

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 will not belong to the utility with the most robots; it will belong to the utility that can make machine action observable, bounded, reversible and accountable.

For review Not approved for publication

Audience Utility CEOs and executive committees

Horizon Immediate decisions · 2026–2028

Research cut-off 29 September 2026

Read time 18 minutes

AT A GLANCE

Five signals that frame the executive decision

- **70% versus 23%** – In a 2026 survey of 25 network operators, 70% reported wide use of AI for maintenance but only 23% for real-time operations. The constraint is increasingly assurance and accountability, not model availability. The IEA cautions that the sample is indicative rather than statistically representative. [\[3\]](#)[\[4\]](#)
- **47 manual steps** – Great Britain's National Energy System Operator described a heavily manual 47-step scheduling process in a completed innovation project that tested an AI scheduling adviser. The example shows how operational complexity can accumulate before any machine is given authority to act. [\[6\]](#)
- **A regulatory window that closes during the planning horizon** – EU AI Act rules for high-risk uses in critical infrastructure are scheduled to apply from 2 December 2027. Governance choices made in 2026 will therefore shape what can be deployed and evidenced in 2027 and 2028. [\[11\]](#)
- **100 km made operationally visible** – An NTT DATA public success story describes a gas-pipeline inspection programme covering about 100 km in sections of about 20 km, connecting visual, thermal and multispectral capture to review and operational follow-up. It is evidence of a connected inspection chain, not a universal ROI benchmark. [\[16\]](#)
- **10,900 openings a year** – The US Bureau of Labor Statistics projects about 10,900 annual openings for electrical power-line installers and repairers over 2025–2035, while emphasizing the

occupation's exposure to high voltage, height, weather and emergency work. This is US evidence, not a global workforce estimate. [\[14\]](#)

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 Pacific Gas and Electric Company 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. [\[3\]\[4\]](#)

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. [\[8\]](#)

The evidence points to a practical sequence. Forecasting, inspection and work preparation can earn trust before automated switching or other

direct interventions. The sequence is governed by consequence and recoverability rather than by how impressive the model appears.

Where AI has scaled in grid operations

Maintenance adoption is three times real-time operations adoption in an indicative 2026 survey of 25 network operators. SHARE REPORTING WIDE USE

- Maintenance **70%**
- Real-time operations **23%**

Source: International Energy Agency, 2026. Sample: 11 transmission and 14 distribution operators; indicative, not statistically representative. [\[3\]](#)[\[4\]](#)

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 talk describes five elements of that operating model: 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. [\[16\]](#)

This gives a quantified end-to-end mechanism without pretending that the case proves universal economics:

100 km physical estate → approximately five 20 km mission sections → multimodal evidence capture → structured review → engineering and maintenance follow-up → nine-month implementation evidence. [\[16\]](#)

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 talk 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. [\[1\]](#)

A transparent benchmark can therefore begin with indexed baselines rather than a promised business case. If the current inspection-to-action cycle is indexed at 100, the stated objective would be an index of 60–80 after deployment. A separate manual-workload index of 100 would target 60–80, while a response-time index of 100 would target 70–85. These ranges must not be added together; they measure different parts of the operating loop. Baseline sampling, safety outcomes and false-positive costs must be established before any benefit is claimed. [\[1\]](#)

From corridor coverage to operational follow-up

The public case connects physical coverage, multimodal evidence and an accountable work path.

2. **~5 sections** about 20 km each
3. **Capture** visual, thermal, multispectral
4. **Review** structured evidence
5. **Follow-up** engineering and maintenance

Reading: a nine-month implementation demonstrated a connected inspection chain; the source does not publish universal ROI. Source: NTT DATA BizTalks. [\[16\]](#)

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 is an

inference from the workforce and operating evidence, and it requires local validation because labour conditions differ by geography and utility. [\[14\]](#) The design should also preserve skills rather than hide them behind a model. Field workers and operators need to understand what the system observed, how confident it is, what action it proposes and how to override or escalate it. Otherwise, automation can reduce the organisation's ability to diagnose failure precisely when fallback is needed.

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 precise classification of a utility use case requires legal analysis; 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, including operating volume, altitude and proximity constraints. [\[12\]](#) In the United States, the FAA's BEYOND programme records approved utility inspections, including BVLOS inspections at power-generation facilities and substations. The FAA was still working toward a more general BVLOS framework in September 2026. [\[13\]](#)[\[18\]](#) These sources do not support a single global deployment template. They support a modular operating model in which airspace approval, privacy, worker consultation, AI classification and asset-criticality controls are configured locally while the evidence model remains consistent.

The governance window sits inside the investment horizon

Documentation, oversight and evidence architecture need to precede the relevant compliance dates.

202620272028

Operating model

EU critical infrastructure [2 Dec 2027](#)

Embedded regulated products [2 Aug 2028](#)

Reading: dates are specific to the EU AI Act implementation page and do not replace use-case legal classification. Source: European Commission.

[\[11\]](#)

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 identified by respondents. It also highlights interdependencies across electricity, gas, water and communications. [\[9\]](#)

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. Its examples include plant-monitoring systems hardened against adversarial input and monitored for conditions outside their validated operating region, as well as diagnostic assistants that provide traceable, auditable rationales. [\[10\]](#)

For a utility, this means the autonomy boundary must be implemented technically. Identity and permissions govern which system may act. Validation establishes the environmental and data conditions in which it may act. Deterministic fallback and safe-state design govern what happens when connectivity, sensors or models fail. Logs must show the evidence, recommendation, approval and execution path. Cybersecurity is therefore part of the authority model, not a perimeter control added to the robot or drone.

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, equipment and operating conditions change. Vendor tools may produce data in formats that do not map cleanly into GIS, EAM, ADMS, SCADA or work-management processes. A successful pilot can still fail as an operating model.

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 navigation, perception and manipulation functions, and an operational evidence workflow. The source reports a scalable methodology and foundation, not quantified business outcomes. [\[15\]](#)

This is consistent with the IEA's observation that proven digital tools frequently stall between site-level deployment and standard practice because each project is justified afresh and methods for transferring evidence are weak. [\[3\]](#) The organisational challenge is to standardise what should repeat, such as assurance, interfaces, evidence schemas and KPI

definitions, while validating what is site-specific, such as hazards, environmental conditions and authorised actions.

IMPLICATIONS AND SCENARIOS *A progression model for 2026–2028*

THE ACCOUNTABLE ACTION LADDER

Authority and evidence, not device sophistication

The proposed analytical artifact is a five-level maturity ladder. It measures the authority and evidence of the operating loop rather than the sophistication of a device. This is a qualitative assessment to be validated with utility operations, safety, cyber and regulatory leaders.

The Accountable Action Ladder *Select a level to inspect how authority, human accountability and evidence change together.*

1Observe **2**Interpret **3**Orchestrate **4**Execute within bounds **5**Prove and adapt

LEVEL 1 Observe

OPERATING CAPABILITY

Machines capture repeatable field data

HUMAN ROLE

Plan missions and review raw evidence

EVIDENCE REQUIRED

Coverage, sensor quality, chain of custody

2026–2028 POSTURE

Necessary foundation, insufficient alone

Nature of the figure: proposed qualitative assessment, to be validated with utility operations, safety, cyber and regulatory leaders.

Three strategic postures follow from the ladder.

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 at Levels 1 and 2 and repeats integration work at each site.

Posture 2 – A central autonomy platform. The utility standardises technology before use cases have established their requirements. This can create architectural consistency, but it risks a platform with weak operational adoption and an authority model too generic for safety-critical work.

Posture 3 – Accountable action domains. The utility selects a small number of repeatable operating domains, such as corridor inspection, substation rounds or storm-damage assessment. It standardises the evidence, interfaces, assurance and escalation path for each domain, then expands across sites after KPI and safety gates are met. This is the recommended direction for 2026–2028.

The recommendation is deliberately narrower than enterprise autonomy. It concentrates investment on Level 3, while preparing the controls needed for bounded Level 4 actions. It also creates evidence that regulators, workers and operators can examine before authority expands.

LEVEL 2 Interpret

OPERATING CAPABILITY

AI detects, classifies and prioritises conditions

HUMAN ROLE

Validate findings and thresholds

EVIDENCE REQUIRED

Accuracy by condition, false positives and misses

2026–2028 POSTURE

Scale where errors are visible and recoverable

Nature of the figure: proposed qualitative assessment, to be validated with utility operations, safety, cyber and regulatory leaders.

Three strategic postures follow from the ladder.

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 at Levels 1 and 2 and repeats integration work at each site.

Posture 2 – A central autonomy platform. The utility standardises technology before use cases have established their requirements. This can create architectural consistency, but it risks a platform with weak operational adoption and an authority model too generic for safety-critical work.

Posture 3 – Accountable action domains. The utility selects a small number of repeatable operating domains, such as corridor inspection, substation rounds or storm-damage assessment. It standardises the evidence, interfaces, assurance and escalation path for each domain, then expands across sites after KPI and safety gates are met. This is the recommended direction for 2026–2028.

The recommendation is deliberately narrower than enterprise autonomy. It concentrates investment on Level 3, while preparing the controls needed for bounded Level 4 actions. It also creates evidence that regulators, workers and operators can examine before authority expands.

LEVEL 3 Orchestrate

OPERATING CAPABILITY

Findings enter asset, work and control workflows

HUMAN ROLE

Approve priorities and coordinate exceptions

EVIDENCE REQUIRED

Asset identity, routing, decision and closure logs
2026–2028 POSTURE

Recommended enterprise focus

Nature of the figure: proposed qualitative assessment, to be validated with utility operations, safety, cyber and regulatory leaders.

Three strategic postures follow from the ladder.

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 at Levels 1 and 2 and repeats integration work at each site.

Posture 2 – A central autonomy platform. The utility standardises technology before use cases have established their requirements. This can create architectural consistency, but it risks a platform with weak operational adoption and an authority model too generic for safety-critical work.

Posture 3 – Accountable action domains. The utility selects a small number of repeatable operating domains, such as corridor inspection, substation rounds or storm-damage assessment. It standardises the evidence, interfaces, assurance and escalation path for each domain, then expands across sites after KPI and safety gates are met. This is the recommended direction for 2026–2028.

The recommendation is deliberately narrower than enterprise autonomy. It concentrates investment on Level 3, while preparing the controls needed for bounded Level 4 actions. It also creates evidence that regulators, workers and operators can examine before authority expands.

LEVEL 4 Execute within bounds **OPERATING CAPABILITY**

Machines take predefined, reversible actions

HUMAN ROLE

Set authority, monitor and intervene

EVIDENCE REQUIRED

Verified envelope, safe state, override and rollback tests

2026–2028 POSTURE

Selective use in low-consequence domains

Nature of the figure: proposed qualitative assessment, to be validated with utility operations, safety, cyber and regulatory leaders.

Three strategic postures follow from the ladder.

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 at Levels 1 and 2 and repeats integration work at each site.

Posture 2 – A central autonomy platform. The utility standardises technology before use cases have established their requirements. This can create architectural consistency, but it risks a platform with weak operational adoption and an authority model too generic for safety-critical work.

Posture 3 – Accountable action domains. The utility selects a small number of repeatable operating domains, such as corridor inspection, substation rounds or storm-damage assessment. It standardises the evidence, interfaces, assurance and escalation path for each domain, then expands across sites after KPI and safety gates are met. This is the recommended direction for 2026–2028.

The recommendation is deliberately narrower than enterprise autonomy. It concentrates investment on Level 3, while preparing the controls needed for bounded Level 4 actions. It also creates evidence that regulators, workers and operators can examine before authority expands.

LEVEL 5 Prove and adapt

OPERATING CAPABILITY

Performance and failures update missions and controls

HUMAN ROLE

Govern change and reauthorise material updates

EVIDENCE REQUIRED

Outcome KPIs, drift monitoring, incident learning

2026–2028 POSTURE

Build now as the assurance layer

Nature of the figure: proposed qualitative assessment, to be validated with utility operations, safety, cyber and regulatory leaders.

Three strategic postures follow from the ladder.

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 at Levels 1 and 2 and repeats integration work at each site.

Posture 2 – A central autonomy platform. The utility standardises technology before use cases have established their requirements. This can create architectural consistency, but it risks a platform with weak operational adoption and an authority model too generic for safety-critical work.

Posture 3 – Accountable action domains. The utility selects a small number of repeatable operating domains, such as corridor inspection, substation rounds or storm-damage assessment. It standardises the evidence, interfaces, assurance and escalation path for each domain, then

expands across sites after KPI and safety gates are met. This is the recommended direction for 2026–2028.

The recommendation is deliberately narrower than enterprise autonomy. It concentrates investment on Level 3, while preparing the controls needed for bounded Level 4 actions. It also creates evidence that regulators, workers and operators can examine before authority expands.

NTT DATA POINT OF VIEW *Proposed positioning, subject to SME validation*

Proposed NTT DATA point of view – subject to SME validation

At NTT DATA, our reading is that Physical AI should be governed as operational agency, not procured as a fleet of smart devices. The starting point would be the operating 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 therefore distributes intelligence according to latency, resilience and risk while keeping one traceable decision record.

The accountable action layer would connect authorised asset and operational context, intelligence at the appropriate execution point, workflow-embedded outputs, risk-calibrated guardrails and KPI-led progression. This adopts the operating logic of NTT DATA's public Autonomous Operations Centers talk without treating its indicative targets as results. [\[1\]](#)

The Accountable Action Ladder provides the progression model. Most utilities can create value now by industrialising Levels 1 to 3 across selected domains. Level 4 should expand only where actions are bounded and reversible, the operating envelope has been validated and human override remains effective. Level 5 should begin at the first pilot because evidence and incident learning cannot be reconstructed later.

Public NTT DATA work provides two relevant proofs of method. The gas-pipeline case links long-range, multi-sensor capture to processing, review and operational follow-up across about 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 support the proposition that scale comes from an operating and assurance system around the machine. They do not establish universal financial returns, and the client identities remain undisclosed in the public sources.

NTT DATA's relevance would therefore sit across four connected responsibilities: integrating asset, field, enterprise and OT context; engineering the edge-to-workflow decision path; establishing simulation, validation, cyber and fallback controls; and measuring operational deltas before expanding authority. These capability claims require confirmation from the Utilities offering owner before publication.

The economics should be expressed as directional deltas tied to a 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. The investment case should also include 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.

WHAT LEADERS SHOULD DO *Five immediate decisions and four boardroom questions*

1. **Choose two accountable action domains.** Select use cases with material operational value, repeatable work, observable outcomes and recoverable errors. Avoid beginning with unrestricted control-room autonomy.
2. **Define the 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.
3. **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 pilots.
4. **Test failure as part of value.** Include degraded connectivity, sensor disagreement, model drift, cyber compromise, airspace interruption and manual fallback in the acceptance plan.
5. **Create a 2026–2028 governance calendar.** Map AI, aviation, privacy, safety and sector obligations by jurisdiction, with the EU

critical-infrastructure date of 2 December 2027 as one visible milestone. [\[11\]](#)

Boardroom questions

- Which physical decisions are currently delayed because evidence, expertise and execution sit in different systems?
- Where could a machine act safely within a bounded, reversible envelope, and where must it only recommend?
- Can the utility reconstruct why an AI-enabled action was proposed, approved and executed after an incident?
- What 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.

Review the evidence

BIBLIOGRAPHY

Primary and public sources, all accessed 29 September 2026

1. [\[1\]](#)

Autonomous Operations Centers – NTT DATA BizTalks – 2026, page undated – Current public positioning and indicative objectives – [Open source](#) – – Accessed 29 September 2026.

2. [\[2\]](#)

Physical AI – NTT DATA BizTalks – 2026, page undated – Current foundry definition and public case index – [Open source](#) – – Accessed 29 September 2026.

3. [\[3\]](#)

Modernising Grids in the Age of Electricity, Executive Summary – International Energy Agency – 21 September 2026 – 2024–2026 data and operator survey – [Open source](#) – – Accessed 29 September 2026.

4. **[4]**

AI-enhanced solutions, Modernising Grids in the Age of Electricity – International Energy Agency – 21 September 2026 – 2026 survey of 25 network operators and current deployments – [Open source](#) – – Accessed 29 September 2026.

5. **[5]**

AI for Energy, Opportunities for a Modern Grid and Clean Energy Economy – US Department of Energy – 29 April 2024 – Near-term grid opportunities – [Open source](#) – – Accessed 29 September 2026.

6. **[6]**

Volta Advanced Scheduling Adviser – National Energy System Operator – Project May 2025 to February 2026, page current at access – Control-room process and proof-of-concept objectives – [Open source](#) – – Accessed 29 September 2026.

7. **[7]**

The Volta programme, Powering the future of dispatch with Tapestry – National Energy System Operator – 28 August 2025 – Current innovation programme – [Open source](#) – – Accessed 29 September 2026.

8. **[8]**

SPP Partners with Hitachi to Develop Advanced AI Solution – Southwest Power Pool – 5 June 2025 – Partnership objectives – [Open source](#) – – Accessed 29 September 2026.

9. **[9]**

2025 ERO Reliability Risk Priorities Report – North American Electric Reliability Corporation Reliability Issues Steering Committee – Approved 22 July 2025 – 2024–2025 stakeholder inputs – [Open source](#) – – Accessed 29 September 2026.

10. **[10]**

Concept Note, Artificial Intelligence Risk Management Framework, Trustworthy AI in Critical Infrastructure Profile – National Institute of Standards and Technology – 7 April 2026 – Profile development scope – [Open source](#) – – Accessed 29 September 2026.

11. **[11]**

AI Act, Regulatory Framework for AI – European Commission – Current implementation page, including July 2026 amendments – 2024–2028 application timeline – [Open source](#) – – Accessed 29 September 2026.

12. **[12]**

Easy Access Rules for Unmanned Aircraft Systems, Revision from June 2026 – European Union Aviation Safety Agency – 30 June 2026 – UAS operations and infrastructure inspection – [Open source](#) – – Accessed 29 September 2026.

13. **[13]**

BEYOND – US Federal Aviation Administration – Current programme page, BEYOND phases 2020–2029 – BVLOS trials and utility waivers – [Open source](#) – – Accessed 29 September 2026.

14. **[14]**

Electrical Power-Line Installers and Repairers, Occupational Outlook Handbook – US Bureau of Labor Statistics – Updated August 2026 – 2025 employment and 2025–2035 projections – [Open source](#) – – Accessed 29 September 2026.

15. **[15]**

Physical AI Orchestrator for Robotics at Scale – NTT DATA BizTalks – 2026, page undated – Public success story – [Open source](#) – – Accessed 29 September 2026.

16. **[16]**

Long-Range UAV Inspection for Gas Pipelines – NTT DATA BizTalks – 2026, page undated – Public success story and implementation parameters – [Open source](#) – – Accessed 29 September 2026.

17. **[17]**

PG&E Corporation and Pacific Gas and Electric Company 2025 Annual Report, Form 10-K – PG&E Corporation and Pacific Gas and Electric Company, filed with the US Securities and Exchange Commission – 2026 filing for year ended 31 December 2025 – 2025 asset and operating data – [Open source](#) – – Accessed 29 September 2026.

18. **[18]**

Testimony of Bryan Bedford, FAA Administrator, Fiscal Year 2027 Budget Request – US Federal Aviation Administration – September

EVIDENCE AND REVIEW ANNEX

Retained for Gate 2 and human validation

Gate 1 was approved by the user on 29 September 2026. The status remains 80% Grounded Draft because Knowledge Owner, Utilities SME, offering, legal, editorial and brand validations have not been confirmed.

Gate 2 quality verdict

PASSED – no open high-severity finding. The approved A–E text, all 18 numbered sources, all 18 claim-ledger entries, the visible status and the public-safe case framing are consistent in the source package and this HTML export.

Block	Result	Coverage and residual note
Factual	Passed	Claims, figures, dates, qualifiers and inference labels checked against the bibliography. Three BizTalks item pages were confirmed through the live Utilities/Physical AI catalogues and indexed page content after direct refetch was blocked by the research tool.
Format and brand	Passed with partial browser coverage	Self-contained NTT DATA HTML; embedded logo; four source-labelled visuals; responsive table labels; no gradient, shadow, emoji or external asset dependency in the rendered page. Static DOM checks and a 15-page fallback print render passed. The runtime could not download a Playwright browser, so final target-browser review remains human.
Catalogue and lexicon	Passed	Autonomous Operations Centers and both named success stories exist under the public Utilities/Physical AI catalogues; capability claims remain labelled as proposed or subject to owner validation.
Public hygiene	Passed	No pricing, named client employee, restricted analyst content, confidential metric, placeholder or unapproved client logo.
Cross-format consistency	Passed	83 approved content units checked; zero missing. Citation sequence, title, figures, status and bibliography numbering match the approved package.

Consolidated human review list

1. **Knowledge Owner and Utilities SME:**validate “operational agency,” “accountable action layer,” the five ladder levels and the recommended Level 3 focus.
2. **Utilities offering owner:**confirm the exact wording and delivery feasibility of the integration, assurance and edge-to-workflow capability statement.
3. **Legal or regulatory reviewer:**refresh EU AI Act and UAS/BVLOS wording and dates immediately before publication, then classify the intended use cases by jurisdiction.
4. **Utilities SME and editorial reviewer:**retain the qualifiers around indicative target ranges and the US-specific workforce evidence.

5. **Editorial and brand reviewers:**approve publication voice and inspect interaction, responsive layout, logo isolation and typography in the target publishing browser.

Open high-severity items: none. **Partial coverage:** browser execution and direct refetch of three BizTalks item URLs were constrained by the runtime; catalogue/index verification and static rendering were used. **Resolved in this pass:** citation linking, mobile table labels, self-contained assets, Gate 1 status wording and cross-format text preservation.

Claim Ledger

#	Material claim	Source	Verification status	Confidence	Publication note
C1	Physical AI enables machines and physical operations to perceive, decide and act in real time	[2]	Verified in public NTT DATA source	High	Approved public wording; SME should confirm use in Utilities context
C2	The Autonomous Operations Centers model combines Physical AI with authorised data and workflows under human oversight	[1]	Verified	High	Public-safe
C3	At least 1,700 GW of renewable projects and 600 GW of batteries awaited connection in 2025	[3]	Verified in IEA 2026 report	High	Preserve geography-neutral global wording
C4	IEA surveyed operators report 70% wide AI use for maintenance and 23% for real-time operations	[3][4]	Verified with sample limitation	Medium-high	Always state n=25 and indicative limitation when space permits
C5	NESO described a 47-step manual scheduling analysis	[6]	Verified	High	Project description, not sector benchmark
C6	SPP set an objective of reducing interconnection analysis time by 80%	[8]	Verified as target	High	Do not describe as achieved
C7	EU high-risk critical-infrastructure AI rules apply from 2 December 2027	[11]	Verified against EC page at cut-off	High	Legal review required before publication; regulatory dates can change
C8	NTT DATA pipeline case covered about 100 km in about 20 km mission sections over nine months	[16]	Verified in public case	High	No client name; do not infer ROI
C9	BLS projects about 10,900 annual openings for US power-line installers and repairers	[14]	Verified	High	US-specific, 2025–2035

#	Material claim	Source	Verification status	Confidence	Publication note
C10	A single US utility reported approximately 18,000 transmission circuit miles, 109,000 distribution circuit miles and 50,900 gas pipeline miles	[17]	Verified in SEC filing	High	Illustrative company scale only
C11	Physical AI's near-term value is greatest where errors are visible and recoverable	[4] plus analysis	Analytical inference	Medium-high	Label as interpretation
C12	Physical AI should be governed as operational agency	Editorial thesis	Proposed NTT DATA point of view	To validate	Knowledge Owner and SME approval required
C13	The Accountable Action Ladder is a useful maturity model	Original artifact	Proposed qualitative assessment	To validate	Validate with operations, safety, cyber and regulatory SMEs
C14	Indicative 20–40% cycle-time and workload reductions and 15–30% response-time reduction can frame measurement	[1]	Verified as indicative objectives	Medium	Never present as achieved impact
C15	NIST is developing a trustworthy-AI profile for critical infrastructure	[10]	Verified	High	Profile was in development at cut-off
C16	US and European drone frameworks allow infrastructure inspection under jurisdiction-specific conditions	[12][13][18]	Verified at cut-off	High	Avoid implying harmonisation or blanket approval
C17	NTT DATA's deployment-factory case connects simulation, orchestration, reusable functions and evidence	[15]	Verified in public case	High	No quantified result claimed
C18	NTT DATA can integrate and govern the accountable action layer across IT, OT, edge and field workflows	[1][15][16]	Partly grounded capability claim	Medium	Offering owner must validate feasibility and exact wording

Depth Ledger

Element	What the piece does	Where	Pass
Reframing thesis	Treats Physical AI as operational agency and an accountable action layer rather than a device category	Standfirst and section D	Yes, pending SME validation

Element	What the piece does	Where	Pass
Debatable and quotable	Makes accountability, not robot count, the source of advantage	Standfirst	Yes
Named concept	The accountable action layer	Throughout	Yes, original for this piece
Original artifact	Five-level Accountable Action Ladder	Section C	Yes, qualitative and labelled
Implications form	Maturity ladder plus three strategic postures	Section C	Yes
Why now	IEA September 2026 report, EU 2027 compliance date, active system-operator programmes	Sections A and B	Yes
Quantified mechanism	100 km corridor, approximately five 20 km sections, evidence-to-follow-up workflow and nine-month implementation	Section B	Yes
Why hard	Data, workflow, safety, cyber, airspace, validation, workforce and cross-site variation	Section B	Yes
Fact and opinion separation	Evidence precedes the visibly labelled proposed NTT DATA view	Sections B and D	Yes
Real cases	Two public NTT DATA cases used with no invented client identity or ROI	Section D	Yes
Economics	Directional deltas and new control costs; targets kept separate from outcomes	Section D	Yes
Diagnosis weight	Approximately 61% of article body precedes the NTT DATA POV	Sections B and C	Yes
Evidence density	18 sources; mix of NTT DATA public assets, IEA, DOE, NERC, NIST, EC, EASA, FAA, BLS, system operators and SEC filing	Bibliography	Yes
Counterargument	Automation may add data, fragility and oversight cost if the loop and authority are not designed	Sections B and D	Yes

Human review requirements

Item	Risk type	Why it matters	Source or reference	Recommended action	Reviewer	Status
Thesis and named concept	Institutional position	“Operational agency” and “accountable action layer” are proposed NTT DATA language	Standfirst and section D	Approve, edit or reject the framing	Knowledge Owner and Utilities SME	Gate 1 user-approved; role review pending

Item	Risk type	Why it matters	Source or reference	Recommended action	Reviewer	Status
Accountable Action Ladder	Domain validity	Authority levels and boundaries affect safety and operating practice	Section C	Validate levels, terminology and recommended Level 3 focus	Operations, safety and OT SMEs	Gate 1 user-approved; role review pending
NTT DATA capability statement	Offering accuracy	Integration, assurance and edge-to-workflow claims must reflect deliverable capability	Section D; [1][15][16]	Confirm exact wording and delivery evidence	Utilities offering owner	Gate 1 user-approved; role review pending
Public case use	Knowledge governance	Cases are public but client identities are undisclosed and outcomes are limited	[15][16]	Confirm both cases may be featured in a BizTalks perspective	Knowledge Owner	Gate 1 user-approved; role review pending
Indicative target ranges	Evidence interpretation	Ranges are objectives, not realised outcomes	[1]	Keep the qualifier and indexed illustration, or remove the ranges	Utilities SME and editorial reviewer	Gate 1 user-approved; role review pending
Regulatory wording	Legal and currency	AI and aviation rules vary and may change after the cut-off	[11][12][13][18]	Legal review and date refresh immediately before publication	Legal or regulatory reviewer	Gate 1 user-approved; role review pending
Workforce inference	Geographic transferability	BLS data are US-specific; labour conditions differ globally	[14]	Retain as a bounded example or add regional evidence for priority markets	Utilities SME and editorial reviewer	Gate 1 user-approved; role review pending
Publication voice and length	Editorial	Extended version is the source master and may require web compression	Full draft	Approve the extended master before any HTML export or derivative	Editorial reviewer and user	Gate 1 user-approved; editorial role review pending
Visual assets and logo	Brand and rights	Publication HTML requires	Content design blueprint	Use design-system embedded logo;	Brand reviewer	Gate 1 user-approved;

Item	Risk type	Why it matters	Source or reference	Recommended action	Reviewer	Status
		approved logo and any imagery must be rights-cleared		source imagery only after rights check		brand role review pending

Source Retrieval Log						
Library	Query or path	File or source	Location	Date or version	Classification	Used for
public-biztalks	User-supplied URL	Autonomous Operations Centers	Public web	Accessed 29 Sep 2026	Public and reusable	Framing, operating model, indicative objectives
public-biztalks	Physical AI foundry	Physical AI	Public web	Accessed 29 Sep 2026	Public and reusable	Canonical definition and case sweep
public-biztalks	Case sweep	Physical AI Orchestrator for Robotics at Scale	Public web	Accessed 29 Sep 2026	Public and reusable	Scaling method
public-biztalks	Case sweep	Long-Range UAV Inspection for Gas Pipelines	Public web	Accessed 29 Sep 2026	Public and reusable	Quantified mechanism and case evidence
biztalks-content	01 INDUSTRIES / 03 Utilities, then site search for Physical AI and autonomous operations	No matching documents returned	SharePoint Rebirth	29 Sep 2026	Access unavailable for evidentiary use	Internal case sweep
value-offering	Shared Documents / GAS AND POWER, then site search for Physical AI and robotics	No matching documents returned	SharePoint UtilitiesAndEnergyBusinessLine	29 Sep 2026	Access unavailable for evidentiary use	Capability validation
brand	Contents, then site search for brand guidelines and logo	No matching documents returned	SharePoint TheMarketingCollective	29 Sep 2026	Access unavailable for evidentiary use	Brand validation

Library	Query or path	File or source	Location	Date or version	Classification	Used for
external primary sources	Grid AI, operations, reliability, regulation, aviation, workforce and filings	IEA, DOE, NERC, NIST, EC, EASA, FAA, BLS, NESO, SPP, SEC filing	Public web	Cut-off 29 Sep 2026	Public and reusable subject to source terms	Diagnosis and verification

Publication architecture

The recommended publication format after Gate 1 is an interactive NTT DATA HTML perspective using the NTT DATA design system: Future Blue, Smart Navy, white and one functional green accent; Noto Sans for reading and Noto Serif for the hero; logo at the top right; no gradients, shadows or decorative imagery.

1. **Hero and thesis.** Title, standfirst, visible 80% Grounded Draft flag and research cut-off.
2. **Evidence contrast.** A two-bar comparison of 70% wide AI use in maintenance and 23% in real-time operations, explicitly labelled as an indicative IEA survey of 25 network operators. [\[3\]](#)
3. **Operating mechanism.** A left-to-right chain showing the quantified pipeline case: 100 km, approximately five 20 km sections, multimodal capture, review and follow-up, with the nine-month programme duration in the caption. [\[16\]](#)
4. **Interactive Accountable Action Ladder.** Five selectable levels reveal the capability, human role, evidence test and deployment boundary. The figure is a proposed qualitative assessment to be validated.
5. **Governance timeline.** 2026–2028 milestones covering current enforcement, the 2 December 2027 EU critical-infrastructure date and the 2 August 2028 product-related high-risk date, with a jurisdiction warning. [\[11\]](#)
6. **Decision module.** Four boardroom questions with downloadable review prompts after publication approval.
7. **Source drawer.** Full bibliography, claim ledger and methodology; no internal source content.

Visual and Source Asset Register

Asset	Retrieved from	Rights/classification	Used in	Status
NTT DATA global logo, inline blue/white variants	Installed NTT DATA design-system package	Bundled corporate asset	Fixed masthead and footer	Embedded; no external dependency

Asset	Retrieved from	Rights/classification	Used in	Status
Evidence visuals	Approved editorial package and cited public sources [3][4][11][16]	Original CSS/data rendering; no third-party imagery	Adoption bars, inspection flow, governance timeline	Source-labelled
Accountable Action Ladder	Approved editorial package	Original qualitative framework, to validate	Interactive maturity module	Source nature labelled

80% Grounded Draft · Research cut-off 29 September 2026 · NTT DATA
Utilities