priority. Field data may indicate that access is limited or that another nearby condition changes the appropriate response. A planning system may recommend a sequence that looks efficient centrally but is difficult to execute safely or effectively on site. The enterprise must reconcile physical intelligence from sensors with physical intelligence from crews.

If those perspectives remain fragmented, the organization risks either delaying action while it reconciles conflicting information or executing a plan that underestimates field constraints. The highest-value operating model is one in which control room intelligence and field intelligence continuously inform each other, creating a more accurate understanding of the physical system than either could provide alone.

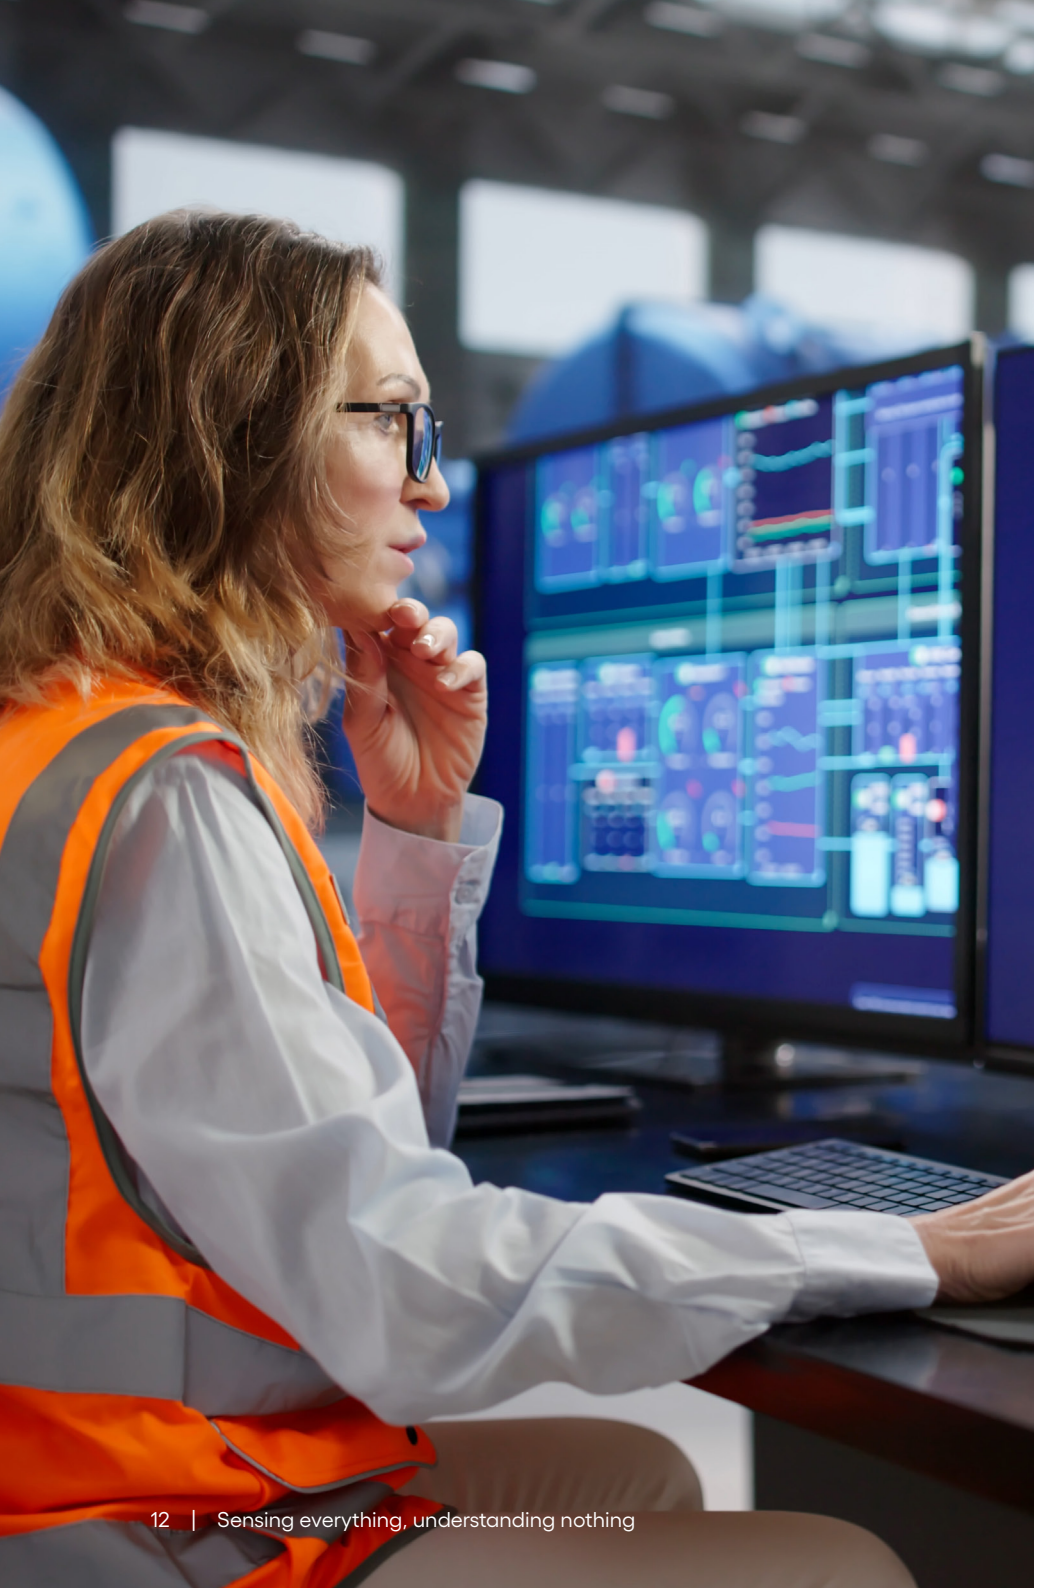

The consequences are systemic

Fragmented intelligence does not fail at the point of sensing, it fails in how signals are connected, interpreted, and acted on. Across utility operations, this gap alters how risks are recognized, prioritized, and managed, often shaping outcomes as much as the events themselves.

Reduced anticipatory capability

Many high-impact events are not caused by one signal crossing one threshold. They emerge from the interaction of conditions. Equipment stress, environmental movement, demand variability, and field constraints can combine into a risk pattern that no single system detects as decisive. Fragmented intelligence limits the enterprise's ability to identify these patterns early. Utilities that have retrospectively analyzed their response to major events consistently identify observational gaps, not sensing gaps but the gaps between sensing systems, as contributors to outcomes that were materially worse than they should have been. The contribution is rarely the primary cause, and it is rarely attributed as such in the public record. It appears instead as the difference between a response that integrated what the utility knew and one that acted on what individual systems happened to be saying.

Weaker prioritization

Utilities operate under constrained resources. Crews, capital, replacement parts, inspection capacity, and operating flexibility are limited. When signals are interpreted separately, prioritization becomes more difficult. Teams may optimize their own queues while the enterprise struggles to determine which actions matter most to system resilience and customer outcomes.

Slower and less consistent response

In a high-pressure operating environment, teams must assemble evidence from multiple systems before deciding. This creates delay and variability. Different regions may interpret similar patterns differently. Experienced personnel may compensate, but institutional dependence on informal expertise creates risk as the workforce changes and operational complexity increases. In critical infrastructure, where the senior workforce is entering a generational retirement, the inability to institutionalize the cross-sensor, cross-asset interpretations that this workforce has learned is not a future concern. It is a present exposure that becomes apparent in the next significant event.

Fragmented learning and regulatory exposure

Every physical event should improve the enterprise's understanding of its infrastructure. A near miss should refine risk interpretation. Failure should improve preventive decision-making. A successful intervention should inform future prioritization. In fragmented architectures, learning remains distributed across monitoring tools, field reports, asset records, and operational systems. The enterprise may document the event, but it does not necessarily convert it into persistent system-level intelligence. Critical infrastructure operators are also subject to some of the most demanding traceability and governance frameworks in any industry. NERC CIP, IEC 62443, DOT pipeline safety regulations, state commission oversight, and the increasingly integrated frameworks that link these requirements demand that the utility demonstrate not only that it sensed its infrastructure but that it understood what it was sensing and acted on that understanding coherently. The reconstruction of this understanding, from fragmented sensing systems, is inadequate to the regulators' demands.

Why visibility alone can create a false sense of readiness

The expansion of observability can create confidence that the enterprise is prepared. Control rooms display more signals. Dashboards show richer asset and environmental data. Inspection data becomes more frequent. Predictive models produce more alerts. Leaders may reasonably conclude that the organization has a clearer view of risk than ever before.

That conclusion may be true and incomplete. A clear view of individual signals does not guarantee a coherent interpretation of system behavior. In fact, more signals can increase cognitive burden if they are not organized through a shared operational context. Teams may spend more time interpreting, filtering, and reconciling alerts. The enterprise may become better at detecting events than at understanding what those events mean in combination.

The next frontier is therefore not only better sensing. It is better sense-making. Utilities need to understand how physical signals relate to each other, how risk evolves across assets and geography, how decisions in one domain affect conditions elsewhere, and how the outcomes of actions should inform future operations. This is the difference between observability and operational intelligence. Operators that have consolidated sensor data into unified analytical platforms report that the analytical coherence of their data has improved substantially, while the perceptual coherence of their infrastructure has not. Analytics describe what the sensors observed. They do not constitute the utility's real-time understanding of what its physical infrastructure is actually doing.

The executive failure mode: Too many signals, too little sense-making

As sensing expands, executives may face an unexpected challenge: the enterprise can generate more warnings than it can interpret. Alert volume rises, dashboards proliferate, risk scores evolve, and model outputs compete for attention. Teams must determine which signals matter, which are noise, which are connected, and which require action. This is not merely a user-interface problem. It is a problem of institutional sense-making.

A fragmented environment often forces leaders to manage by escalation. The most visible alert, the most urgent complaint, the most experienced voice, or the most conservative interpretation can dominate the decision. That may be appropriate in some cases, but it is not a scalable model for an increasingly instrumented infrastructure. As the physical system becomes more dynamic, the organization needs a more disciplined way to synthesize signals and understand their relationships.

The cost of weak sense-making is not only operational. It affects capital allocation, resilience strategy, public confidence, and workforce productivity. If the enterprise cannot consistently identify which conditions are truly linked, it may invest in the wrong interventions, over-prioritize visible risks, under-prioritize emerging ones, and exhaust teams with avoidable reconciliation work. In critical infrastructure, where the consequences of incoherent perception are measured in lives, assets, and public trust, this is not a productivity concern. It is an institutional exposure that accumulates until an event reveals it.

What maturity begins to look like

A mature utility would treat physical intelligence as a shared institutional asset rather than a collection of monitoring outputs. Observations from grid telemetry, asset systems, environmental models, field inspections, pipeline data, substation monitoring, and customer demand would contribute to a common understanding of system behavior. This does not mean all decisions become centralized. It means distributed teams operate from a more coherent interpretation of the physical environment.

In this model, the enterprise becomes better at recognizing patterns that span domains. Environmental conditions can be understood in relation to asset vulnerability. Asset behavior can be interpreted against demand and operational history. Field constraints can inform central prioritization. The outcome of each intervention can improve future interpretation. Over time, the enterprise develops a richer understanding of how its infrastructure behaves under stress, variation, and change. The requirement is not to replace the rigorous engineering and safety frameworks that already govern individual systems, because those frameworks are working and must continue to do so. The requirement is to provide an institutional layer above them that ensures their combined interpretation is coherent, that events are recorded in a unified evidence trail, and that the enterprise can demonstrate consistent understanding across its physical infrastructure in a way that regulators, the public, and the utility's own workforce can trust.

For leaders, the strategic question becomes whether the organization is building a permanent capacity for system-level understanding or merely adding new layers of sensing. The former creates resilience and institutional learning. The latter can create more visibility without solving the underlying challenge of fragmented operational intelligence.

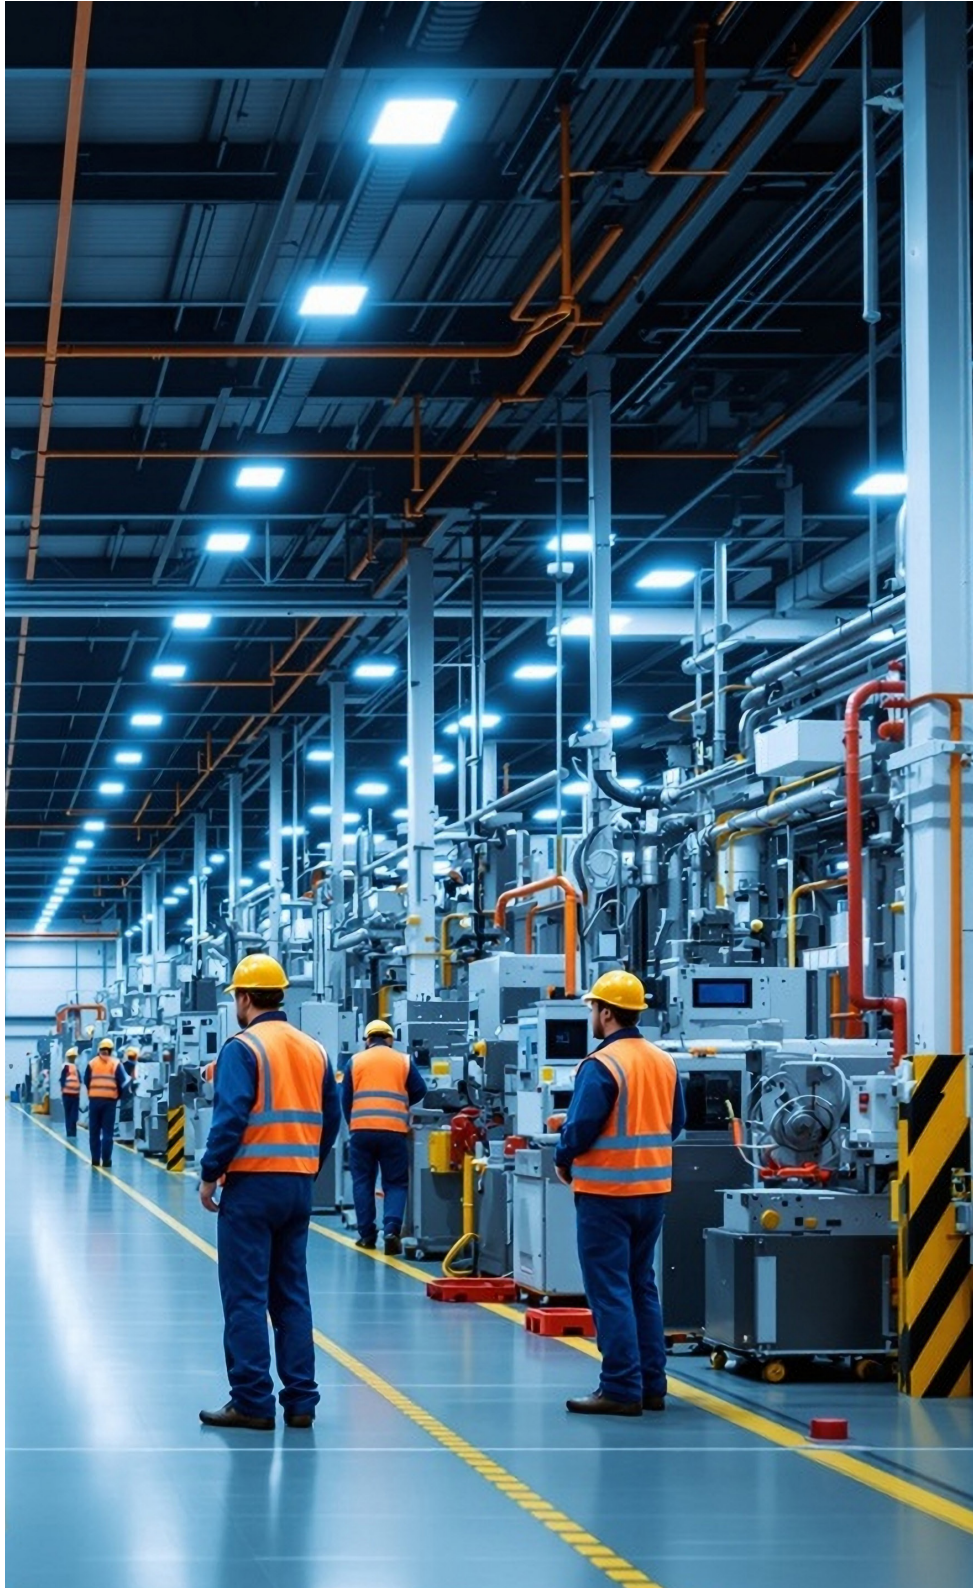

The way forward

Looking forward, energy and utilities enterprises will need to move from isolated sensing toward integrated understanding. The goal is not to replace specialized monitoring systems. Those systems are essential. The goal is to allow their observations to contribute to a shared, persistent, and governed interpretation of the physical operating environment.

Organizations that build this capability will be better able to anticipate complex conditions, prioritize interventions, coordinate responses, and learn from events across the enterprise. They will reduce the gap between what the infrastructure is telling them and what the organization understands. They will also improve explainability because decisions will be grounded in a clearer account of how multiple signals were interpreted together, which is the form of evidence that regulators increasingly expect and which post-hoc reconstruction cannot reliably provide.

The future of operational intelligence in energy and utilities will not be defined by sensing density alone. It will be defined by the enterprise's ability to turn distributed observations into system-level understanding. Utilities that make this shift will operate with greater resilience and confidence, and will meet their obligations to safety, reliability, and regulatory standing from a position of institutional coherence rather than from one of reconstruction. Those that do not will remain constrained by a persistent paradox. They will sense more than ever, yet still struggle to know what it means at the moments when understanding matters most.

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