

The autonomy gradient:

an adoption architecture for physical AI

Turning physical intelligence into business value: an executive guide for enterprise adoption

September, 2026



Executive summary

Enterprise AI has moved from experimentation to a broader phase of adoption. However, an execution gap remains, and the numbers tell the story. Nearly nine in 10 organizations now use AI in at least one business function, yet only about one-third have begun scaling it across the enterprise. Roughly two-thirds remain in experimentation or pilot.¹

Model capability alone is no longer sufficient. Production physical AI depends on how the wider system comes together. Nowhere is this more evident than in physical AI, where intelligence must operate on factory floors, in warehouses, across energy networks and within critical infrastructure. The cost of a wrong or delayed decision is measured in scrap, downtime and safety.

The paper makes its case in six parts:

- **The physical AI inflection point:** What has changed technically and commercially.
- **Why good pilots fail:** The seven friction points that separate a successful demonstration from dependable production.
- **The physical AI adoption architecture:** Six interconnected layers, including the connectivity layer most frameworks omit, that make physical AI reliable at scale.
- **The Autonomy Gradient:** A six-level framework, A0 through A5, for deciding how much autonomy each workflow should carry, with a scoreable rubric leaders can apply in minutes.
- **Physical AI in action:** Four NTT DATA engagements across manufacturing, energy and infrastructure, from foundation readiness to conditional autonomy.
- **A leader's decision framework:** An executive checklist, a staged rollout plan with gate criteria and a materiality test for deciding where physical AI pays.

The central idea is deliberately contrarian. Autonomy is not the destination. It is a design choice. The strongest physical AI programs do not aim for maximum autonomy. They determine how much autonomy each workflow needs based on its context, risk and the role of human expertise. The Autonomy Gradient provides a common language for making these choices, much as SAE J3016 created a framework for driving automation in the automotive industry.²

1. The physical AI inflection point

1.1 From digital intelligence to physical intelligence

The first wave of enterprise AI generated content, analyzed information and improved decisions made at a desk. The next wave extends intelligence into the physical world, where work is performed and consequential decisions are made every second. This transition is giving rise to physical AI.

Physical AI is AI embedded into the perception, reasoning and actuation loop of a physical system. It interprets unstructured sensor inputs such as video, audio, vibration and telemetry. It reasons about the physical environment and acts, whether by guiding an operator or commanding a machine.

Physical AI is not rebranded industrial IoT, SCADA or condition monitoring. These systems typically collect and display data based on predefined rules. Physical AI, by contrast, interprets situations it was not explicitly programmed for, adapts as conditions change and participates in the decision itself.

Physical AI is a significant step beyond traditional automation. Conventional machine vision performs predefined tasks in controlled environments.

Foundation models change the economics. They understand changing environments, adapt to new conditions and improve with additional data, so a new use case can be introduced in weeks rather than months.

Consider a forklift assembly line. Traditional automation verifies individual assembly checkpoints. Physical AI observes the complete workstation, understands every stage of the process and provides real-time guidance whenever a deviation is detected. As production requirements evolve or new product variants are introduced, the system adapts without a rebuild.




1.2 How physical AI works

Physical AI turns information into action through a continuous cycle. It captures signals from the physical world and processes them close to the source. AI then interprets what is happening and routes the resulting insight to the person or system that must act. The system learns from the outcome and feeds that learning back into the cycle.

The cycle only works when every link, from sensor and network to edge compute, model and workflow, is engineered as one system. That is the premise of the Adoption Architecture discussed in Section 3.

Physical AI is often equated with robotics, but robots are only one element. Its real strength lies in connecting AI, systems and workflows, so timely decisions happen with or without a robot in the loop.



1.3 Why 2026 marks a pivotal year

Four measurable shifts converged going into 2026.

01

Capital moved decisively

Gartner forecasts worldwide AI spending of \$2.59 trillion in 2026, up 47% year-over-year. AI infrastructure is expected to account for more than 45% of that spending.³

02

The model class arrived

Commercially usable vision-language-action and world foundation models are making physical AI more practical. Examples include NVIDIA Cosmos, Isaac GR00T, physical Intelligence's pi-zero family and Google DeepMind's Gemini Robotics. However, long-horizon reliability remains a challenge, making autonomy a critical design decision shaped by context, risk and human expertise. Stanford's Fei-Fei Li, whose ImageNet work seeded modern computer vision, calls spatial intelligence "the next frontier of artificial intelligence" and argues that language models alone remain ungrounded in the physical world.⁴

03

The market formalized

Analyst firms now recognize it as a distinct category with a full-stack view spanning devices, compute, simulation, models, security, connectivity and services. NVIDIA founder and CEO Jensen Huang framed the moment at CES 2026: "the ChatGPT moment for physical AI is here," the point at which machines begin to understand, reason and act in the real world.^{5, 6}

04

Edge infrastructure matured

Sensor and camera costs have declined steadily while edge inference hardware has improved in performance per watt with each generation. This moves real-time inference from a special project toward a defensible line item.

The opportunity is significant, and so is the complexity. Many companies can demonstrate physical AI in controlled settings. Far fewer can integrate it with existing technology, core systems and frontline processes at scale. The difference between demonstration and dependable production is where this paper begins.



2. Why good pilots fail

2.1 The adoption gap

The technology is ready and the business case is strengthening, yet only a minority of initiatives progress to large-scale implementation. For every 33 AI proofs of concept an enterprise starts, only around four reach production. In 2025, 42% of US companies abandoned most of their AI initiatives, up from 17% a year earlier.^{7,8}

Today's models can interpret images, analyze sensor data and support complex decisions. The challenge begins when those capabilities must operate inside real production environments. Infrastructure, processes and compliance requirements vary across facilities. What performs well at one site often needs significant adaptation elsewhere. Closing that gap is an architectural problem, not a modeling problem.



2.2 The seven friction points




Figure 1. The seven friction points between pilot and production.

Friction 1. Edge infrastructure readiness

Physical AI enables timely decisions, which requires inference to run close to where data is generated. Most production environments were never designed for distributed computing, and 'close to the source' means different things for different workloads. The latency envelope, not preference, dictates where inference must physically sit, from on-machine controllers for millisecond safety responses to plant-edge servers for in-line inspection.

Friction 2. Connectivity and network readiness

The link between sensing and edge compute is a network. In most brownfield industrial sites, it is also the weakest link. Metal-dense plants, moving assets, outdoor yards and RF-hostile process areas can challenge consumer-grade wireless. Wired connections can be costly or impractical for mobile use cases. Safety-critical traffic also cannot share best-effort infrastructure with everything else. Deterministic behavior, coverage, mobility and traffic separation must be engineered deliberately.

Friction 3. IT and OT integration

Physical AI needs to connect information from both IT and OT environments. These systems have often evolved separately, making it difficult to share data, context and insights across them. Successful adoption requires a clear path for information to move across machines, operational platforms and enterprise systems.

Friction 4. Data residency and latency

Many decisions cannot wait for a round trip to a centralized cloud, and regulation increasingly requires that they do not make the trip at all. GDPR and the EU Data Act govern European operational data, and the EU AI Act imposes obligations on high-risk AI used in safety functions and critical infrastructure. China's PIPL and data export rules constrain manufacturing footprints there. Architectures must support local processing without losing central oversight.

Friction 5. Safety and quality certification

Introducing AI into production requires a level of assurance that a demonstration does not. Physical AI must meet the safety, quality and regulatory requirements of the environment in which it operates.

Relevant requirements include IEC 61508, ISO 13849, ISO 10218, ISO/TS 15066 and ISO 3691-4, along with the EU AI Act and NIST AI Risk Management Framework.

Friction 6. OT cybersecurity

Physical AI introduces new security challenges into environments that were often designed for isolation. New sensors, edge systems and data flows create new points to secure across the IT/OT boundary. Security reviews can become a significant barrier to moving industrial AI projects from pilot to production.

Friction 7. Workforce and operating model


Technology alone does not transform anything. Teams need to understand how AI fits into their existing responsibilities, when they should rely on its recommendations and when human judgment should take precedence.

When AI complements rather than obscures human expertise, adoption becomes easier to sustain.

2.3 Scaling requires an architectural foundation

These seven barriers are closely connected. Solving one while overlooking another rarely produces lasting results.

A pilot demonstrates technical capability. Scale demands repeatable, resilient execution across every location. This shifts the conversation from model performance to architectural strategy. NTT DATA addresses this through a structured approach that unifies infrastructure, connectivity, AI, governance and business functions in a single framework: the [Physical AI Adoption Architecture](#).



3. Physical AI Adoption Architecture

3.1 Six layers, one operating environment

Scaling physical AI requires every component to work together within one connected operating environment. The NTT DATA Physical AI Adoption Architecture consists of six interconnected layers; each serves a distinct purpose, and together they enable physical AI to operate reliably at scale.




Figure 2. The Physical AI Adoption Architecture.

Layer 1. Sensing

Every informed decision begins with reliable data. Cameras, sensors and connected industrial devices capture the information that powers physical AI. Sensor fusion brings these inputs together, creating a more complete view of the physical environment.

Successful adoption depends on thoughtful placement rather than increasing the number of devices. Every sensor should capture information that contributes directly to a defined objective.

Layer 2. Connectivity

Connectivity links sensing, computing and action. It is especially important in brownfield environments, where existing networks may not have been designed for real-time AI workloads.

The right connectivity approach depends on the environment and workload. [Private 5G](#) can provide secure, fast and reliable mobile connectivity for applications such as autonomous mobile robots, outdoor yards and large facilities. Wi-Fi can support dense, fixed deployments, while Time-Sensitive Networking can support workloads that require highly predictable wired performance.

Layer 3. Edge compute

Many decisions must be made within seconds. Processing information close to its source improves responsiveness and strengthens resilience. Edge computing allows AI to operate directly inside production facilities without relying entirely on centralized cloud infrastructure.

Layer 4. Model

A mature model layer is organized as a curated model zoo, with open-source and commercial models documented through model cards covering intended use, data provenance, performance and limitations. Edge deployments combine task-specific models for low-latency detection with core models for contextual reasoning. Models also need to represent spatial relationships, physical context and the state of the environment to support reliable decisions in dynamic settings. Production lifecycle management includes pretrained, fine-tuned and site-adapted models, simulation, synthetic data, quantization, distillation, fleet versioning and drift detection.

Layer 5. Workflow integration

Many physical AI initiatives lose momentum at exactly this stage. Valuable insights remain isolated unless they become part of everyday execution. Maintenance teams, quality engineers, production supervisors and core applications all need timely access to relevant insight. Integration with MES, ERP, SCADA and existing workflows ensures AI strengthens established processes instead of creating parallel ones. For workflows that require action, intelligence must also translate into planning, control and execution through the appropriate operational systems, machines or human interfaces. This is where measurable outcomes begin to emerge.



Layer 6. Governance

Trust is the precondition for adoption, and governance is how trust is manufactured. In production, governance is a set of concrete mechanisms, not principles: a model registry with an approval workflow; ground truth sampling with acceptance thresholds; human override logging with audit-trail retention; a rehearsed incident and rollback procedure; model change control tied to the site's quality management system, mapped to the EU AI Act and NIST AI RMF where they apply.

Governance is also the layer that operationalizes the Autonomy Gradient. Each autonomy level in Section 4 carries a defined governance requirement, so moving a workflow up the gradient is a controlled, evidenced decision rather than a gradual drift.



3.2 Security and reliability across the layers

Security by design means protecting devices, models and data flows across the IT and OT environment. IEC 62443 provides an established framework for structuring these protections.

Reliability is equally important. Physical AI becomes part of production, so it must meet production-level availability expectations. This requires redundancy where a single failure could interrupt production. It also requires clear fallback modes so systems continue safely when a component fails. For example, an inspection workflow can queue frames during a network interruption rather than silently passing products.

Availability targets should be agreed for each workflow and measured like any other operational KPI.

3.3 From implementation to maturity

Maturity develops progressively as architecture, governance and working practices evolve. The table below lets a leadership team locate its organization in under a minute.

	Stage 1. Instrumented	Stage 2. Assisted	Stage 3. Orchestrated	Stage 4. Adaptive
Sensing and data	Key assets instrumented; data siloed by system	Sensing designed per decision; data contextualized	Unified namespace across lines	Cross-site data products feed models continuously
Connectivity	Best-effort plant network	Engineered coverage for pilot areas	Deterministic and sliced by traffic class	Multi-site standard design, centrally operated
Edge and model operations	Single pilot node; manual model updates	Managed nodes; versioned models	Fleet operations; drift monitoring live	Automated retraining triggers; simulation in the loop
Workflow integration	Standalone dashboard	Alerts into email or chat	Events land in MES, QMS and maintenance systems	Closed-loop actions with confirmation
Governance	Ad hoc validation	Approval workflow for go-live	Registry, override logging, rollback tested	Autonomy-level framework governs every workflow
Typical autonomy	A1 Monitored	A2 Advisory	A2 to A3 by workflow	A3 to A5 where justified

Table 1. The Physical AI Maturity Model across the six architecture layers.




Figure 3. Maturity progression.

4. The Autonomy Gradient

4.1 How far is far enough?

As physical AI becomes part of day-to-day work, one question continues to shape every adoption journey:

Where should autonomy stop?

Technology now makes high autonomy possible almost everywhere, which pulls attention toward maximizing it. Large-scale adoption follows a different path. Every setting has its own priorities, constraints and acceptable level of risk.

The most effective physical AI initiatives are driven by outcomes rather than technology alone. Some workflows justify full autonomy because human presence is impractical or unsafe. Others create more value with people in the loop. The economics point the same way. Stanford economist Erik Brynjolfsson calls the fixation on fully human-replacing automation “the Turing Trap,” arguing that AI which augments people rather than imitating them ultimately generates far more value.⁹ This progression can be understood through the Autonomy Gradient.

The Autonomy Gradient is a six-level framework, A0 through A5, for determining the right level of autonomy for each workflow based on four measurable determinants: human presence, decision latency, judgment content and operational consequences. Autonomy is treated as a design choice per workflow, not a destination for the enterprise, in the way SAE J3016 treats driving automation as defined levels rather than a binary.²




Figure 4. The Autonomy Gradient.

4.2 The six levels of the Autonomy Gradient

Level name	Who acts, who approves	Latency envelope	Governance artifact required	Industrial example
A0 Manual	Humans perform and decide; AI absent	Human speed	Standard operating procedures	Craft assembly, exception handling
A1 Monitored	AI observes and reports; human decides and acts	Human speed	Data and detection quality validation	Asset condition dashboards, energy monitoring
A2 Advisory	AI recommends; human decides and acts	Seconds to minutes	Recommendation accuracy tracking; override logging	Assembly guidance, inspection assist, maintenance prioritization
A3 Supervised execution	AI acts; human approves before effect or reviews a reversible action immediately after	Sub-second to seconds	Approval workflow; full audit trail; rollback path	Auto-divert suspect parts pending review, drone flight plans approved per mission
A4 Conditional autonomy	AI acts within a defined operating envelope; human handles exceptions	Milliseconds to seconds	Envelope definition; exception escalation; periodic revalidation	AMR fleets in mapped zones, Beyond Visual Line of Sight (BVLOS) inspection on approved corridors
A5 Full autonomy	AI acts without human involvement; audit only	Milliseconds	Certification against the applicable safety regime; continuous assurance monitoring	Safety interlocks, subsea and hazardous environment operations

Table 2. The six levels of the Autonomy Gradient.

Higher autonomy does not necessarily mean higher value. As the examples in Section 5 show, some of the strongest returns come from A2 and A3 use cases rather than full autonomy at A5. A single plant may use A2 for inspection assistance and A4 for an AMR fleet, with each workflow governed according to its specific requirements.

4.3 A scoreable rubric: the four determinants




Figure 5. The four determinants of appropriate autonomy.

Rate each determinant from 1 to 5 for the workflow in question; higher scores favor higher autonomy.

Determinant	Score 1	Score 5
Human presence	Humans naturally present and attentive at the point of work	Human presence impractical or unsafe
Decision latency	Minutes or longer; deliberation is affordable	Sub-second; only automation can act in time
Judgment content	Depends on experience and context	Fully specifiable rules and thresholds
Operational consequence	Severe or irreversible	Minor and easily reversible

Table 3. Scoring the four determinants. A total of 4 to 8 indicates A1 to A2; 9 to 13, A2 to A3; 14 to 17, A3 to A4; and 18 to 20, A4 to A5, subject to the governance artifacts in Table 2.

4.4 What the gradient looks like across industries



Figure 6. Physical AI across industries.

Manufacturing is primarily at A2 and A3 today, with use cases such as inspection assistance, assembly guidance and automated part diversion. A4 is emerging in intralogistics through autonomous mobile robot (AMR) fleets.

Energy and utilities span a wider range, from A1–A2 for condition monitoring to A3–A4 for drone-based inspection. For drones, aviation regulations and approvals can determine how much autonomy is possible.

Logistics and warehousing have some of the most mature A4 applications. AMR and AGV fleets can operate autonomously within defined areas, with people handling exceptions.

Critical infrastructure generally stays at A1–A2 for high-consequence decisions such as switching and control. At the same time, inspection can reach A4–A5 when operating in environments where sending people presents a greater risk.

4.5 Progression along the gradient

Progress along the gradient is rarely linear. Early initiatives improve visibility and strengthen decision support. As experience grows, higher levels of autonomy become appropriate. Robotist Rodney Brooks' annual reviews of autonomy predictions document how consistently announced autonomy outruns delivered autonomy¹⁰; levels are earned through governance and operating evidence, not announcements.

The **Autonomy Gradient** provides a practical framework for making these decisions.

Physical AI in action

The following NTT DATA deployments show how physical AI moves from architecture to implementation, and how different workflows call for different levels of autonomy.

5.1 Turning quality control into real-time intelligence

Industry: Manufacturing

Autonomy level: A2 Advisory

A leading manufacturer faced a familiar hurdle where product quality depended heavily on operator experience and consistent adherence to standard operating procedures. The objective was to improve consistency without disrupting established manufacturing processes.

NTT DATA developed an AI-powered Workstation Twin: a computer-vision digital representation of the physical workstation that continuously monitors assembly activity. Live video lets the system identify deviations, detect quality issues as they occur and guide operators immediately. Quality assurance moved from post-production inspection into the assembly process itself.

The solution improved visibility into workstation performance and recurring deviations. Earlier identification of deviations helped reduce rework and warranty exposure. Real-time guidance supported more consistent adherence to operating procedures, contributing to stronger product quality. Overall, quality assurance shifted from reactive inspection toward real-time guidance, while integration with existing workflows created a foundation for broader physical AI adoption across production lines.



5.2 Reimagining power infrastructure inspection for Grid Sky Way

Industry: Energy

Autonomy level: A4 (Conditional autonomy)

Electric power infrastructure extends across vast geographic regions, much of it difficult to access. Severe weather, natural disasters and workforce shortages make the manual inspection that network reliability depends on increasingly hard to sustain. Grid Sky Way was developed to address these constraints through a nationwide drone flight network dedicated to electric power infrastructure in Japan.

NTT DATA enabled intelligent drone operations by combining autonomous flight with centralized route and airspace management. The platform consolidates flight information, weather, drone ports and operational controls, giving inspection teams one coordinated view of the entire network.

Under Japan's drone regulations, the deployment enables Level 4 autonomous flight, allowing drones to operate beyond visual line of sight over inhabited areas. This creates a pathway toward higher autonomy through regulatory approval. Drones enable faster infrastructure assessment when landslides prevent field teams from reaching affected assets. Aerial assessment helps utilities prioritize restoration while keeping field personnel out of hazardous areas. More broadly, the deployment demonstrates how higher autonomy is earned through the right combination of technology, operational controls and regulatory approval.

5.3 From aerial coverage to actionable intelligence

Industry: Infrastructure

Autonomy level: A3 Supervised execution

River inspection becomes especially critical after floods and other natural disasters. Conventional patrols can take one to two hours to cover long distances. Some affected areas may be unsafe to reach.

NTT DATA addressed this challenge by coordinating multiple autonomous aircraft through a single inspection system. A high-speed VTOL drone flies rapid corridor patrols while a rotary-wing drone performs detailed inspection where needed. Both operate autonomously through the airpalette UTM platform, which coordinates flight paths, manages multiple drones and streams inspection data to operators over LTE connectivity.

The approach reduced patrol time for a 30-kilometer river section to approximately 30 minutes. Teams could move from broad situational awareness to detailed investigation without deploying additional field crews. On the Autonomy Gradient, the deployment represents A3 supervised execution, with operators planning and approving missions while the aircraft fly autonomously and are reviewed in real time.



5.4 Making infrastructure ready for physical AI

Industry: Manufacturing

Autonomy level: Foundation (pre-deployment): architecture readiness for the gradient

Battery manufacturing combines high-speed production with equipment from multiple OEMs, creating a complex environment for data, monitoring and AI. The manufacturer needed an architecture that could connect production systems and provide a reliable foundation for physical AI.

NTT DATA designed a scalable architecture that unifies production visibility and creates AI-ready data pipelines across the manufacturing environment. The architecture also incorporates resilience and data integrity through redundancy and store-and-forward capabilities while addressing cybersecurity and regulatory requirements. By supporting intelligence closer to production, the architecture enables real-time decision-making without relying entirely on centralized cloud processing.



6.The economics of physical AI

6.1 Four drivers of measurable outcomes

Four mechanisms convert the architecture into financial performance.



Figure 7. Four drivers of measurable outcomes.

Real-time decisions.

Production environments generate thousands of events daily. Any delay between observation and response costs schedule, quality and safety. Edge intelligence collapses that delay, so minor issues are corrected before they become major ones.

Data sovereignty.

Local inference satisfies residency and security requirements, from GDPR to sector rules, while preserving end-to-end visibility.

Rapid iteration.

Foundation models evolve with changing production requirements, so new products and process changes are incorporated in weeks, keeping innovation in step with the business.

Scalable expansion.

The greatest benefits arrive after the first successful rollout. A proven architecture is a repeatable blueprint. Each subsequent site deploys faster, cheaper and more predictably than the last.



6.2 The materiality test

An honest return-on-investment conversation starts with materiality. Physical AI does not clear the investment bar everywhere. On low-volume lines, in processes that are already stable and well controlled, and at sub-scale sites, the full cost of deployment spans sensing, connectivity, integration and ongoing model operations. The cost can exceed the benefits accrued. A capability that is impressive in a demonstration can still be immaterial to the P&L.

The discipline cuts both ways. The same four determinants that set the right autonomy level also expose workflows where the right level of investment is zero, at least for now. Deciding where not to deploy is part of the decision framework, not a failure of ambition. It is what keeps a portfolio's returns compounding rather than diluting.

6.3 Rethinking return on investment

Physical AI is often evaluated through workforce reduction, which captures only a small part of its potential. The broader view considers how better decisions improve every stage of production: earlier defect detection, greater process visibility, faster issue resolution and higher first-pass quality. These improvements compound. For executive leaders, the conversation shifts from the cost of implementing AI to the long-term advantage of an intelligent operating model.

The greatest return from physical AI comes from better decisions, faster implementation and repeatable expansion. Operational excellence becomes the output of an architecture designed for continuous improvement.



7.A leader's decision framework

7.1 Turning strategy into action

The next step is a roadmap that aligns technology investment with operational priorities. Every company starts from a different maturity level. NTT DATA's structured approach begins by assessing the six layers of the Adoption Architecture. It then prioritizes workloads by appropriate autonomy level, so each initiative matches business needs, acceptable risk and long-term objectives.

7.2 Sequence workflows, then scale

Physical AI does not need to arrive everywhere at once. The strongest results come from staged expansion with explicit gates.



Table 4. A staged rollout with gate criteria.

7.3 Executive checklist

- Before beginning, or before expanding, a physical AI program, answer five questions:
- Have you identified the workflows where physical AI can deliver measurable value and scored them against the four determinants?
 - Is your architecture ready across all six layers: sensing, connectivity, edge, model, integration and governance?
 - Which workflows belong at A2 Advisory, and which can justify A3 and above, with the governance artifacts each level requires?
 - Can successful deployments be replicated across sites and business units at declining cost per site?
 - Do you have a partner able to integrate IT, OT, connectivity and AI into one scalable operating model?

Why NTT DATA

Adopting physical AI requires expertise across IT, AI, edge computing, connectivity, governance and core systems. It also requires a partner that understands both digital platforms and physical environments.

Many firms can provide individual capabilities. NTT DATA brings them together in one connected offering. This spans global private 5G and connectivity built on the NTT Group's network heritage, data center and edge infrastructure operated at global scale, deep IT and OT integration across MES, ERP and industrial platforms, and Japanese industrial expertise demonstrated through platforms such as Grid Sky Way and airpalette UTM. Together, these capabilities align with the six layers of the Physical AI Adoption Architecture presented in this paper.

NTT DATA supports progress through three strategic pillars:

Modernize the foundation

Prepare infrastructure, connectivity, data and core systems to support physical AI at scale.

Embed AI into operations

Integrate intelligence into workflows where it can improve quality, safety and productivity. Apply the level of autonomy that each workflow requires.

Scale with trust and control

Expand physical AI responsibly across facilities. Maintain transparency, accountability and governance as adoption grows.

Closing thoughts

Intelligence is no longer confined to dashboards, reports and centralized systems. It is becoming part of the environments where products are built, assets are maintained and critical decisions shape outcomes every day. Success will depend on strong architecture and clear-eyed judgment about where autonomy creates impact and where it creates risk.

That judgment is what this paper tries to equip. The six-layer adoption architecture makes the engineering tractable. The Autonomy Gradient makes the autonomy decision explicit, scoreable and governable. The staged rollout turns both into a plan with dates and gates.

The future of AI will be realized where intelligence meets action. With the right architecture, the right level of autonomy for each workflow and a trusted partner, intelligent operations will become a lasting competitive advantage. Those who invest today will lead the next generation of intelligent industries.

References

- ¹ [McKinsey & Company, "The State of AI in 2025: Agents, innovation, and transformation," McKinsey Global Survey, November 2025.](#)
- ² [SAE International, "J3016: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles."](#)
- ³ [Gartner, "Gartner Forecasts Worldwide AI Spending to Grow 47% in 2026," press release, May 19, 2026.](#)
- ⁴ [F.-F. Li, "From Words to Worlds: Spatial Intelligence Is AI's Next Frontier," essay, November 2025.](#)
- ⁵ [IDC, "IDC Market Glance: Physical AI — Embodied Intelligence, 2Q26," document AP54610726, June 2026.](#)
- ⁶ [NVIDIA, CES 2026 keynote and press announcements, J. Huang remarks on physical AI, January 2026.](#)
- ⁷ [S&P Global Market Intelligence, enterprise AI initiative survey findings on initiative abandonment, 2025.](#)
- ⁸ [IDC and Lenovo, "The AI CIO Playbook," 2025 \(proof-of-concept to production conversion analysis\).](#)
- ⁹ [E. Brynjolfsson, "The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence," Daedalus 151\(2\), Spring 2022.](#)
- ¹⁰ [R. Brooks, "Predictions Scorecard," annual reviews of robotics and autonomous-vehicle predictions, rodneymrooks.com, 2018–2025.](#)

Disclaimer

Case study outcomes are described qualitatively and are indicative rather than audited figures. One engagement is anonymized at the client's request. Third-party research findings are cited from the publicly reported sources listed in the References and remain the property of their respective publishers.

Visit nttdata.com to learn more.

NTT DATA is a \$30+ billion business and technology services leader in AI and digital infrastructure. We accelerate client success and positively impact society through responsible innovation. As a Global Top Employer, we have experts in more than 70 countries. NTT DATA is part of NTT Group.

