

BIZTALKS · UTILITIES GAS & POWER | EXECUTIVE PERSPECTIVE · SEPTEMBER 2026

Physical AI Will Create an Accountable Action Layer for Utilities

Physical AI will matter to utilities as a new operating layer that can perceive conditions, prepare decisions and act across dispersed assets.



“ The advantage in Physical AI belongs not to the utility with the largest fleet of robots or remote sensors, but to the organisation that can make machine agency observable, bounded, reversible and accountable across dispersed critical infrastructure.”

A. At a glance

Five signals that frame the executive decision

Executive Signal

Operational Significance and Evidence

70% vs 23% Adoption Gap

In an indicative 2026 IEA survey of 25 network operators, 70% report wide AI use for maintenance, but only 23% for real-time operations. The primary constraint is assurance and accountability, not model availability [3][4].

47 Manual Steps

Great Britain's NESO documented a heavily manual 47-step scheduling process in its Volta innovation project, illustrating how operational complexity accumulates before machines can be granted action authority [6].

Closing Regulatory Window

EU AI Act high-risk rules for critical infrastructure apply from 2 December 2027. Operating model and evidence choices made in 2026 directly dictate what can be legally evidenced and deployed in 2027–2028 [11].

100 km Inspection Chain

A public NTT DATA case study demonstrates multimodal inspection across a 100 km gas corridor in 20 km sections, proving that value comes from connecting capture to engineering follow-up, not imagery volume alone [16].

10,900 Annual Openings

The US Bureau of Labor Statistics projects ~10,900 annual openings for electrical power-line installers over 2025–2035 amid hazardous conditions, reinforcing automation as expertise augmentation rather than headcount elimination [14].



B. What the evidence says

Why operational agency is becoming a board-level question

The operating environment is outrunning manual coordination

Grid expansion is not keeping pace with the transformation of electricity systems. In September 2026, the International Energy Agency reported that at least 1,700 GW of advanced-stage renewable projects and 600 GW of utility-scale batteries were awaiting connection in 2025. It also recorded 2024 congestion costs of USD 12 billion in the United States and EUR 4.3 billion in the European Union. The IEA attributes the delivery constraint to permitting, supply chains, labour availability and equipment costs, while also noting that distributed resources, bidirectional flows and concentrated loads make the network harder to operate [3].

Physical AI does not remove those constraints. It changes the operating response available to a utility while capital projects, regulatory decisions and workforce development proceed on slower clocks. In NTT DATA's public definition, Physical AI gives machines, robots and physical operations the ability to perceive an environment, decide what matters and act in real time. The Utilities application combines computer vision, robotics, IoT and Edge AI with authorised utility data and existing enterprise and operational-technology workflows, while retaining human oversight where risk requires it [1][2].



The US Department of Energy has separately identified near-term AI opportunities across grid planning, permitting, operations and reliability, and resilience, while warning that deployment must not introduce new risks to the grid or individuals. This supports a business-led scope without treating every AI use case as equally mature or consequential [5].

That final condition is decisive. Utilities already use remote sensing, predictive models, drones and automation. Physical AI becomes a distinct executive issue when these components are connected into a loop that can move from observation to intervention. The board is then governing operational agency: which systems may prepare or execute an action, within which limits, using which evidence and under whose accountability.


The scale of the physical estate makes this more than a technology architecture question. One 2025 SEC filing shows a major US dual-fuel investor-owned utility operating approximately 18,000 circuit miles of transmission, 109,000 circuit miles of distribution and 50,900 miles of gas distribution and transmission pipelines. These figures describe one utility, not the global sector, but they illustrate the number of dispersed assets, environments and operating conditions that a single control model may need to address [17].



The adoption gap reveals where value and risk separate

The IEA’s 2026 work places AI applications on a spectrum: forecast, detect, diagnose, screen and prioritise, simulate, and optimise. Inspection and maintenance have scaled first because the data can often be reviewed, errors can be detected before they directly affect the system and a planner, engineer or operator remains in the decision path. Real-time control is moving more slowly because the evidence standard rises when a recommendation can immediately change the physical system [4].

The 70% maintenance versus 23% real-time adoption gap among the surveyed operators should not be read as a forecast for all utilities; the IEA’s sample comprised 11 transmission and 14 distribution operators. It is still a useful strategic signal. The near-term value pool is concentrated where Physical AI can reduce exposure, improve evidence and shorten coordination without placing an opaque model in charge of a safety-critical outcome [3][4].



Source: International Energy Agency (2026). Indicative sample of 25 transmission and distribution network operators [3][4].



This pattern appears in live operator programmes. Great Britain's NESO has used the Volta programme to explore AI-supported scheduling and dispatch. One component examined whether an adviser could recommend and test scheduling scenarios for a process described as a heavily manual 47-step analysis. NESO's published objectives were time savings, reduced security-of-supply risk and greater planning flexibility; they were objectives of a proof of concept, not reported realised benefits [6]. A separate Volta announcement emphasized verification and validation as part of the programme's partner structure [7].

Southwest Power Pool's June 2025 partnership set an objective of reducing generator-interconnection study analysis time by 80%. That is a programme target, not a result, but it shows that system operators are applying industrial AI first to bounded analytical work whose output can be reviewed before it changes the grid [8].

Physical AI must close the loop rather than multiply data

The weak version of Physical AI creates more imagery, alarms and dashboards. The strong version connects field evidence to a decision, routes the decision into an established workflow and records what happened. NTT DATA's Autonomous Operations Centers model describes five elements of that operating loop: unified operational context, intelligence at the execution point, workflow-embedded outputs, critical-infrastructure guardrails and KPI-led progression [1].



The public gas-pipeline case makes the mechanism concrete. A roughly 100 km corridor was divided into mission sections of about 20 km. Long-range UAVs captured visual, thermal and multispectral evidence; the programme connected capture, processing, review and operational follow-up. The implementation ran for nine months with a delivery team of seven to eight full-time equivalents. The case reports improved visibility of vegetation and third-party threats and a foundation for more automated inspection, but does not publish avoided-cost, incident-reduction or cycle-time outcomes [16].



Exhibit: End-to-end operational loop demonstrated across a 9-month implementation. Source: NTT DATA BizTalks [16].

The control point is the transition from evidence to action. An anomaly detected in an image has limited value if it cannot be matched to an asset, assessed against current work, assigned to a responsible team and closed with traceable evidence. In the IEA’s formulation, data quality, network observability and integration with operational systems determine whether AI can move beyond a model demonstration [4].

NTT DATA’s public work provides indicative planning ranges for a more connected operating loop: 20–40% reductions in operational cycle time, 20–40% reductions in manual workload and 15–30% reductions in decision and response time. The source explicitly labels these as objectives dependent on baseline performance, infrastructure, process maturity and deployment scope. They should be used to define a measurement plan, not quoted as achieved outcomes [1].



The workforce case is exposure and capacity, not headcount elimination

Physical AI has a credible role where field work is hazardous, repetitive or difficult to reach. The US Bureau of Labor Statistics describes electrical power-line work as involving high voltage, height, confined spaces, weather and extended emergency shifts. It projects 10% employment growth from 2025 to 2035 and about 10,900 openings a year, many to replace people leaving the occupation [14].

Those data argue against framing automation as a simple labour-substitution programme. The near-term case is to allocate scarce expertise differently: machines perform repeatable observation or enter selected hazardous environments; people validate evidence, manage exceptions and carry accountability for consequential decisions. This design preserves critical field skills rather than burying them behind opaque algorithms [14].



Regulation and airspace rules make governance part of architecture

The 2026–2028 horizon is unusually consequential because technology programmes and regulatory obligations will mature together. The European Commission states that EU AI Act rules for high-risk uses including critical infrastructure will apply from 2 December 2027, with later dates for certain AI embedded in regulated products. The executive implication is that documentation, human oversight, risk management and evidence retention cannot be retrofitted at the end of a pilot [11].

Drone operations add another jurisdiction-specific layer. EASA's June 2026 rules include a predefined risk assessment tailored to routine automated surveillance and infrastructure inspection under specified conditions [12]. In the United States, the FAA's BEYOND programme records approved utility BVLOS inspections, while the FAA targets a general BVLOS framework [13][18].

Cybersecurity moves with the decision, not with the device

Physical AI expands the operational attack surface because sensors, models, communication links, orchestration software and machines participate in one decision chain. NERC's 2025 reliability-risk report lists grid transformation, cyber vulnerabilities, resource adequacy, energy policy and extreme events among the leading risk categories [9]. NIST's April 2026 concept note for a trustworthy-AI profile in critical infrastructure applies a full-lifecycle approach across IT, OT and industrial control systems [10]. For a utility, this means the autonomy boundary must be implemented technically: identity and permissions govern which system may act; validation establishes operating limits; deterministic fallback and safe-state design govern what happens when models fail; and audit logs preserve the complete evidence path.

Why scaling remains difficult

The hard problem is repeatability across sites without assuming that sites are identical. Asset types, failure modes, connectivity, work rules and permissions vary. Models drift as seasons and operating conditions change. The public NTT DATA Physical AI Orchestrator case addresses this problem through a deployment-factory approach: physics-based simulation and digital-twin synchronisation, central mission orchestration, reusable robotic functions, and an operational evidence workflow [15].



C. Implications and scenarios

The Accountable Action Ladder: five levels of operational agency

To evaluate where Physical AI creates value without unacceptable risk, NTT DATA proposes the Accountable Action Ladder. This five-level maturity model evaluates the authority and evidentiary rigor of the operating loop rather than the physical sophistication of the robotic device:

Level	Operating Capability	Human Role	Evidence Required	2026–2028 Posture
Level 1 Observe	Machines capture repeatable field data.	Plan missions and review raw evidence.	Coverage, sensor quality, chain of custody.	Necessary foundation; insufficient alone.
Level 2 Interpret	AI detects, classifies and prioritises anomalies.	Validate findings and alert thresholds.	Accuracy by condition, false positives and misses.	Scale where errors are visible and recoverable.
Level 3 Orchestrate	Findings enter asset, work and control workflows.	Approve priorities and coordinate exceptions.	Asset identity, routing, decision and closure logs.	Recommended enterprise focus for 2026–2028.
Level 4 Execute within bounds	Machines take predefined, reversible actions.	Set authority limits, monitor and intervene.	Verified envelope, safe state, override & rollback tests.	Selective use in low-consequence domains.
Level 5 Prove & adapt	Performance and failures update missions and controls.	Govern change and reauthorise material updates.	Outcome KPIs, drift monitoring, incident learning.	Build now as the underlying assurance layer.

Exhibit: The Accountable Action Ladder. Qualitative maturity model for governing machine agency. Source: NTT DATA [1][2].

Three strategic postures for utility leadership

Posture 1

More pilots

Business units acquire devices and models for local use cases. This can produce fast demonstrations, but evidence, interfaces and safety cases fragment. The organisation remains stuck at Levels 1 and 2, repeating bespoke integration at every asset.

Posture 2

A central autonomy platform

The utility standardises enterprise robotics before operational use cases have established their requirements. This creates architectural consistency on paper, but risks building a platform with weak operational adoption and an authority model too generic for safety-critical field work.

Posture 3

Accountable action domains (Recommended)

The utility selects a focused set of repeatable operating domains, such as corridor inspection, substation rounds or storm-damage assessment. It standardises evidence schemas, interfaces, assurance and escalation paths for each domain, expanding authority only after KPI and safety gates pass.

The recommendation for 2026–2028 is deliberately narrower than broad enterprise autonomy. It concentrates capital on Level 3 (workflow orchestration), while establishing the technical controls and evidence records required for bounded Level 4 actions.



D. NTT DATA point of view

Physical AI must be governed as operational agency, not procured as a fleet of smart devices

The starting point must be the operational decision: what condition must be perceived, what response can be prepared, which action may be executed and who remains accountable. Technology selection follows that chain.

This is hard because utilities must join systems that were designed for different purposes and lifecycles. A drone may capture evidence in minutes, while asset records, work orders and approvals move on different identifiers and clocks.

A model can rank anomalies, but operations still needs a defensible threshold, an escalation route and a safe response when the model or connection fails. The right architecture distributes intelligence according to latency, resilience and risk while maintaining an unbroken, audit-ready decision record [\[1\]](#).



Four operational capabilities for critical infrastructure

Public NTT DATA deployments provide two proven proofs of method. The gas-pipeline case links long-range, multi-sensor capture to processing, review and operational follow-up across 100 km [16]. The Physical AI Orchestrator case links simulation, mission control, reusable robotic functions and field evidence in a deployment-factory methodology [15]. Together they demonstrate that scale comes from an operating and assurance system around the machine.

Core Framework

NTT DATA proposes that utility leadership anchor their Physical AI roadmap across four connected operational capabilities:

- | | |
|--|--|
| <div>1</div> <div>Unified Operational Context</div> <div>Integrating asset data, GIS topology, work management, and OT telemetry into an authoritative, synchronized operational view.</div> | <div>2</div> <div>Edge-to-Workflow Decision Engineering</div> <div>Connecting edge AI perception with automated ticket generation, work-order routing, and human escalation pathways.</div> |
| <div>3</div> <div>Physics Simulation & Fallback Controls</div> <div>Hardening autonomy through digital-twin simulation, deterministic safe-state transitions, and technical boundary enforcement.</div> | <div>4</div> <div>Outcome-Led Governance & Metrics</div> <div>Measuring realized operational deltas (cycle-time reduction, exposure hours avoided) before expanding machine agency.</div> |

A transparent economic framework: directional deltas vs control costs

The business case for Physical AI should be expressed as directional operational deltas tied to a documented baseline: less time between evidence and action, less repetitive review, fewer unnecessary entries into hazardous environments, greater inspection coverage per planning cycle and more consistent closure evidence. Crucially, the investment case must formally budget the new costs of model validation, connectivity resilience, mission supervision, airspace compliance, cyber controls and workforce training. A deployment should advance only when the net operating delta and the safety case remain positive under realistic failure conditions [1].

E. What leaders should do

Five immediate leadership decisions for 2026–2028

-
- 01 Choose two accountable action domains**
Select use cases with material operational value, repeatable work, observable outcomes and recoverable errors (e.g., feeder corridor patrols or substation inspection). Avoid beginning with unrestricted control-room autonomy.

 - 02 Define authority before selecting the machine**
Specify what the system may observe, recommend, prepare and execute; the conditions that suspend authority; and the named human accountable for each decision class.

 - 03 Fund the evidence path as core infrastructure**
Require common asset identities, time-stamped evidence, model and rule versions, approvals, execution records, overrides and closure outcomes across all pilots.

 - 04 Test failure as part of value**
Include degraded connectivity, sensor disagreement, model drift, cyber compromise, airspace interruption and manual fallback in the formal acceptance criteria.

 - 05 Create a 2026–2028 governance calendar**
Map AI, aviation, privacy, safety and sector obligations by jurisdiction, with the EU critical-infrastructure application date of 2 December 2027 as a non-negotiable compliance anchor [\[11\]](#).



Four boardroom questions for executive governance

Q1 Which physical decisions are currently delayed because evidence, expertise and execution sit in disconnected systems?

Q2 Where could a machine act safely within a bounded, reversible envelope, and where must it only recommend?

Q3 Can the utility reconstruct why an AI-enabled action was proposed, approved and executed after an operational incident?

Q4 What tangible operational evidence would justify expanding authority from one site or asset class to the next?

The decision for 2026–2028

Industrialise accountable action domains before expanding machine authority. Keep authority bounded, evidence visible and operational accountability explicit.



Bibliography

Research cut-off: 29 September 2026 · All primary and public sources accessed and verified: 29 September 2026.

Official publications verified across IEA, US DOE, NESO, SPP, NERC, NIST, European Commission, EASA, US FAA, US BLS, SEC filings and NTT DATA BizTalks.

- [1] NTT DATA BizTalks (2026). Autonomous Operations Centers. Current public positioning, operating model and indicative planning ranges. URL: <https://www.biztalksnttdata.com/talk/autonomous-operations-centers/> (Accessed: 29 September 2026).
- [2] NTT DATA BizTalks (2026). Physical AI. Foundry definition, perception-decision-action loop and case index. URL: <https://www.biztalksnttdata.com/foundry/physical-ai/> (Accessed: 29 September 2026).
- [3] International Energy Agency (IEA) (21 September 2026). Modernising Grids in the Age of Electricity, Executive Summary. Grid expansion constraints, connection backlogs and congestion costs. URL: <https://www.iea.org/reports/modernising-grids-in-the-age-of-electricity/executive-summary> (Accessed: 29 September 2026).
- [4] International Energy Agency (IEA) (21 September 2026). AI-enhanced solutions, Modernising Grids in the Age of Electricity. Survey of 25 network operators, application spectrum and 70% vs 23% adoption gap. URL: <https://www.iea.org/reports/modernising-grids-in-the-age-of-electricity/ai-enhanced-solutions> (Accessed: 29 September 2026).
- [5] US Department of Energy (DOE) (29 April 2024). AI for Energy, Opportunities for a Modern Grid and Clean Energy Economy. Grid planning, operations, reliability, and risk-managed AI deployment. URL: <https://www.energy.gov/cet/articles/ai-energy> (Accessed: 29 September 2026).
- [6] National Energy System Operator (NESO) (2025–2026). Volta Advanced Scheduling Adviser. Control-room 47-step manual scheduling process and proof-of-concept objectives. URL: <https://www.neso.energy/about/innovation/our-innovation-projects/volta-advanced-scheduling-adviser> (Accessed: 29 September 2026).
- [7] National Energy System Operator (NESO) (28 August 2025). The Volta programme: Powering the future of dispatch. Innovation programme partnership, verification and validation frameworks. URL: <https://www.neso.energy/about/innovation/our-innovation-projects/volta-advanced-scheduling-adviser> (Accessed: 29 September 2026).
- [8] Southwest Power Pool (SPP) (5 June 2025). SPP Advanced AI Solution for Interconnection Studies. Generator interconnection study time reduction target of 80%. Public announcement. (Accessed: 29 September 2026).
- [9] North American Electric Reliability Corporation (NERC) (22 July 2025). 2025 ERO Reliability Risk Priorities Report. Grid transformation, cyber vulnerabilities and multi-sector interdependencies. URL: https://www.nerc.com/globalassets/our-work/reports/ero-reliability-risk-priorities-report/2025_risc_ero_priorities_report.pdf (Accessed: 29 September 2026).

[10] National Institute of Standards and Technology (NIST) (7 April 2026). Concept Note: Trustworthy AI in Critical Infrastructure Profile. AI Risk Management Framework applied across IT, OT and industrial control systems. URL: https://www.nist.gov/system/files/documents/2026/04/08/Concept%20Note_%20Development%20of%20the%20NIST%20AI%20RMF%20Trustworthy%20Use%20of%20AI%20in%20Critical%20Infrastructure%20Profile.pdf (Accessed: 29 September 2026).

[11] European Commission (2026). AI Act: Regulatory Framework for AI. High-risk critical infrastructure compliance timeline (application date: 2 December 2027). URL: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai> (Accessed: 29 September 2026).

[12] European Union Aviation Safety Agency (EASA) (30 June 2026). Easy Access Rules for Unmanned Aircraft Systems. Predefined risk assessment for automated infrastructure surveillance and drone flights. URL: <https://www.easa.europa.eu/en/document-library/easy-access-rules/online-publications/easy-access-rules-unmanned-aircraft-systems> (Accessed: 29 September 2026).

[13] US Federal Aviation Administration (FAA) (2020–2029). BEYOND: UAS Integration Programme. Beyond Visual Line of Sight (BVLOS) utility infrastructure inspection trials and waivers. URL: https://www.faa.gov/uas/programs_partnerships/beyond (Accessed: 29 September 2026).

[14] US Bureau of Labor Statistics (BLS) (August 2026). Electrical Power-Line Installers and Repairers, Occupational Outlook Handbook. Workforce projections: 10% growth 2025–2035, ~10,900 annual openings. URL: <https://www.bls.gov/ooh/installation-maintenance-and-repair/line-installers-and-repairers.htm> (Accessed: 29 September 2026).

[15] NTT DATA BizTalks (2026). Physical AI Orchestrator for Robotics at Scale. Physics simulation, digital twin sync and deployment-factory methodology. URL: <https://www.biztalksnttdata.com/story/physical-ai-orchestrator-for-robotics-at-scale/> (Accessed: 29 September 2026).

[16] NTT DATA BizTalks (2026). Long-Range UAV Inspection for Gas Pipelines. Multimodal capture, 100 km corridor, five 20 km sections and 9-month delivery evidence. URL: <https://www.biztalksnttdata.com/story/long-range-uav-inspection-for-gas-pipelines/> (Accessed: 29 September 2026).

[17] US Securities and Exchange Commission (SEC) (2026). Major US Dual-Fuel Investor-Owned Utility 2025 Form 10-K. Illustrative utility asset scale: 18k transmission, 109k distribution, 50.9k gas miles. Public regulatory filing. (Accessed: 29 September 2026).

[18] US Federal Aviation Administration (FAA) (September 2026). Testimony of Bryan Bedford, FAA Administrator, FY2027 Budget Request. Status of general BVLOS commercial rulemaking and utility integration. URL: <https://www.faa.gov/testimony/testimony-bryan-bedford-faa-administrator-hearing-committee-appropriations-subcommittee> (Accessed: 29 September 2026).

