Physical AI at scale. Manufacturing & Warehousing

From pilots to governed autonomy in physical operations

Contents

1 **Physical AI at scale: what it actually takes** Page 6

2 **Physical AI at scale: reference guide for practitioners** Page 31

3 **Physical AI glossary: a working vocabulary for operational leaders** Page 40

Authors

Prasad Satyavolu

Senior Managing Director, Accenture Global Lead - Manufacturing, Physical AI & Robotics

Prasad Satyavolu shapes Accenture's global manufacturing agenda, with a key focus on growing the firm's Physical AI business across the Americas, EMEA and APAC. He brings more than 30 years of executive experience in industry and professional services at the intersection of cyber physical systems and the convergence of digital and physical automation. He has led large-scale transformations across manufacturing, product development, supply chain and operations. His work is focused on advancing re-industrialization, from major CapEx programs and competitive products/platforms to Physical AI-enabled manufacturing.

LinkedIn

Christian Souche

**Managing Director,
Accenture Global Lead - Robotics**

Christian leads Robotics & Physical AI activities at Accenture , working at the intersection of human, business and technology trends to imagine and launch the next generation of robots in industrial and non-industrial domains. Christian has extensive experience managing R&D to incubate emerging services and products to accelerate service activation at scale.

LinkedIn

> **"Physical AI is the point where AI stops advising the enterprise and starts helping the enterprise act safely, physically, and continuously across the full operational value chain. It is about designing operations for flexibility, agility, and resilience (DFAR) while raising the bar on reliability, safety, and maintainability."**
>
> Prasad Satyavolu

About this series

Physical AI is reshaping how organizations design, operate and continuously improve physical environments. While this first volume focuses on manufacturing and warehousing, Physical AI extends far beyond industrial operations into cities, homes, infrastructure, healthcare, retail, mobility and other domains of the physical economy.

The Physical AI at Scale series explores these domains through a common lens: building trusted representations of reality, enabling intelligent action in the physical world and orchestrating people, machines, and systems at scale.

This first book examines how Physical AI is transforming manufacturing and warehousing operations, from digital twins and simulation to robotics, orchestration, and autonomous operations.

Physical AI at scale: what it actually takes

Where does Physical AI start and where does it stop? Does it begin when a company designs a new product, when it designs the factory that will make it, when it defines the robotics and automation systems that will run it, or when that product is operating in a customer environment? The answer is that Physical AI does not belong to one point in the lifecycle. Its full value emerges only when it is considered across the company's full physical value chain: from product design and industrialization to production, operation, service, maintenance, and customer experience.

Physical AI is the convergence of artificial intelligence, simulation, digital representations of reality, autonomous systems, and orchestration technologies that enable physical operations to perceive, understand, decide, act, and continuously learn.

It combines world models, digital twins, operational data, sensors, robotics, intelligent machines, enterprise systems, and human expertise into a closed-loop system that connects the digital and physical worlds. Physical AI helps organizations translate business intent into safe, governed action across the physical value chain, from design and industrialization to production, logistics, operations, service, and infrastructure management.

At its core, Physical AI increases the autonomy, adaptability, and resilience of physical systems. It enables organizations to understand what is happening, simulate what could happen, decide what should happen next, coordinate execution across people and machines, and continuously improve through feedback and learning.

63%
agree on the vision of an AI-driven, highly automated operation.

Without that transversal view, organizations capture only fragments of the Physical AI opportunity. They may improve a design workflow, automate a cell, or optimize a facility, but they will not unlock the compounding value that comes from connecting intent, evidence, execution, and learning across the full lifecycle. Physical AI at scale is defined by whether operational decisions, how to plan, design, operate, and optimize in the physical world, are made within clear constraints, supported by evidence, and governed with safe override and rollback paths. Without that, organizations are not scaling Physical AI but automation in pockets.

It starts with more than a technical problem. Major capital programs[1] run 30–40% over time and budget.[2] New facilities routinely miss their targets; one plant designed for 150 units per day initially achieved 20.[3]
Most organizations in manufacturing, consumer, and infrastructure industries have already invested in AI, digital twins, robotics, and automation. The next step is harder: making those technologies work together as part of a governed physical operating system. That requires simulation-ready digital twins that capture behavior and constraints, AI systems that can reason and act within clear boundaries, and robots and automation that can be coordinated across heterogeneous environments. The technology exists and is increasingly accessible, with the annual growth rate of Physical AI estimated at close to 50%. What remains difficult is control: turning business intent into safe, repeatable action across mixed operations, and learning fast enough to stay ahead of disruption. A recent Accenture survey of 552 factory managers confirms the gap. While 63% agree on the vision of an AI-driven, highly automated operation, only 38% are actively building toward it today. The constraint is the absence of a system that connects the pieces.

This document is about what it actually takes to close that gap: the system logic, the structural requirements, and the leadership decisions that determine whether Physical AI scales. It focuses on the physical execution layer, how to design, govern, and scale Physical AI across manufacturing, consumer, and infrastructure operations. For the broader enterprise AI strategy and the case for systemic AI across the full manufacturing lifecycle, see Accenture's companion report: Systemic AI at the root of manufacturing performance (2026).

What Physical AI is and is not

Physical AI is a system capability, not a collection of tools.

It is the technology that enables physical systems to increase their level of autonomy: translating business intent (targets, constraints, safety rules, and operating policies) into safe, repeatable action in the physical world, while continuously learning from outcomes under clear decision rights, traceability, and human oversight where required.

Recent advances in AI foundation models are accelerating the development of Physical AI. Vision-language models (VLMs) enable systems to interpret complex physical environments and reason about what they observe using natural language. Vision-language-action models (VLAs) extend these capabilities by connecting perception and reasoning directly to physical action, enabling robots, autonomous systems, and intelligent machines to understand instructions, adapt to changing conditions, and execute tasks in the real world. Together, these models reduce the need for extensive task-specific programming and help physical systems learn more generalized behaviors across manufacturing and warehousing environments.

At its core, Physical AI works as a loop: decisions are tested virtually before they reach the floor; autonomous behaviors are trained and validated within clear constraints; and execution across people, machines, and robots is coordinated at the pace of operations.

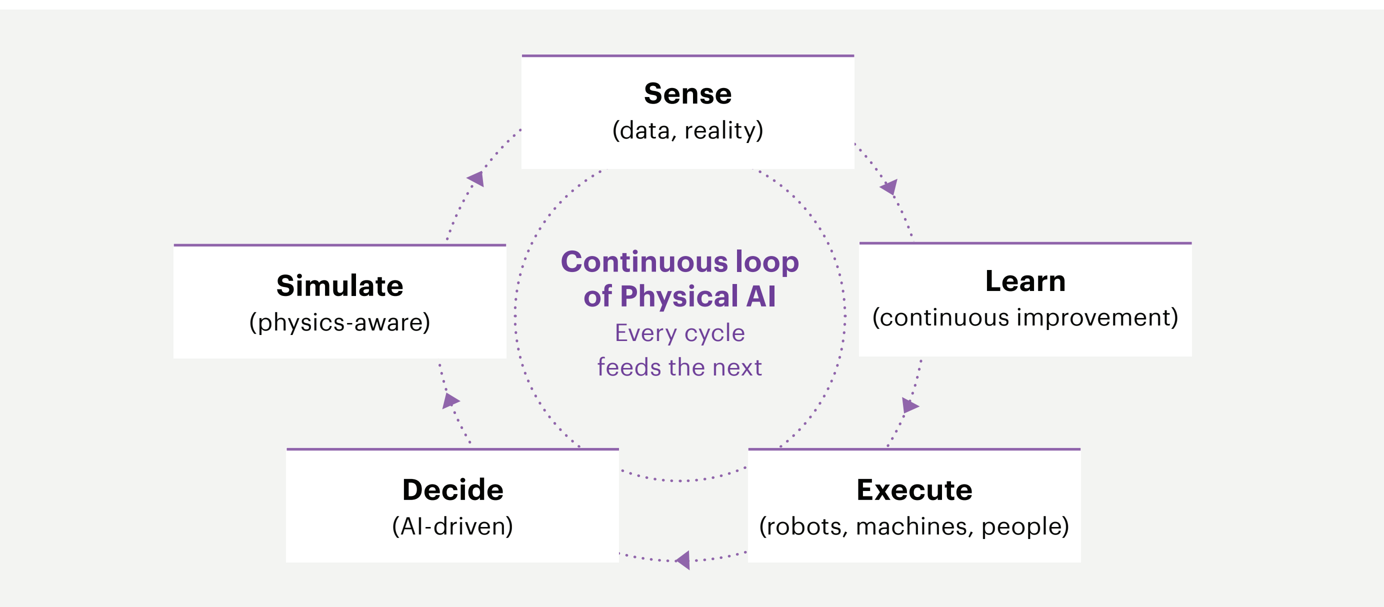

This is what distinguishes Physical AI from traditional automation, analytics, or digital twin initiatives. Traditional automation relies on coding autonomous system behavior in advance. Physical AI shifts that model toward training system behavior: using simulation, data, and real-world feedback to help systems learn how to sense, decide, act, and improve under changing physical conditions. The unit of value is therefore not a robot, a model, or a dashboard. It is the end-to-end closed loop that allows an operation to increase autonomy safely and continuously.

> **Physical AI is more than a robotics deployment, a digital twin program, or an AI initiative.**
>
> **Each can play a role, but none is sufficient on its own.**

Where you are is where you start

Organizations approaching Physical AI come from different positions. Legacy infrastructure, fragmented systems, aging assets, high dependency on human expertise and growing labor shortages are the norm. The starting point does not determine whether you can move; it shapes where you enter and what you prioritize first.

Most organizations enter from one of three positions, often with some overlap:

01 **Greenfield.**
Organizations building new have the greatest design freedom, often under increasing pressure to relocate or regionalize production closer to demand. Proximity to market can reduce transport-related carbon emissions, lower tariff exposure, and improve resilience against supply disruption. The critical decisions are therefore architectural and strategic: which systems to integrate from day one, how to establish data continuity across product design, factory design, automation, operations, and service, and how to avoid building the fragmentation that older facilities are now trying to escape.

02 **Modernizing.**
Organizations modernizing existing operations are making selective investments against a live environment. The challenge is sequencing: where to intervene without destabilizing what works, and how to build reusable foundations rather than accumulating more point solutions.

03 **Foundational.**
Organizations earlier in the journey (with fragmented toolchains, limited sensor coverage, or inconsistent data quality) are building the preconditions for physical AI first. The work here is establishing data integrity, clarifying decision rights, and identifying the one or two operational domains where a controlled deployment can demonstrate what the system is capable of.

Physical AI maturity diagnostic

These starting positions apply across sectors. A greenfield data center and a greenfield automotive plant face the same architectural decisions. A modernizing hospital and a modernizing warehouse face the same sequencing challenge. What they are ultimately building toward is also similar: a highly autonomous and flexible physical operation, whether factory, warehouse, facility, or infrastructure system, that can adjust layouts, flows, schedules, and tasks in very short time frames, with limited human intervention and clear human control where it matters. The destination is a closed-loop system where learning happens before disruption, execution is governed rather than improvised, and agility becomes an operating capability rather than a one-off response.

Physical AI is not a single technology implementation. It is the convergence of intelligent infrastructure, simulation, automation, orchestration, governance, workforce transformation, and robot intelligence.

Use the assessment questions below to map your current state across each dimension. The output provides a fact-based view of your maturity today and identifies the capabilities required to unlock the next level of value, autonomy, resilience, and operational performance.

For each dimension, select the stage that best describes your organization today, not your future aspirations. Be evidence-based and choose the stage that is consistently demonstrated across the operation.

Assess your Physical AI maturity

Physical AI maturity model part 1: Foundation

Stage 1 — Connected	**Stage 2** — Predictive	**Stage 3** — Autonomous within bounds	**Stage 4** — Self-optimizing
Intent and constraints: How clearly is business intent translated into rules the operation can act on?			
Targets and KPIs live in plans and performance reviews. The operation runs on experience and judgment, not codified rules.	KPIs are visible across the operation and feed AI-generated recommendations. Humans interpret those recommendations and decide what to do.	Decision boundaries are explicit and embedded: what the system can decide autonomously, what it must escalate, and under what conditions.	Operating intent is dynamic. The system adjusts its own objectives within defined limits as conditions change, without human reconfiguration each time.
Infrastructure: What can your physical and technical environment actually perceive, process, and execute?			
Assets are connected and generating data. Automation is deployed but programmed, not adaptive. The infrastructure supports visibility, not autonomous execution.	Sensor coverage supports predictive use cases. Scalable AI computing is operational. The infrastructure supports continuous analysis but not real-time autonomous execution.	Autonomous assets operate reliably within defined environments. Real-time inferencing runs at the machine and robot level. Validated for autonomous execution, not just monitoring.	Compute, sensing, and robotics adapt dynamically to operational conditions. Robot fleets operate under unified coordination. The infrastructure reconfigures itself as demands change.

Physical AI maturity model part 2: Technology

Stage 1 — Connected	**Stage 2** — Predictive	**Stage 3** — Autonomous within bounds	**Stage 4** — Self-optimizing
Design and simulate: How clearly is business intent translated into rules the operation can act on?			
Operational data is collected and reported. Teams review what happened. Changes are not simulated before they are committed.	Digital twins are operational for key assets or processes. Simulation is used to test changes before deployment. Predictive analytics generates forward-looking insight.	Outcomes from live operations feed back into models continuously. Scenario replay validates readiness before any change reaches the floor. Synthetic data covers situations too rare or unsafe to replicate live.	Simulation and live operations are synchronized in real time. The operation continuously generates its own training data. What one site learns transfers to others without manual re-engineering.
Autonomy and coordination: How are decisions made and executed across your operation?			
Automation executes programmed tasks. All decisions are made by people. No coordination layer connects assets, systems, and workflows into a single operating picture.	AI recommendations are integrated into planning and scheduling. A control tower provides cross-functional visibility. Coordination is human-led but AI-informed.	Autonomous agents execute bounded workflows without human approval at each step. A neutral orchestration layer coordinates people, machines, robots, and enterprise systems in real time.	Autonomous planning, scheduling, and execution are coordinated as one system across sites and functions. Local decisions are continuously aligned with enterprise constraints without human mediation.
Robot and agent intelligence: How capable and coordinated is intelligence at the robot and autonomous agent level?			
Robots and automation equipment execute fixed, programmed routines. Behaviour is deterministic: the robot does exactly what it was told to do, nothing more. No on-device reasoning. No coordination between machines. Human intervention is required for any deviation from the programmed path.	Robots operate with basic perception and limited adaptability. AI models inform robot behaviour like collision avoidance, path optimisation, anomaly detection but reasoning happens off-device and upstream. Robots are smarter, but still instruction-followers. Coordination between machines is rule-based, not dynamic.	On-device intelligence enables robots to reason and adapt locally without waiting for upstream instruction. Robots handle variation, unexpected conditions, and edge cases within defined envelopes. Coordination between robots, AMRs, cobots, and enterprise systems happens in real time through a neutral orchestration layer. This is the RoboBrain tier, where robot-level AI becomes a meaningful capability differentiator.	Robot intelligence is fleet-wide and self-improving. Individual robots contribute observations to shared models. The fleet learns collectively, what one robot encounters updates how others behave. Robots negotiate tasks, priorities, and handoffs autonomously. Human oversight is concentrated at fleet governance level, not individual machine level.

Physical AI maturity model part 3: Organization

Stage 1 — Connected	**Stage 2** — Predictive	**Stage 3** — Autonomous within bounds	**Stage 4** — Self-optimizing
Governance and change control: How does your organization control what the system can do and how it changes			
Governance is defined on paper. Ownership and decision rights exist but are not embedded in how the operation runs day to day.	Governance is operational. AI model management is active. Changes follow a defined process. Accountability for AI-generated decisions is assigned.	Safety envelopes are validated and tested. Human override mechanisms have been exercised under realistic conditions and are trusted by the workforce. Nothing reaches production without an evidence gate.	The system monitors its own compliance, flags drift, and initiates controlled rollback without human intervention. Auditability is continuous. Governance scales automatically as the program expands.
Security and resilience: How does your operation protect itself and recover when something goes wrong?			
OT and IT are connected with a cybersecurity baseline in place. Security is perimeter-focused. Recovery from disruption is manual and depends on who is available.	Security covers the expanded attack surface of connected assets, AI models, and digital twins. Monitoring is active. Recovery procedures are documented and tested.	The system degrades safely under failure or attack: critical operations continue, faults are isolated, autonomous behaviors suspend cleanly. Resilience is validated, not assumed.	Threat detection, response, and recovery are partially automated. Resilience is tested continuously through simulation. The system's behavior under failure is part of the governance
Talent and workforce: How has your workforce's role changed as the operation has become more intelligent?			
The workforce knows change is coming. AI literacy programs are underway. Roles have not materially changed. Training is classroom-based and disconnected from daily work.	Engineers, operators, and planners use AI tools in their daily workflow. Skills are developing but unevenly distributed. Workers are adapting to AI-assisted decision-making with varying confidence. Workers begin using connected worker tools (mobile devices, digital work instructions, real-time AI feeds) to access insights at the point of work. Adoption is uneven and tool-dependent.	Roles have structurally changed. Operators supervise autonomous systems. The workforce helped design how their roles evolved. Trust in the system is demonstrated through use, not assumed. Connected worker is embedded in operations. Operators receive real-time alerts, contextual AI guidance, and system status through wearables or devices. The human worker is integrated into the autonomous loop, interacting with it continuously.	Human expertise is concentrated on exception handling, continuous improvement, and expanding the system's operating envelope. Insights from human overrides feed systematically back into models. The connected worker is a data source as well as a decision point. Human actions, overrides, and observations feed back into models automatically. The boundary between human and autonomous system is fluid and governed.

Why now?

Four shifts have converged to make Physical AI viable at scale; and they are not reversible.

01

Hardware has become AI-native.

A new generation of robots, autonomous systems, sensors, and edge devices has moved beyond the limited perception and cognitive capacity of earlier automation. Advanced sensor coverage, onboard compute, and embedded intelligence now allow physical systems to perceive more context, process more information locally, and act on Physical AI models closer to the point of execution. This changes the role of hardware: it is no longer only an actuator at the end of a programmed process, but a trainable, connected part of the closed loop.

02

AI has moved from insight to agency.

AI systems are no longer limited to recommendations. With appropriate governance, they can make and execute decisions within defined boundaries. NIO operates body shops with 300 robots and 12 workers producing 20 cars per hour.[5] BMW has reported a 400% efficiency improvement from deploying humanoid robots in production.[6]

03

Operational environments are irreversibly heterogeneous..

Mixed systems, legacy infrastructure, multiple technology vendors, and human variability are the norm across manufacturing, retail, healthcare, hospitality, and data-intensive environments alike. Scaling in this reality requires a neutral coordination layer rather than tighter coupling to any single platform or vendor.

04 Simulation has become the training and commissioning environment.

Simulation is not the starting point of Physical AI, but it is a critical component once organizations move from programmed automation to trained behavior. High-fidelity environments are used for virtual commissioning, scenario replay, readiness validation, and the generation of synthetic data for robot and autonomy training. This makes simulation an evidence and learning environment: a place to test how systems behave before they are deployed physically, and to expand training coverage for situations that are rare, unsafe, or expensive to capture in live operations. In one documented example, proposal review cycles that previously took six to eight months were reduced to weeks.[4] Equally important is the rise of egocentric data collection. Robots, autonomous vehicles, wearable devices, cameras, and other intelligent systems can now capture first-person observations while performing real work. Combined with simulation environments and synthetic data generation, this creates richer datasets for training Physical AI systems, improving their ability to navigate dynamic environments, understand workflows, and adapt safely to changing operating conditions.

05 Compute has reached the threshold Physical AI requires.

Purpose-built AI infrastructure (GPU clusters for physics-accurate simulation, edge hardware for real-time inference at the machine and robot level, and high-bandwidth connectivity between physical and digital environments) is now available at the scale and cost point that makes the closed loop viable at operational cadence. Five years ago it was not.

Together, these shifts make a fundamentally different operating model available and they are creating a race across industries. As with the shift from mobile phones to smartphones, or from analogue to digital cameras, the transition will not remain confined to one sector. Companies that build intelligent operations will be able to improve quality, productivity, efficiency, agility, and resilience at a pace others cannot match. Those that treat Physical AI as another technology program risk falling behind structurally, not temporarily. The organizations moving fastest are not doing so because they have better technology access. They are doing so because they have made the structural and governance decisions that let technology work as a system.

The operational framework for Physical AI

Physical AI becomes real when the closed loop is designed around an operational framework. At scale, the framework requires nine elements working together:

Intent and constraints	the “what” and the “must not.” Targets, operating policies, safety rules, cost and service constraints, and explicit decision rights: who or what can decide, when, and under which conditions. Without this layer, autonomy operates without boundaries.
Simulation-ready representation	the evidence engine. A digital twin with the fidelity required to test layouts, control logic, workflows, and failure modes before they are executed physically. Used for virtual commissioning, scenario replay, and readiness validation.
Hardware and sensing layer	the physical interface. Robots, machines, sensors, cameras, edge devices, controllers, and autonomous systems provide the perception and actuation capacity that allows Physical AI to operate in the real world. The capability of this layer determines what the system can see, understand, reach, move, inspect, and adjust. Without sufficient sensor coverage, reliable actuation, and hardware designed for integration, autonomy remains theoretical.
Compute layer	the processing foundation. Physical AI depends on the right distribution of compute across cloud, data center, edge, robot, and machine levels. Simulation and model training may require high-performance infrastructure, while real-time inference, safety decisions, and local coordination often need low-latency edge compute close to the physical process. The compute architecture defines where intelligence runs, how fast decisions can be made, and how the system scales from one cell to many sites.
Decision models and autonomy	the bounded “how.” Optimization and AI agents that propose or execute actions inside defined safety envelopes. Autonomy is trained and validated against constraints before it operates in production.
Orchestration layer	the system coordinator. A neutral coordination layer that turns decisions into action across people, machines, robots, and enterprise systems at operational cadence. This is where heterogeneous environments scale and where most programs that stall are missing a layer. Orchestration also determines how the compute foundation is used in practice: which decisions run centrally, which require low-latency edge inference, and how workloads are distributed as the program scales from one line to many sites.
Telemetry and learning	the feedback loop. Instrumentation that captures outcomes, variance, exceptions, and overrides, and feeds learning back into simulation, policies, and models. Without this, the system cannot improve.

Governance and change control	the safety rails. Evidence gates, auditability, rollback and override paths, and monitoring that make the system safe, explainable, and improvable over time.
Security and resilience	the trust boundary. Physical AI expands the attack surface because AI models, robotics, OT systems, enterprise systems, sensors, and edge devices become connected parts of one operating loop. Security must cover identity and access, model and data protection, OT cybersecurity, network segmentation, device hardening, monitoring, and recovery procedures. Resilience matters equally: the system must degrade safely, isolate faults, maintain critical operations, and recover without turning autonomy into operational risk.
Talent, upskilling, and change management	the human foundation. Physical AI changes roles, not just tools. Operators, planners, engineers, and maintenance teams need new skills, new workflows, and a clear understanding of how the system makes decisions and when they can override it. Research shows that more than three-quarters of frontline workers are dissatisfied with AI training, and many are uncertain about their place in the future operation. Organizations that scale Physical AI involve workers in designing how their roles will evolve, train people in the flow of real work rather than in abstract settings, and measure success based on how humans and AI perform together.[5]

Physical AI does not scale when these elements are handled as separate workstreams. It scales when they are designed together, with governance built in from the start. The business case is practical: operations that are more autonomous, productive, flexible, resilient, safe, and less constrained by labor availability. The impact shows up in better quality, higher throughput, faster ramp-up, lower downtime, stronger capital productivity, and faster response to change. But those outcomes depend on the full system working together: hardware, compute, simulation, autonomy, orchestration, security, governance, and people. Programs usually stall when orchestration is missing or when governance is treated as a late-stage control instead of a design requirement.

Physical AI is an ecosystem challenge

Physical AI is rarely constrained by the performance of a single robot, AI model, warehouse management system, or digital twin. Modern manufacturing and warehousing operations run on a complex ecosystem of specialized platforms, technologies, and partners spanning engineering, product lifecycle management, planning, manufacturing execution, logistics, automation, simulation, cloud infrastructure, and industrial operations. Increasingly, each layer is becoming intelligent in its own right through embedded AI, autonomous agents, and domain-specific decision support. The challenge is no longer generating intelligence within individual systems.[10]   new products or expanding to new sites is faster than it has ever been, the organization no longer rebuilds from scratch at each step.

The challenge is coordinating intelligence across the ecosystem. Physical AI therefore succeeds when organizations can transform isolated sources of intelligence into a connected operating capability that continuously senses, decides, executes, and learns across the value chain.

From system integration to systems of action

For decades, industrial enterprises focused on connecting systems and exchanging data. Physical AI raises the bar. The goal is not simply integration, but coordinated action. Manufacturing disruptions rarely remain isolated to a single function. An engineering change can affect production schedules, material availability, quality controls, maintenance priorities, logistics flows, workforce allocation, and customer commitments simultaneously. Organizations therefore need more than systems of record; they need systems of action that connect intelligence directly to execution.

The competitive advantage is ecosystem orchestration

As AI capabilities proliferate across software platforms, robot fleets, suppliers, logistics providers, and technology partners, local optimization becomes insufficient. Every participant in the ecosystem may be intelligent, but without coordination the overall operation becomes more complex rather than more effective. Leading organizations are therefore extending orchestration beyond enterprise boundaries to create shared operational context across partners, suppliers, maintenance providers, and logistics networks. Equally important is governance. Human accountability, decision rights, escalation thresholds, auditability, and security controls must be embedded into the ecosystem from the beginning. Organizations that successfully orchestrate platforms, partners, agents, and physical operations transform digital investments into measurable operational outcomes. Over time, this becomes a compounding advantage: learning from one site, workflow, or deployment can be reused across the broader network, allowing the enterprise to move faster, scale more efficiently, and continuously raise performance, resilience, and agility.

The ecosystem will shape the next wave of Physical AI

Physical AI will not be built by any single company, platform, or technology provider. Its evolution depends on a broad ecosystem of industrial software providers, simulation platforms, robotics companies, infrastructure leaders, cloud hyperscalers, hardware manufacturers, research institutions, and industry practitioners. As autonomous capabilities mature, the most significant innovations will emerge not from individual technologies but from how ecosystems work together to create interoperable, trusted, and scalable systems of action. Accenture is actively collaborating with clients, technology partners, industrial leaders, and ecosystem innovators to help define this next generation of Physical AI. While this first volume focuses on manufacturing and warehousing, future volumes in this series will explore how Physical AI is reshaping cities, homes, infrastructure, mobility, healthcare, and other domains of the physical economy. These future perspectives will be informed by ongoing collaboration with our ecosystem partners, bringing together diverse expertise, practical experience, and emerging innovations from across industries.

The investment logic

Physical AI creates value by moving learning upstream and reducing variance downstream. The economic case is not only higher throughput or lower downtime. It is also a lower cost to serve: the total cost of fulfilling demand reliably across production, warehousing, maintenance, labor, inventory, transport, quality, and service recovery. When operations can predict disruption earlier, rebalance capacity faster, reduce rework, improve maintenance reliability, and coordinate material flow across sites, they reduce the hidden costs that accumulate between the order and the customer promise. The same logic applies to safety and maintenance: fewer incidents, near misses, and uncontrolled interventions reduce disruption, exposure, and recovery cost, while predictive and increasingly autonomous maintenance shifts effort from reactive intervention to planned, evidence-based support.

> **Manufacturers deploying AI systemically report 30–40% reductions in unplanned downtime, 25–30% gains in maintenance capacity, and over 30% faster new product introduction. Companies pursuing autonomous operations project a 5% increase in EBITA and a 7% improvement in return on capital employed. The pattern is consistent across industries and starting positions.**[7]

01 **Virtualize before committing.**

Virtual commissioning is not new, but Physical AI raises what must be commissioned virtually. More sophisticated robots, autonomous systems, and flexible production environments require physics-based simulation engines that can test motion, perception, interaction, control logic, workflows, and operating scenarios before anything is changed physically. One global industrial technology company designed and tested an entire plant in the virtual world before laying a single brick, cutting lead times by 78% and time-to-market by 33%, while raising productivity by 14%.[8]

02 **Train autonomy under real-world conditions.**

Autonomous behaviors are developed, validated, and bounded using simulation, synthetic data, controlled pilots, and the hardware environment in which they will operate: compute infrastructure, robots, IoT devices, sensors, and connected machines. This reduces deployment risk and shortens the path from pilot to production by proving not only that the model works, but that it works with the physical systems that must execute it.

03 **Orchestrate execution as a system.**

Scheduling, maintenance, quality, and fleet coordination are managed as interconnected decisions, not isolated optimizations. Local improvements that degrade end-to-end performance are not improvements.

The outcome is more than efficiency. It is repeatability: decisions that can be explained, audited, and improved over time; operations that recover faster from disruption; and a foundation that can be replicated across factories, sites, and networks without rebuilding from scratch at each step. The learning from one site becomes reusable across many. Over time, that repeatable operating intelligence becomes part of the company’s core intellectual property. Competitiveness has long depended on hard assets, workforce capability, brand, and market access. Physical AI adds a new layer: the models, simulations, operating rules, and orchestration logic that encode how the company designs, runs, improves, and scales its physical operations. That know-how is difficult to copy because it is built through experience, validated in context, and continuously refined through use.

What it takes to make this work

Across industries and starting positions, the programs that scale share the same underlying structure as one integrated system:

A simulation-ready representation of the operation that supports what-if analysis, readiness checks, and evidence-based approvals before changes reach the floor. **KION**, working with Accenture, deployed digital twins of industrial warehouse environments to simulate and validate autonomous robot operations, testing configurations in physics-accurate simulations before implementing them in live facilities.

Autonomy that is reliable under constraints with explicit boundaries, validated safety envelopes, and human oversight where required. Accenture and **Schaeffler** developed a proof of concept demonstrating facility planning via digital twin, humanoid robot fleet testing, and live operations optimization with KPI comparison across simulated scenarios enabling site managers to identify issues before they affected production. Autonomy without governance is not a feature; it is a liability.

An orchestration layer that turns decisions into action coordinating people, machines, robots, and enterprise systems in real time, across vendors and platforms. Accenture's Physical AI Orchestrator provides this layer: a vendor-agnostic control plane that connects simulation, AI training, and live operations into a single closed-loop system across complex, mixed enterprise environments. **Belden** has deployed this to enable real-time worker safety monitoring and product quality inspection.

An operating model that makes this sustainable with clear ownership, decision rights, change control, and continuous improvement rhythms built in from the start. This includes the human dimension — and it is where many programs quietly fail. Research across 85 frontline workers in six industries found that more than three-quarters were dissatisfied with their AI training, and many were uncertain whether they would have a place in the future factory. The organizations that scale Physical AI address this directly, in three ways: they reduce uncertainty by involving workers early in mapping how roles will evolve; they train people in the flow of real work rather than in abstract settings; and they measure success based on how humans and AI perform together, not on technology adoption rates alone. What we consistently see across programs: workforce enablement is treated as an engineering requirement, not an afterthought.[9]

What we consistently see across programs: workforce enablement is treated as an engineering requirement, not an afterthought.[9]

Accenture's Physical AI Orchestrator

The operational framework described in this document is not a conceptual model. Accenture has built it. The Physical AI Orchestrator is Accenture's enterprise-grade orchestration product: a modular, vendor-agnostic control plane that connects digital twins, physics-based simulation, AI models, and real-world execution across physical operations. It is designed to operate across complex, mixed enterprise environments: multiple OEMs, legacy OT systems, heterogeneous robot fleets, and existing enterprise platforms.

The Orchestrator addresses the gap that most Physical AI programs hit at scale: the absence of a neutral coordination layer that can turn decisions into action across a heterogeneous operation without requiring single-vendor lock-in. It connects simulation, AI training, and live execution into one closed loop, but its value is broader than coordination. It provides the digital brain for physical operations: a modular architecture that connects data fabric, enterprise and OT systems, simulation-ready assets, physics-based digital twins, AI models, robot and IoT connectors, and governance controls. This allows companies to ingest and structure engineering and operational data, assemble reusable plant, warehouse, machine, and robot assets, test scenarios virtually, orchestrate execution in real time, and feed outcomes back into continuous improvement. The governance infrastructure (evidence gates, auditability, rollback paths, change control, and security controls) makes autonomous operation safe, explainable, and scalable.

In practice, the Orchestrator supports the full journey from design to live operation. It helps teams ingest and integrate CAD, PLM, OT, IoT, MES, SCADA, WMS, and asset-management data; build collaborative layout and process twins; use simulation engines for process, robotics, and scenario testing; monitor and orchestrate operations through AI agents and real-time alerts; and continuously optimize performance through reporting, benchmarking, and value capture. This is what turns Physical AI from a set of pilots into a software-defined operating model that can scale across lines, sites, and networks.

Core pillars of the Physical AI Orchestrator

Planning and Simulation

Planning and Simulation enables high-fidelity virtual modeling, performance prediction, and scenario validation before physical deployment. It de-risks operations by testing control logic, robotic behaviors, layouts, and operational policies using physics-grounded, bidirectional digital twins.

This pillar spans reality capture through live-linked simulation, creating the foundational environment in which Physical AI strategies are designed, validated, and continuously refined.

This capability is the first and most critical pillar, forming the bedrock for all downstream Physical AI intelligence and orchestration.

Robobrain

RoboBrain provides a unified AI and execution layer that manages the full lifecycle of robot intelligence—from data collection and simulation-based training to validated deployment and continuous runtime improvement. It turns robots from individually programmed assets into governable, reusable, and continuously learning enterprise systems.

This pillar focuses on training and operationalizing robotic brains on top of foundational and domain-specific models, building on active work with teams such as Kion.

Robobrain enables robots to learn, adapt, and make intelligent decisions in real time by continuously improving from simulation and operational feedback while operating within defined safety and governance constraints.

Orchestration

Unified Orchestration Platform (UOP) orchestrates rather than automates in silos. It coordinates execution across heterogeneous environments, operates within defined constraints, provides governance, traceability, evidence gates, and closed-loop control.

It unifies digital twins, robotic intelligence, vision systems, IoT data, and autonomous agents into a single, scalable, bidirectional runtime.

This pillar also ensures tight integration between Enterprise, Shopfloor and fleet management systems, enabling real-time feedback loops through vision and sensor data. It is the control plane that synchronizes planning, intelligence, and execution—ensuring Physical AI systems operate coherently, safely, and at enterprise scale.

The physical AI orchestrator solution: enabling AI to operate in the real world

Planning/ Engineering Twin

Operations Twin

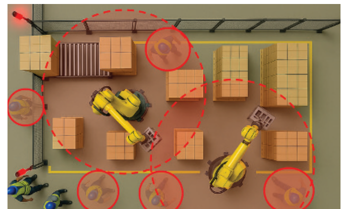

Vision AI

Robotics Validation

Where the physical AI orchestrator helps

Optimize new designs (greenfield) and confidently execute modernizations (brownfield)

Design + virtual commissioning	**Train + safely deploy physical AI**	**Run + continuously improve operations**
Before commissioning, validate • Layouts • Automation logic • Robotic behaviors In physics-aware simulation	• Train and validate robotics/ Physical AI behaviors in simulation • Then deploy with lifecycle controls • Orchestrate across mixed environments (IT/OT/robotics/ AI)	• Extend into real-time operations with closed-loop monitoring, decision support, and adaptive control as conditions change
• **More predictable launches and changes** • **Faster time to value for automation and autonomy**	• **Reduce time-to-training and time-to-deployment** • **Safer operations and a trained workforce**	• **Higher operational performance and cost efficiency**

What is in the physical AI Orchestrator solution

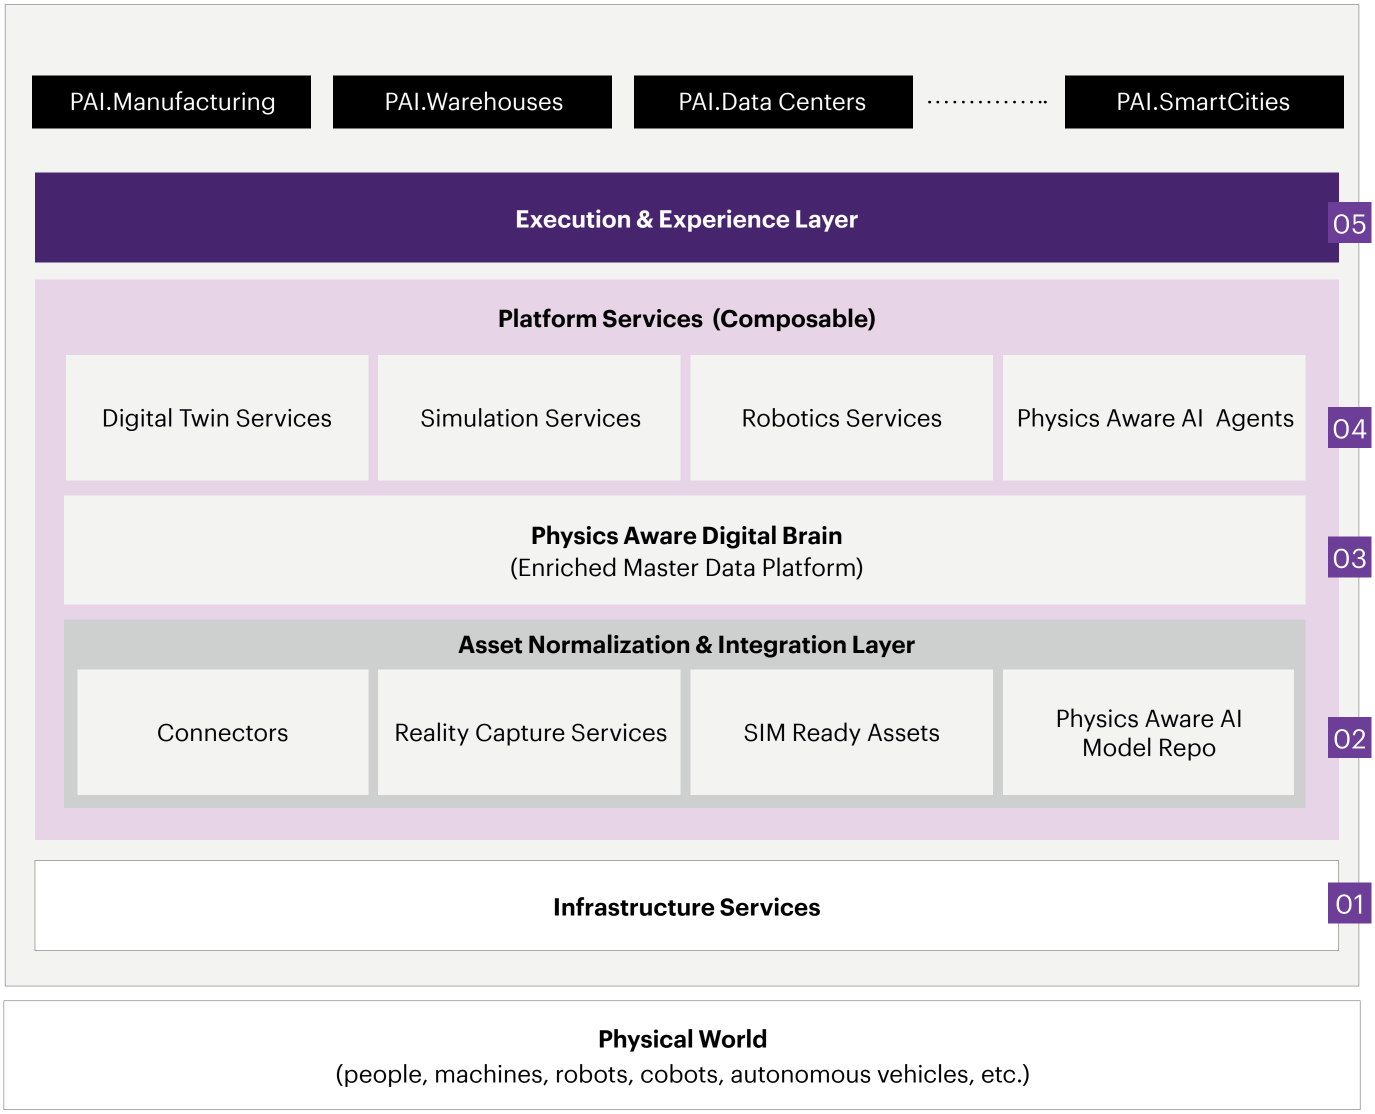

01 Provides the compute, data, and connectivity backbone required to run high-fidelity simulations, AI processing, and real-world system integration at scale.

02 Converts real-world data, engineering models, and AI assets into simulation-ready digital models — forming the foundation for accurate digital twins and physics-based intelligence.

03 Serves as the central decision and orchestration engine, allowing the system to sense what's happening, reason about options, act in the real world, and continuously learn across digital and physical environments.

04 Provides modular services for digital twins, simulation, and robotics, using physics-aware AI agents to evaluate scenarios and choose the best actions across both virtual and real environments.

05 Gives engineers and operators a single, shared workspace to design, test, and run factories, interacting directly with AI-driven digital twins of the physical world.

What leaders need to decide

Physical AI is a leadership decision. Organizations that scale successfully make a small number of choices explicit before the program begins:

Define the outcome and autonomy ambition. What are we designing for: flexibility, agility, resilience, reliability, safety, maintainability, cost to serve, or some combination of these? What level of autonomy is actually required to achieve that ambition?or cost to serve? And what level of autonomy is required to achieve it?

Set the automation boundary. Which decisions, processes, and operations can be automated, under which conditions, and where must human approval remain mandatory?

Design the human role. As autonomy increases, where will people approve, supervise, intervene, override, or improve the system? What skills and workflows need to change?

Govern override, rollback, and change. How will teams stop, reverse, test, and release changes safely? If the system cannot be overridden or rolled back, it is not production-ready.

Define the evidence gates. What must be proven before moving from simulation to pilot, and from pilot to scale? Evidence gates make scaling possible without multiplying risk.

Monitor performance and learning. How will leaders know the system is improving over time, not merely operating? Which signals will show whether safety, reliability, autonomy, and business performance are moving in the right direction?

These are organizational decisions. Made early and explicitly, they turn Physical AI from a series of pilots into an operating capability.

What Physical AI looks like when it works

Two to three years into a well-executed Physical AI program, the operation looks and feels different, regardless of industry setting.

Digital and physical work in sync. Every critical asset and process has a digital counterpart that is used continuously for planning and daily decision-making. Issues are predicted and resolved before they become disruptions. Meetings focus on improvement, not firefighting.

The workforce has changed alongside the technology. Operators supervise and guide autonomous systems rather than performing repetitive manual tasks. Planners use AI recommendations as readily as they once used spreadsheets. Maintenance technicians carry both wrenches and tablets. Critically, the people on the floor trust the system because they helped design it, they understand how it works, and they have seen it get better over time.

Performance gains are measurable and compounding. Lower cost per unit, higher throughput, faster commissioning, more reliable delivery, fewer safety incidents. And because the foundation is in place, launching new products or expanding to new sites is faster than it has ever been, the organization no longer rebuilds from scratch at each step.

> **This is the operating model that Physical AI makes possible. The technology to get there exists. The path is proven. What it actually takes is the decision to build it as a system and the discipline to govern it that way from the start.**

Physical AI at scale: reference guide for practitioners

This section translates the leadership agenda into scaling evidence: what must be proven before moving from simulation to pilot to production, and how the operational framework shows up in different industry contexts. It is written for leaders and their teams who need to connect strategy, governance, and execution across physical operations.

Start with the gate or industry pattern that matches your current challenge.

Evidence gates: what to prove before you scale

Physical AI scales through deliberate steps, with clear evidence required before each move. The first gate is organizational: leaders need a north star, a governance model, and a delivery structure before teams start testing autonomy, simulation, or orchestration. Skipping this foundation rarely saves time. It usually moves risk downstream, where it becomes more expensive to fix.

Gate 1: North star and governance readiness

The north star is defined: the business outcomes, operating objectives, and target level of autonomy are explicit.

Governance is in place: decision rights, risk ownership, safety accountability, escalation paths, and change-control principles are agreed.

The team structure is clear: business, operations, engineering, IT, OT, security, data, and workforce/change leaders have named owners and working rhythms.

The first domain is chosen deliberately: scope, value case, constraints, affected roles, and success metrics are agreed before technology work begins.

If these questions cannot be answered, the program is not ready for a technical gate. Physical AI needs a mandate and an operating model before it needs another model, twin, robot, or dashboard.

Gate 2: Technology readiness: simulation and orchestration

The decision or process is defined: inputs, outputs, frequency, constraints, owner, and expected action are explicit.

The simulation environment is ready: the digital twin can reproduce baseline behavior well enough to compare alternatives and validate readiness.

The orchestration layer is defined: the system can coordinate decisions across people, machines, robots, enterprise systems, and operational workflows.

The hardware, compute, and integration path is understood: robots, sensors, IoT devices, edge/cloud compute, OT/IT interfaces, and security requirements are mapped before touching the floor.

If this gate is weak, the pilot will test integration gaps instead of business value. It is far cheaper to find those gaps in simulation, orchestration design, and technical readiness than to discover them on the floor.

Gate 3: Pilot safety and performance

The system behaves predictably under normal and stressed conditions.

Human-AI interaction has been validated: where humans approve, supervise, intervene, override, or improve the system is defined and tested.

Exception handling is in place: fallback modes exist and have been exercised.

Evidence is captured: what happened, why, what the system learned, and what must change before expansion.

A pilot that produces no structured evidence is still only an experiment. This gate makes sure the learning is captured, not just the headline performance metric.

Gate 4: Replication readiness

Interfaces are standardized: OT/IT integration patterns, data contracts, and robot or fleet adapters are documented and repeatable.

Operating rhythms exist: review cadence, performance monitoring, and continuous improvement processes are embedded, not ad hoc.

Change control is established: who can change what, how changes are tested, and how they are released.

Security, safety, and audit requirements are met and verified across the target rollout environment.

Gate 4 is where scaling programs often underestimate the work. Replication is not repetition. A second line, site, or region only scales if the interfaces, governance, learning loops, and ownership model are genuinely portable.

Metrics to track early

Track a small set of early indicators for each domain rather than measuring everything. The most useful signals are usually commissioning duration, ramp stability, throughput variance, unplanned downtime variance, safety incidents and near misses, schedule adherence, quality escapes, mean time to recover, and override rate, including why overrides happen. Override patterns are especially revealing: a high override rate usually means either the system is not yet trustworthy or the governance model has not been communicated clearly enough.

Industry reference: the operational framework in practice

The operational framework looks different by domain. The examples below show where programs typically start, what the closed loop looks like in practice, and what a realistic first scope might be. They are pattern illustrations drawn from recurring experience across manufacturing, consumer, healthcare, hospitality, and infrastructure industries, not individual case studies.

Retail: virtual backroom optimization
Primary lever: simulation-ready representation, used to make layout and workflow decisions with evidence rather than store-by-store guesswork.
Where it starts: backrooms that are congested, hard to see into, and dependent on local workarounds.
Closed loop in practice: build a high-fidelity digital twin of the backroom; run repeatable simulations under consistent operating parameters; compare alternative layouts and workflows side by side; replay scenarios to validate decisions before physical changes are made.
Typical starting scope: one store format with one to two pilot locations. Codify the resulting playbook before rolling out across regions.
Manufacturing and warehousing: Physical AI as the operating system for operations
Primary lever: orchestration layer and unified digital twin for system-level decisions.
Where it starts: siloed engineering and operations data, fragmented tooling, and inconsistent decisions across design, commissioning, and run phases.
Closed loop in practice: consolidate key data sources into a single digital twin; use simulation and AI-driven insights to test changes before deployment; coordinate execution across assets and teams as one system.
Typical starting scope: one line, cell, or warehouse zone with clear decision rights and evidence gates established from the start.

Industrial safety and quality: virtual safety fences for human-robot collaboration
Primary lever: governance and safety rails combined with bounded autonomy.
Where it starts: safety risks in shared human-robot spaces; stop-the-line incidents; inconsistent enforcement of safety zones across shifts.
Closed loop in practice: use digital twins and computer vision to monitor worker movement in real time; define safety envelopes explicitly; automatically pause or adjust robotic operations when a human enters a hazardous zone; learn from near-miss patterns to refine zoning and policies over time.
Typical starting scope: one high-risk zone with explicit thresholds defined for alerts versus automatic intervention.
Asset-heavy industries: autonomous drone inspection
Primary lever: telemetry and learning, combined with end-to-end workflow orchestration.
Where it starts: dangerous and time-consuming manual inspections; limited coverage; delayed insight into asset conditions.
Closed loop in practice: automate the inspection workflow end to end — scheduling, autonomous flights, AI analytics; stream telemetry for rapid issue detection; integrate findings into asset management systems to drive maintenance decisions.
Typical starting scope: a repeatable route set for one asset class with a clear alerting and escalation policy before expanding coverage.
Logistics and distribution: warehouse layout generation and optimization
Primary lever: simulation-ready representation and constraint-aware decision models.
Where it starts: slow redesign cycles, late discovery of bottlenecks, and difficulty evaluating labor and automation tradeoffs as SKU and volume profiles shift.
Closed loop in practice: generate constraint-aware layouts from operational requirements; build them into an immersive digital twin; simulate throughput, congestion, and cycle time; iterate rapidly while preserving KPI targets before committing to physical changes.
Typical starting scope: one process area with defined constraints around space, safety, equipment, and service levels.
EV and advanced manufacturing: digital-first greenfield build and launch readiness
Primary lever: simulation-ready representation and evidence gates for readiness validation.
Where it starts: large new facility builds with many specialized equipment suppliers; high integration, commissioning, and ramp-up risk.
Closed loop in practice: construct a high-fidelity factory digital twin; integrate equipment models into a unified virtual environment; run simulations to identify spatial conflicts and workflow issues before physical installation; rehearse the factory repeatedly prior to go-live.
Typical starting scope: one critical value stream with evidence gates defined for “ready to install” and “ready to run” before either happens physically.

Warehousing and intralogistics: vendor-agnostic fleet management and virtual commissioning
Primary lever: orchestration layer for neutral coordination, combined with simulation-based validation.
Where it starts: mixed robot fleets; heavy on-site commissioning effort; downtime during changes; integration complexity with WMS.
Closed loop in practice: connect fleet management to a simulation environment; test maps, routes, tasks, and charging layouts virtually; validate performance before deployment; reduce vendor coupling by standardizing interfaces across the fleet.
Typical starting scope: one facility zone with one to two robot types. Expand to the full site once KPIs have stabilized and interfaces are confirmed portable.
Metals and process manufacturing: AI-powered melt shop control and disruption recovery
Primary lever: telemetry and learning combined with constraint-aware decision models.
Where it starts: capital-intensive operations where small disruptions cascade into yield and throughput losses.
Closed loop in practice: mirror live operations in a shared digital twin; replay scenarios to validate operating windows; use an AI agent to evaluate constraints and recommend corrective actions; execute approved scheduling changes to stabilize flow under pressure.
Typical starting scope: one constrained subsystem with clear human approval points defined before the system is trusted to recommend under disruption conditions.
Data centers: autonomous operations and physical infrastructure management
Primary lever: telemetry and learning combined with orchestration across physical and digital infrastructure.
Where it starts: manual inspection cycles, reactive maintenance on critical cooling and power infrastructure, limited real-time visibility across rack density and thermal load, and growing pressure to reduce energy consumption and downtime risk.
Closed loop in practice: build a digital twin of the physical facility including cooling, power, and networking infrastructure; deploy autonomous robots and sensor networks for continuous inspection and anomaly detection; use AI agents to optimize thermal management and energy distribution in real time; coordinate maintenance scheduling and hardware-handling workflows as a governed system with defined decision rights and override paths.

Summary

Physical AI is moving intelligence beyond the screen and into the physical world. It combines AI, simulation, digital twins, sensors, robotics, autonomous systems, enterprise platforms, and human expertise to help physical operations perceive, reason, decide, act, and continuously learn. While much of the early attention has focused on robotics, Physical AI is fundamentally a broader transformation: the creation of intelligent physical systems capable of translating business intent into safe, governed action across manufacturing, warehousing, logistics, infrastructure, and other physical environments. Building on advances in agentic AI, simulation technologies, cloud infrastructure, and autonomous systems, Physical AI is becoming a foundational capability for the next generation of industrial operations.

Manufacturing and warehousing represent some of the most advanced adoption environments for Physical AI today. Organizations are using digital twins to model facilities before investment decisions are made, simulating operations before changes reach the floor, coordinating fleets of robots and autonomous systems, and creating closed-loop environments where operational data continuously improves performance. These capabilities allow leaders to reduce downtime, accelerate commissioning, improve quality, increase throughput, strengthen resilience, and respond more rapidly to disruption. Physical AI shifts the focus from isolated automation initiatives toward intelligent operations that continuously improve over time.

However, technology alone is not enough. Physical AI only creates value when organizations design the full system. This includes simulation-ready digital representations, sensing infrastructure, compute capacity, autonomous decision-making, orchestration layers, governance frameworks, security controls, and workforce enablement. The organizations making the greatest progress view Physical AI not as a collection of individual use cases, but as an operating model. They connect planning, engineering, production, quality, maintenance, logistics, and workforce decisions into a coordinated system that can sense, decide, execute, learn, and adapt continuously.

The challenge is increasingly becoming one of orchestration. Manufacturers operate across a complex ecosystem of platforms, technologies, suppliers, logistics providers, and service partners. As AI and autonomous agents become embedded throughout this ecosystem, value will come not from individual systems acting independently, but from their ability to coordinate decisions and execute outcomes together. Physical AI therefore requires organizations to move beyond system integration and towards systems of action, where intelligence flows across enterprise functions and ecosystem boundaries while remaining governed, secure, traceable, and human-led.

This playbook explores what it takes to scale Physical AI successfully in manufacturing and warehousing. It presents the operational framework, governance principles, technology foundations, ecosystem considerations, and leadership decisions required to move from pilots to production. It also introduces Accenture's perspective on orchestration, closed-loop operations, and the convergence of physical, agentic, and systemic AI. Most importantly, it argues that Physical AI is not simply another technology cycle. It is the emergence of a new operating model for industrial enterprises, one that transforms intelligence into action and turns every operational cycle into an opportunity to learn, improve, and compete more effectively.

Physical AI glossary: a working vocabulary for operational leaders

The language of Physical AI is still taking shape. Terms such as “digital twin,” “orchestration,” and “autonomous operations” are used differently by vendors, analysts, and engineering teams. That ambiguity has a cost: slower decisions, misaligned programs, and investments that solve the wrong problem.

This glossary creates a shared operating vocabulary for AI-enabled physical systems used in production, material flow, logistics, and infrastructure operations. It is written for leaders accountable for safety, throughput, service levels, capital efficiency, and operational resilience.

Each definition focuses on three practical questions: how decisions are made, where control sits, and how the system behaves when conditions change.

Current as of August 2026.

3D reconstruction — is the process of generating accurate three dimensional models from sensor data such as cameras, LiDAR, or depth sensors. In industrial environments, it is used to digitally capture factories, warehouses, work cells, and equipment layouts with spatial precision. These models form the geometric foundation for digital twins, simulation, safety analysis, and robotic navigation. For plant leaders, the value lies in reducing guesswork when planning changes to layouts, automation, or material flow. High quality reconstruction allows teams to validate decisions virtually before committing capital or disrupting operations. Poor reconstruction quality directly increases downstream risk in simulation, commissioning, and autonomous execution.

AGV (Automated Guided Vehicle) — is a mobile automation system that follows predefined paths using markers, wires, or fixed guidance infrastructure. AGVs are commonly used for repetitive material transport in highly structured environments. Their behavior is predictable and reliable but limited in flexibility compared to autonomous mobile robots. AGVs are typically managed through centralized fleet or warehouse management systems.

AI Agent — is a software entity responsible for a specific decision or action domain, such as assigning work, prioritizing orders, or recommending routes. It perceives input from systems and sensors, evaluates constraints and objectives, and produces decisions or recommendations. In Physical AI systems, agents often work together, each focused on a narrow problem space. For Plant Managers, the key is not that agents exist, but what authority they are given. Well-designed agents operate within strict guardrails and escalate when confidence is low. This ensures autonomy improves throughput and stability rather than creating unpredictable behavior on the floor.

AI Auditability — is the ability to trace, explain, and govern AI driven decisions and actions. It includes visibility into data sources, model behavior, decision logic, and system outcomes. In physical operations, auditability is essential because AI decisions can affect safety, quality, and compliance. For executives, auditability enables accountability and regulatory confidence. It allows organizations to understand why a decision was made and whether it aligned with policy and constraints. Without auditability, scaling AI into physical execution becomes operationally and legally unacceptable.

Asset Library — is a reusable catalog of standardized digital representations of machines, robots, conveyors, and tooling. Assets include geometry, physics properties, behaviors, and operational metadata. Libraries enable consistent reuse across simulation, design, and operations. For executives, asset libraries reduce engineering effort and improve consistency across sites. They are essential for scaling digital twins and simulation programs without rebuilding models from scratch. Poorly governed libraries quickly become obsolete and undermine trust in simulation results.

Autonomous supply chains — self-orchestrating systems that leverage artificial intelligence (AI) and data to optimize operations, reduce costs, and improve resilience against disruptions. represents a significant shift from traditional automation, focusing on real-time decision-making and adaptability in a rapidly changing market environment. For leadership, this improves resilience and service levels under volatility. The challenge lies in integration across partners, systems, and organizational boundaries. Autonomy without shared governance increases risk rather than reducing it.

BOM (Bill of Materials) — is a structured list of components, materials, and subassemblies required to build a product. It defines quantities, relationships, and often versions across engineering and manufacturing views. BOM accuracy is critical to planning, costing, and execution. For plant managers, errors in the BOM propagate directly into shortages, rework, and missed deliveries. In Physical AI contexts, BOM data feeds simulation, planning, and material orchestration. Governance across EBOM, MBOM, and service BOM is essential for scale.

Collaborative robot (cobot) — robot designed for safe operation near humans, typically with force limiting and sensing. A collaborative robot is designed to operate safely in proximity to humans. Cobots use force limiting, sensing, and safety certified control logic to prevent injury. They support flexible automation by augmenting human work rather than fully replacing it. Cobots are widely used in assembly, inspection, packaging, and material handling.

Cross Site Simulation — is the practice of simulating operations across multiple plants, warehouses, or distribution nodes to understand network level outcomes. Instead of optimizing one site in isolation, it models how constraints and decisions in one facility affect upstream and downstream performance elsewhere. It typically includes interactions such as demand allocation, capacity shifts, transportation lead times, and inventory positioning—so leaders can test scenarios before committing. For a COO, this matters because many “local improvements” simply move bottlenecks to another site or increase cost to serve. Cross site simulation supports decisions like where to build what, when to shift volume, how to rebalance labor, and how to mitigate disruptions without overcorrecting. The credibility of results depends on consistent master data, realistic constraints, and clear agreement on what KPIs are being optimized (service, cost, resilience).

Digital thread — is a continuous, contextual data flow across the product and operating lifecycle. It connects design, engineering, manufacturing, and operations data into a unified view. The digital thread ensures traceability, consistency, and data continuity across stages. It is the backbone for digital twins and closed loop optimization. A data stream which communicates contextualized data across the various stages of the value chain. This data can be used to make more informed decisions at any point in time, and powers modern technologies such as digital twins; end-to-end data continuity across design build operate (connects engineering and operations).

Digital Twin — is a high-fidelity virtual representation of a real or planned physical asset, system, or process, continuously informed by real-world data. It is not just a 3D model; it includes behavior, constraints, performance characteristics, and operational context. In practice, digital twins are used to simulate scenarios, validate decisions, and predict outcomes before making changes on the shop floor. For operations leaders, the value of a digital twin lies in risk reduction and decision confidence. It allows teams to evaluate layout changes, automation strategies, and control logic without disrupting production. When connected to live data, the twin becomes a learning system, helping organizations understand why outcomes differ from plans and how to improve them over time.

Discrete event simulation (DES) — flow simulation where state changes at events (arrivals, machine cycles) to analyze throughput and bottlenecks. Discrete event simulation models of systems where state changes occur at specific events such as arrivals, machine cycles, or failures. It is commonly used to analyze production flow, queues, and bottlenecks. DES enables comparison of alternative layouts, schedules, and control policies. It is a core tool for throughput and capacity optimization.

Egocentric Data Collection — the capture of first-person observations from robots, autonomous vehicles, wearable devices, cameras, or human operators as tasks are performed. Unlike static operational datasets, egocentric data records how work unfolds in real environments and provides valuable training data for Physical AI systems. When combined with simulation and synthetic data, it helps improve perception, decision-making, and adaptive behavior in dynamic operating conditions.

Embodied Intelligence — is the broader concept that intelligence arises from physical form and interaction, not abstract reasoning alone. It emphasizes the role of perception, movement, and feedback. This principle underpins advanced robotics. For executives, embodied intelligence explains why purely digital AI cannot be directly transferred to physical operations. It highlights the need for training in realistic environments.

ERP (Enterprise Resource Planning) — systems, which manage planning, finance, procurement, and enterprise wide execution. They provide a system of record for orders, inventory valuation, and financial performance. ERP operates at business timescales rather than real time control. For COOs, ERP defines constraints and commitments that physical operations must respect. Integration with OT systems is essential for closed loop execution. ERP alone cannot manage physical variability without downstream systems.

Fleet Management System (FMS) — is the software system responsible for executing fleet management decisions. It assigns missions, manages traffic, and monitors execution. The FMS enforces operational rules at runtime. For leadership, the FMS is a critical control point for safety and efficiency. It must integrate with WMS, MES, and orchestration layers.
Weak integration limits scalability.

Humanoid robot — a general-purpose robotic platform designed to operate in environments built for people. Humanoid robots combine mobility, perception, manipulation, and reasoning rather than optimizing for one fixed task. Their potential value is highest where tasks are variable, labor is scarce, and redesigning the environment is difficult. They should be treated as a staged strategic bet, not a default automation answer: safety performance, task fit, training data, and operational reliability need to be proven before broad deployment.

Hyper-automation — is an approach to improving end to end business and operational processes by combining multiple automation methods—such as workflow automation, AI, and RPA—so improvements scale across functions rather than staying trapped in one team's tool. In operations, the key is that hyper automation targets the *process* (order to cash, plan to produce, maintenance execution), not just a single task like "auto generate a report." For COOs, this matters because isolated automation often increases complexity: it creates more handoffs, exceptions, and brittle dependencies. Hyper automation requires orchestration, governance, and clear ownership so automated actions remain aligned with policy, quality standards, and safety constraints. The most valuable outcomes are reduced decision latency, fewer manual touches, improved compliance, and more predictable performance under stress. If not governed, hyper automation becomes "automation sprawl" that is expensive to maintain and hard to audit.

Industry 4.0 — Industry 4.0, also called the "fourth Industrial revolution", combines new digital technologies such as the Internet of Things (IoT), cloud computing, AI (artificial intelligence) and machine learning with physical manufacturing and engineering, to improve the products we make and how we make them.

Inverse kinematics — the calculation of the joint positions a robot needs to place its tool, gripper, or end effector at a desired position and orientation. It is a basic requirement for reliable robotic motion because most tasks are defined as physical outcomes, not as joint movements. In production, inverse kinematics needs to account for reach, joint limits, collision avoidance, speed, and safety. Poor handling can lead to failed picks, slow cycle times, or unexpected motion paths.

IT/OT fusion (IT/OT convergence) — the integration of enterprise IT systems (ERP, analytics, cybersecurity) with operational technology systems (MES, automation, robotics). The goal is to enable real-time, data-driven decisions across planning and execution layers. This convergence is a prerequisite for Physical AI at scale. From a COO perspective, IT/OT convergence is as much an operating model change as a technical one. It requires clear system-of-record definitions, cybersecurity discipline, and ownership clarity. When done well, it enables faster response to disruptions; when done poorly, it introduces instability and risk.

LLM (Large Language Model) — is the reasoning engine behind many agentic systems. It interprets instructions, plans actions, and generates structured outputs. From a business point of view, it is a component—not the solution. Value emerges only when it is constrained, integrated, and aligned with operational realities.

Model predictive control (MPC) — control method optimizing actions over a future horizon under constraints.

Multi agent orchestration — is the discipline of coordinating multiple AI agents so that local decisions support global performance goals. It defines how agents share information, resolve conflicts, and respect enterprise constraints such as safety, labor rules, and service commitments. Without orchestration, agents optimize locally and often degrade end-to-end outcomes. In practice, this is one of the hardest problems in scaling autonomy. For COOs, effective multi-agent orchestration is what prevents “smart subsystems” from creating dumb factories. It ensures that autonomy remains aligned with KPIs like throughput, OTIF, and cost-to-serve.

Occlusion handling — is the capability to reason and act when objects or regions of interest are partially hidden from sensors—common in bins, racks, pallets, and cluttered workcells. In warehouses and plants, occlusions are not edge cases; they are normal operating conditions driven by packaging, stacking, lighting, and human activity. For robots and perception systems, occlusions can cause misidentification, failed grasps, navigation uncertainty, or safety perception gaps. For plant leadership, the impact shows up as cycle time variability, reduced pick success rate, and increased human interventions that erase ROI. Strong occlusion managing typically relies on multi sensor inputs (vision + depth/LiDAR), temporal reasoning ("what was there a moment ago"), and conservative safety behavior when uncertainty rises. If this is not addressed explicitly, "works in demo" autonomy often fails in real operations.

OEM Agnostic Integration — ability to connect and manage robots from multiple manufacturers within a single orchestration layer.

OpenUSD / USD — common 3D scene description format used to share assets and simulation environments.

Orchestration Layer — the connective "system layer" that governs how decisions are coordinated across systems and time horizons, from planning to execution. It defines priorities, resolves conflicts, and ensures actions comply with safety, labor, and business rules. For COOs, orchestration is the difference between scalable autonomy and uncontrolled automation. Without it, systems optimize locally and degrade end to end performance.

PDM (Product Data Management) — Product data sources referenced alongside PLM for conversion into USD with structure/metadata preserved.

Photogrammetry — reconstructing 3D from overlapping photos (common for sites/objects).

Physical AI — AI conducting human intent in the real world. Enabled by AI systems that perceive, reason, and act via sensors and actuators. In manufacturing and warehousing context this includes the design, commissioning, building, testing, running, and maintenance of autonomous operations with humans in the lead.

Physical AI Orchestrator — A modular, enterprise grade orchestration solution connecting digital twins, physics-based simulation, AI models, and real-world execution across physical operations, developed by Accenture. It enables Physical AI to move from isolated deployments to a software defined operating model that can scale across sites, systems, and technologies.

Powered Industrial Vehicle (PIV) — is an industrial vehicle—commonly forklifts, tuggers, or pallet trucks—that moves materials and inventory within facilities. PIVs are essential to operations but are also a significant source of safety risk due to speed, blind spots, mixed pedestrian traffic, and variable operator behaviors. In modern safety monitoring scenarios, PIVs may be tracked and contextual alerts can be delivered to operators or supervisors when risk conditions emerge. For plant leadership, the goal is reducing near misses and incidents while maintaining flow and productivity. The key is that monitoring must translate into actionable interventions (speed reduction zones, routing adjustments, training insights), not just dashboards. When done well, PIV safety programs improve both safety KPIs and operational stability by reducing disruptive events and investigations.

Product and operating value chain — all the steps that a business uses to create a product or service, from initial design through to delivery to the customer and return and recycling.

Reality capture — capturing the real environment (factory/warehouse) into a digital model using scans/images.

Rebalancing — shifting production/labor/logistics across sites dynamically based on insights from connected systems and simulation.

Retrieval-augmented generation (RAG) — is a design pattern where an AI model retrieves enterprise knowledge to ground its outputs. With RAG, the AI does not rely on generic training data; instead, it pulls answers from our own approved operational sources—work instructions, SOPs, maintenance logs, quality records, engineering specifications, and current production data. From a plant manager's standpoint, this is what separates a risky AI experiment from a system that can be trusted on the shop floor. It ensures that when the AI advises or acts, it is aligned with how this plant runs today, not how a textbook describes operations or how another company works.

Robot Foundation Model — is a large, reusable AI model trained across many tasks and environments. It provides general capabilities that can be fine-tuned for specific robotic applications. This approach reduces training time and increases scalability.

Robot learning — is the mechanism through which physical systems transition from pre programmed automation to self improving operational assets. Accenture emphasizes learning approaches that combine simulation, synthetic data, and real world feedback to accelerate deployment while minimizing risk. Training robots in digital environments allows organizations to scale experimentation, manage edge cases, and optimize performance before physical rollout. Robot learning also enables faster reconfiguration of operations when layouts, products, or demand patterns change. Within Accenture's autonomous operations strategy, robot learning is essential for achieving resilience, adaptability, and continuous improvement at scale.

Robot safety zone — dynamic or static zones that restrict robot motion when humans enter.

Robotic Arm Manipulator — automation referenced for complex manipulation, coordinated within an orchestrated environment.

Route Optimization — is the process of determining the most efficient path or sequence for moving people, vehicles, materials, or robots through a facility or network. In warehouse and manufacturing environments, route optimization balances objectives such as travel distance, congestion avoidance, priority handling, and energy use against operational constraints like safety zones, equipment availability, and service windows. For plant managers, poorly optimized routes are a hidden source of throughput loss — congestion, conflicting traffic, and suboptimal sequencing compound over thousands of daily movements. AI-driven route optimization continuously recalculates paths based on real-time conditions rather than fixed rules, enabling faster response to disruptions such as equipment failures, traffic spikes, or priority changes. The governance requirement is equally important: changes to routing logic must be validated before deployment, and override and fallback behaviors must be defined to ensure safe operation when the optimization system encounters conditions outside its training.

SDV (Software-defined vehicles) — vehicles that rely on software to manage various aspects of their functionality. For example, engine control, infotainment systems, and autonomous driving. The software can be updated automatically to continuously improve driver's experience, add convenience, and improve safety.

Servitization — is the shift from selling physical products to delivering outcomes and services enabled by those products over their lifecycle. Instead of focusing primarily on one time asset sales, organizations monetize availability, performance, efficiency, or usage through service contracts. This model is typically enabled by connected products, embedded software, and continuous data collection from assets in operation.

SimReady (Simulation Ready Assets) — are the foundation for reliable Physical AI, ensuring that digital twins behave like their real world counterparts. Accenture views simulation ready assets as a prerequisite for trustworthy decision making, virtual commissioning, and AI training. When assets include accurate geometry, physics, and operational data, organizations can confidently validate changes before executing them physically. Standardized, reusable simulation assets also enable collaboration across engineering, operations, and ecosystem partners. In Accenture's digital engineering and manufacturing work, simulation ready twins accelerate capital projects, reduce rework, and support continuous optimization of physical operations.

Smart connected products — physical items that contain intelligent, "smart" components (i.e., sensors) and connectivity parts like antennae. Technology unlocks new value for manufacturers and customers by enabling a hyper-personalized experience for the customer. Smart connected products let the product and user share and analyze data. They can also connect data from business systems, outside sources and other related products.

Smart factory — a manufacturing facility that uses technologies such as artificial intelligence (AI), robotics, machine learning, automation, cloud, edge technology, and digital twins, to improve production processes. Smart factories demonstrate higher productivity, cost efficiency, and better product quality.

Smart Material Flow — is an operating concept where real time visibility, decision intelligence, and (often) digital twins are used to keep material moving efficiently as complexity and variability increase. It connects sensing (where inventory is), orchestration (what should move next), and execution (how it moves via people, conveyors, forklifts, or robots). In practice, it targets the most common flow killers: staging congestion, poor prioritization, misrouted inventory, and lack of synchronized replenishment. For plant leaders, the value is improved throughput without adding square meters or headcount—because smarter flow reduces waiting and rework. Smart material flow also improves resilience by enabling rapid rerouting when aisles are blocked, machines go down, or priorities shift. The concept only works if it is integrated with systems of record (WMS/MES/ERP) and governed by clear service and safety constraints.

Synthetic Data — is data generated artificially—often through simulation—to expand coverage of scenarios that are rare, unsafe, or expensive to capture in real operations. In Physical AI contexts, it is used to train and validate perception and decision models across broader variability than historical plant data provides. For operations leaders, synthetic data matters because it reduces the "unknown unknowns" that show up during deployment: unusual lighting, rare obstructions, odd part placements, or atypical traffic patterns. The risk is realism—if synthetic data does not reflect true physics and operating conditions, models can learn the wrong behaviors with high confidence. Therefore, synthetic data must be validated against real distributions and used as augmentation, not fantasy replacement. When governed well, it accelerates safe learning and reduces pilot risk; when governed poorly, it produces brittle systems that fail in production.

Tool-Using Agent — is designed to interact directly with operational and enterprise systems. It acts as a bridge between AI reasoning and real execution across MES, ERP, CMMS, or data platforms. The business value lies in reducing integration friction and manual system navigation. Control over tool access is critical to prevent unintended system changes.

USD (Universal Scene Description) — is a standardized format for representing complex 3D environments as structured scenes—including geometry, object hierarchies, materials, coordinate systems, and metadata—so multiple tools can share the same "source of truth" for a digital environment. In industrial simulation, USD is valuable because it supports assembling large factory or warehouse scenes from reusable assets (machines, racks, robots, conveyors) without losing structure. This enables teams to build simulation and digital twin environments that remain consistent across engineering, visualization, and physics based testing workflows. For COOs and plant leaders, USD matters because it reduces rework and fragmentation: instead of rebuilding environments for each tool or partner, the organization can reuse and govern scene content as a scalable asset base. The operational payoff is faster iteration—layout changes, commissioning tests, and robotics validation can be executed with less translation overhead and fewer data losses. The governance requirement is discipline: naming conventions, version control, and metadata standards must be enforced, otherwise scene reuse collapses and teams revert to one off models. In short, USD is not "graphics"—it is an interoperability backbone for scalable simulation and digital twin programs in physical operations.

Virtual commissioning — is the practice of testing automation logic, equipment behavior, and system integration in a simulated environment before physical deployment. It combines digital twins, control logic, and realistic operating scenarios to expose issues early. This shifts problem-solving from the plant floor to the engineering phase. For operations leaders, the value is very concrete: fewer surprises during startup, shorter ramp-up times, and lower rework costs. Virtual commissioning also enables safer experimentation, allowing teams to validate changes without putting people, equipment, or production at risk. Over time, it becomes a standard gate for any significant operational change.

Virtual Safety Fence — is a dynamic, software defined safety boundary around robots, vehicles, or hazardous zones that triggers slow/stop or restricted behaviors when humans enter defined risk areas. Unlike fixed physical barriers, virtual fences can change based on operating mode, task context, or real time location sensing. For plant management, the purpose is to enable productive shared spaces without compromising safety—especially where flexibility is required and reconfigurations are frequent. The operational risk is that virtual fences depend on reliable sensing and validation; if perception fails, safety integrity is threatened. Therefore, virtual fencing must be engineered with conservative assumptions, clear fail safe behaviors, and alignment to the site's safety standards and certification requirements. Done well, it reduces incidents and unlocks more flexible layouts; done poorly, it either becomes unsafe or is so conservative that it destroys productivity.

Vision-Language-Action (VLA) Model — is a model class that connects what a system sees (vision), what it can understand or be instructed with (language), and what it can do (actions). In physical operations, this enables more flexible task execution, e.g., interpreting high level instructions and grounding them in the current scene context. The operational opportunity is faster changeovers and reduced engineering effort for "every new SKU equal reprogramming." However, the operational risk is equally real: language flexibility can introduce ambiguity, so VLA models must be bounded by strict safety and policy constraints before they can influence actuation. For plant leadership, the right deployment posture is often advisory or supervised at first, with clear escalation paths when confidence is low. VLA becomes valuable when it improves flexibility without compromising predictability, auditability, and safety performance.

Vision-Language Model (VLM) — an AI model capable of jointly understanding images, video, and natural language. In Physical AI environments, VLMs enable systems to reason about scenes, equipment, workflows, anomalies, and human instructions using both visual and language inputs. They provide the perception and contextual understanding layer that increasingly supports robotics, inspection, safety monitoring, and autonomous operations.

Warehouse management system (WMS) — is the execution system that governs inventory location, picking, replenishment, receiving, shipping, and the workflows that keep warehouse operations coordinated. It acts as the system of record for where inventory is and what tasks should happen next, typically under service and accuracy constraints. For plant and supply chain

Warehouse management system (WMS) — the execution system that governs inventory location, picking, replenishment, receiving, shipping, and the workflows that keep warehouse operations coordinated. It is typically the system of record for where inventory is and what tasks should happen next. WMS discipline determines whether robotics and automation can scale reliably, because robots cannot compensate for poor inventory accuracy or inconsistent process execution. Integration is critical: robot fleets and automation systems need to align with WMS tasking and confirmations to avoid competing versions of operational truth.

Worker Agent — a specialized digital role focused on a narrow, well-defined task, such as monitoring quality data, validating production reports, summarizing maintenance logs, or flagging exceptions for review. Worker agents reduce administrative load and help skilled employees focus on higher-value work. Their strength is reliability and repeatability, not broad autonomy. They should operate with clear task boundaries, traceable outputs, and human escalation when judgment is required.

References

1. Accenture (2025): **Blueprint for success**.
2. Benefit ranges shown reflect project experiences, client examples, and industry research. Actual outcomes depend on starting maturity, scope, and operating context.
3. Accenture (2025): **Rethinking the course to manufacturing's future**.
4. Accenture experience; actual timelines may vary
5. Accenture (2025): **Rethinking the course to manufacturing's future**.
6. Accenture (2025): **Rethinking the course to manufacturing's future**.
7. Accenture (2026): **Systemic AI at the root of manufacturing performance**.
8. Accenture (2026): **Systemic AI at the root of manufacturing performance**. Example refers to a global industrial technology company; outcomes are specific to that program and will vary by context.
9. Countryman, T., Oosterhuis, I., Wheless, J., and Afzal, R. "**The Best Manufacturers Build AI with Workers, Not for Them**." Harvard Business Review, May 2026. Based on an unpublished study of video diaries from 85 frontline workers across six industries.
10. Accenture (2026): **Orchestrating the ecosystem Turning digital investments into industrial outcomes**

About Accenture

Accenture helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed for organizations across industries. Our strategy is to be the reinvention partner of choice for our clients and lead in the safe, widespread adoption of AI, and to be the most client-focused, AI-enabled, great place to work in the world. We bring together the talent of our approximately 799,000 people with proprietary assets and platforms, deep process and industry expertise, and leading ecosystem relationships to deliver end-to-end solutions and measurable outcomes at scale. Through our Reinvention Services, we offer broad expertise across Cybersecurity, Digital Core, Finance, Industry and Enterprise, Song, Supply Chain and Engineering, and Talent, with advanced capabilities in AI and Data, Industry and Process, and Technology. We serve approximately 9,000 clients and generated approximately $70 billion in FY25 revenue. Visit us at **accenture.com**