accenture

Systemic AI at the root of manufacturing performance

The new growth infrastructure for manufacturers

Authors

Tracey Countryman

Supply Chain and Engineering Global Lead, Accenture

Tracey Countryman leads Accenture's global Supply Chain and Engineering business, shaping the firm's strategy, investments and innovation agenda across industries. She brings a 30-year career in manufacturing and supply chain transformation and has helped build and scale Accenture's manufacturing and supply chain capabilities over time. Her experience spans network strategy, supply chain planning, manufacturing transformation and large-scale programs that industrialize digital and AI capabilities from foundational data and platforms through plant-level execution.

LinkedIn

Prasad Satyavolu

Manufacturing and Operations Global Lead, Accenture

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

Roland Mayr

Industry and Enterprise, Industrials, Global Lead, Accenture

Roland Mayr leads Accenture's global Industrials industry, advising executive teams on reinvention with digital, data and AI across complex asset structures and process chains. With over three decades at Accenture, he brings deep automotive experience and has held senior leadership roles across the firm, including within Accenture Song and global client account leadership. He is known for a strong customer focus, helping companies move beyond products to deliver relevant, meaningful and sustainable experiences.

LinkedIn

Luis Luque

Operational Technology Security Global Lead, Accenture

Luis Luque brings more than 25 years of experience at the intersection of industrial technology and cybersecurity. He leads a cross-industry global team focused on securing critical infrastructure, from legacy manufacturing and utilities environments to AI data centers and robotics infrastructure, building risk-based, AI-driven capabilities for sensitive industrial settings. Prior to joining Accenture, he co-founded Cimation (acquired by Accenture), where he helped shape industrial IT, automation and ICS security services.

LinkedIn

The era of AI potential is over. We're now in the era of AI scale.

A global industrial technology company designed and tested one of its plants entirely in the virtual world before laying a single brick. Concentrated, committed investment delivered decisive results: That company cut lead times by 78% and time-to-market by 33%, raised productivity by 14% and reduced carbon emissions by 28%. Net climate impact will vary by energy mix and the AI compute footprint over time.

Across facilities operating today, audited results show that organizations deploying agentic AI in the supply chain recover from disruption 60% faster, carry 22% less inventory and reduce cost of goods sold by 5%. A leading French automaker generates over $140 million per year from its AI program, while a Chinese appliance manufacturer estimated its AI program ROI over the next several years at over 100%—with AI now embedded in 30% of all employee KPIs.

These results raise a practical question: How do manufacturers move from isolated wins to an operating model that scales AI across a network? To answer it, we spoke with 36 manufacturing executives and experts across Europe, North America and Asia-Pacific. We wanted to understand how manufacturers were embracing AI and rethinking how they did business with these new technologies. What leaders at the top companies showed us was far more than a series of AI pilots: it was a whole new way of building their organizations. All stats and company examples are drawn from the interviews unless otherwise cited.

Here's what we learned. Top performers have moved beyond fragmented AI deployments to what we call **systemic AI**: intelligence that operates as the production system itself, not alongside it. It unites decision-making with physical execution across machines, lines, plants and supply networks.

Unlike earlier Industry 4.0 programs, which digitized assets and automated islands of work, systemic AI links key performance indicators (KPIs), decision rights and escalation paths so the system can act repeatedly and safely across sites without rebuilding from scratch each time.

Earlier enterprise systems required years of specialist implementation before value appeared, then plateaued. AI gets more capable with every operational cycle: it's a system that improves as it runs, with each cycle making the system faster and less expensive to extend.

The proof is already compelling. Manufacturers deploying AI systemically report significant efficiency gains, including 30–40% reductions in unplanned downtime, 25–30% gains in maintenance capacity and over 30% faster new product introduction.

AI is becoming a primary growth engine for these manufacturers, with faster product launches, more resilient supply chains and competitive advantages that compound across every plant in the network.

> Accenture Research indicates companies pursuing autonomous operations project improvements on financial metrics as well, including a 5% increase in EBITA and a 7% improvement in return on capital employed.[1]

Through our conversations with these leaders, we surfaced five dimensions indicative of success across their companies and actionable today for building toward systemic AI. They are:

- **Integrate planning, production, quality and logistics** so decisions flow end to end
- **Build shared data, platforms and guardrails** so you don't rebuild for every plant
- **Redesign decision rights and operating rhythms** so AI is part of daily work
- **Connect physical and agentic AI** to create the closed loop
- **Design for humans in the lead** to define accountability as autonomy scales

Each of the dimensions requires continuous, concurrent attention. Each removes a different bottleneck that would otherwise break the AI “loop” of sensing, deciding, executing and learning. Miss one and the system stalls. Get them right, however, and AI becomes your competitive advantage by compounding value across every site in your network.

And this advantage, once built, is structural. Once a company can replicate systemic AI capabilities from one site to the next, every new deployment gets faster, widening the gap between companies that can scale AI across their network and those that can’t.

As one vice president of digital manufacturing platforms put it, “The gap is not ambition—it is the willingness to fund the operational infrastructure that makes AI durable."

This report examines systemic AI in detail, demonstrating what a mature capability looks like across the lifecycle of a plant. It then looks more deeply at the five dimensions of successful systemic AI in manufacturing and offers leaders actions they can take now to build a competitive advantage through AI.

The manufacturers who will reach systemic AI first are not the ones with the cleanest data or the most aligned workforce today. They are the ones who start with a clear vision of scaling AI. Every month spent waiting for perfect conditions is a month that movers will compound their advantage.

Beyond the pilot, across the cycle

In our interviews, manufacturing leaders largely focused their AI efforts on operations and maintenance, where data is richest and return on investment (ROI) is easiest to see. But that focus misses the bigger opportunity.

Manufacturing value accrues across the full lifecycle of a plant, from design through construction, commissioning, ramp-up and decades of operation. An AI program that targets only the operating phase is likely to stall. Fragmented data, unclear ownership and workforces not yet designed for AI aren't late-stage problems; they run through the entire lifecycle.

Executives pointed to the biggest step-change upstream: design, build and commissioning. AI applied here locks in cost and performance before capital is committed, catches mistakes before teams rework physical assets and compresses time-to-market. Digital twins help teams test factories before construction begins; virtual commissioning reduces risk at the most fragile handoff to operations.

As one leader told us,

> **"If your product lifecycle, from initial design to serial manufacturing, is one year, and you cut it in half, that is a huge competitive advantage."**

The same logic applies inside facilities that are already running. For most manufacturers, the harder and more immediate question is how to make smarter capital decisions in the plants they already own. Questions such as when to modernize, where to debottleneck, how to sequence upgrades—AI is now changing the calculus on all three.

Full-cycle value draws a sharp line between AI pilots and systemic AI: a “closed loop” in which AI continuously senses, decides, executes and learns. The loop depends on three converging AI capabilities: generative AI to accelerate knowledge work; agentic AI to coordinate decisions and actions across functions; and physical AI to embed intelligence in the real world. Together, they let manufacturers sense disruption, respond in real time and improve with every cycle, creating compounding value rather than one-off wins.

Defining systemic AI in manufacturing

Manufacturers face whole-cycle problems. Systemic AI addresses them across two fronts: maturity and lifecycle. First, it changes how AI matures from disconnected pilots into a repeatable operating capability. Second, it extends value creation across the full-plant lifecycle, from design through decades of operation.

Systemic AI maturity

Most large manufacturers have pilots running across plants, functions and geographies. Those pilots rarely translate into sustained, scaled performance because teams build them as one-off point solutions. They rely on custom integrations, inconsistent definitions and local governance, requiring each new deployment to start fresh. One plant's "scrap" can be another plant's "rework," and teams often label the same signals differently across lines, which blocks reuse even when the model logic is sound.

Systemic AI replaces that pattern with an operating capability that teams can deploy, govern and improve across sites. The comparison below shows the shift from the starting point for most manufacturers to the destination leaders build toward (Figure 1). Companies typically progress through three stages.

In the pilot stage, teams prove value in a handful of use cases, usually within a plant or a function. They depend on manual workarounds and heroics to keep models running, and they struggle to reuse data and logic elsewhere. As they progress, leaders invest in shared foundations.

They standardize KPIs and data models, connect information technology (IT) and operational technology (OT) data flows, define decision rights and escalation paths and put model oversight into day-to-day operations. Teams can replicate use cases across sites with decreasing effort.

By the time it reaches systemic AI, the organization runs closed loops in production at scale. Agents coordinate decisions across planning, production, quality, maintenance and supply. Governance and guardrails keep autonomy safe, and learning from one line propagates across the network.

Figure 1

Aspect	Piloted AI—Where Most Manufacturers Are	Systemic AI—Where Leaders Are Heading
What It Looks Like	Disconnected use cases. Value is local and fragile.	AI woven into the operating model. Value compounds continuously.
How AI Works	Point solutions stand alone. Each pilot requires its own data and rebuild.	Composite agents connect shop floor to boardroom in closed loops.
Who Leads	Individual champions like a data scientist, a plant manager or CTO.	Leadership treats AI as a strategic asset. CEO sets the ambition.
Data	Siloed, inconsistent, largely paper-based or locked in legacy MES/ERP.	Unified IT/OT data fabric across sensors, MES, ERP and PLM in real time.
Scale	Hard to replicate beyond isolated pilots. Each plant starts over.	Use cases replicate across plants, geographies, functions with decreasing marginal effort.
Governance	Unclear ownership. IT/OT accountability contested. ROI hard to prove.	AI governance embedded: clear ownership, model oversight, measurable KPIs.

Use cases rarely limit progress. Foundations do. Companies that reach systemic AI treat AI as infrastructure. They build an industrial data fabric, define governance and guardrails, assign ownership for outcomes and run performance management against shared KPIs. That sequencing turns local proof into network-wide reuse.

Maturity matters because it determines where value can show up. Piloted AI concentrates benefits in operations, while systemic AI extends those benefits upstream to lock in cost, speed and resilience across the entire lifecycle of a plant. One Chinese manufacturer considers R&D and production one of their strongest opportunities for scaling AI, citing a reduction in design lead time from 30 to 3 days, while a US aerospace manufacturer found that enabling AI in the notoriously challenging eBOM to mBOM handoff eliminated traceability errors that had previously propagated unchecked into production.

Systemic AI across the plant lifecycle

Not every use case is ready for systemic AI on day one. Suitability depends on risk, the level of ongoing human supervision required and the maturity of third-party platforms and vendors. Many AI programs start in operations, maintenance and quality because teams can instrument assets, capture signals and link outcomes to downtime, scrap and throughput. Systemic AI keeps that value and expands AI into the earlier phases that set cost and performance before production starts.

Manufacturers create value across six phases: design, build, test, commission, operate and maintain. Systemic AI strengthens decisions in each phase and preserves learning from one phase to the next. Time-to-value differs sharply across lifecycle phases, and upstream applications often create disproportionate strategic advantage because they lock in cost and performance before production starts (Figure 2). In practice, upstream wins take longer to build, but they lock in performance and compress future ramp-ups, which is why leaders treat them as a source of durable advantage.

Figure 2

Design	Build	Test	Commission	Operate	Maintain
80%	**+20%**	**Weeks**	**3-18**	**5-15 pts**	**-23%**
Engineering lead time	Productivity	vs. months of physical trials	Months for line ramp-up	OEE improvement	Unplanned downtime
Chinese Appliances & Robotics Manufacturer: 30 days to 3 days for engineering drawings	A Global Industrial Technology Company: +30% volume flex, +40% space efficiency	Synthetic data validates robots before shop floor deployment	A Major European Automaker: new model to full-rate production in 3 months	North American Semiconductor Leader: 30% quality error reduction	Industry average: AI-powered predictive maintenance

In **design**, AI accelerates engineering work, improves requirements quality and strengthens the engineering-to-manufacturing handoff. In **build** and **test**, digital twins and synthetic data let teams validate layouts, workflows and automation logic before they commit capital or disrupt a running line. In **commissioning**, virtualization dramatically reduces the potential for integration failures across equipment, manufacturing execution systems (MES), enterprise resource planning (ERP) and control systems and shortens ramp-up.

In **operate** and **maintain**, closed-loop systems (implemented with appropriate ongoing supervision and monitoring) reduce downtime, improve yield and help plants respond to quality drift and demand changes within the shift. Leaders also need a sequencing lens that reflects where value concentrates in discrete manufacturing and how quickly it shows up. The executives we spoke with identified five priority value areas with the highest concentration of investment and the most traceable metrics, along with typical time to return (Figure 3).

Figure 3

	Domain	Value at Stake	Why It Matters	Deployment Readiness	Time to ROI
1	Supply Chain Planning **Upstream multiplier**	**60%** Time-to-recover reduction[2]	Highest long-term impact. Creates resilience, not just efficiency.	**Platform-based**	**12-24 months**
2	Quality Intelligence **Most proven**	**30%+** Defect reduction	Most direct line from AI deployment to traceable cost savings.	**High readiness**	**6-12 months**
3	Asset Performance **Universally applicable**	**23%** Less unplanned downtime	Deployable in parallel with any other initiative. Immediate measurable return.	**High readiness**	**6-18 months**
4	Engineering and New Product Introduction **Competitive weapon**	**2x faster** Product development	Separates next-gen manufacturers from those in slow product cycles.	**Moderate-high**	**18-36 months**
5	Procurement and Input Costs **Fastest payback**	**3-8%** Cost improvement on spend	50-70% of COGS is direct materials. Data is clean. Tools are ready.	**Highest readiness**	**3-6 months**

In our conversations, leaders described the strongest step-change opportunity upstream. When a manufacturer cuts the cycle from initial design to serial production in half, it changes competitive position through faster launches and faster learning. Teams achieve that compression when they use AI to translate engineering intent into manufacturable plans, validate designs in simulation and reduce commissioning risk before teams touch physical assets.

AI and the brownfield CapEx decision

For most manufacturers, the most consequential AI-enabled capital question is not how to build the perfect new factory. It is how to make smarter decisions inside the plants they already own.

Brownfield CapEx is structurally hard: production cannot stop, equipment is interdependent and the cost of a wrong decision is irreversible. AI changes this by making it possible to simulate before committing. A proposed layout change, a debottlenecking intervention, a new line configuration—all can be tested in a digital twin before a single piece of metal moves. One executive described an AI-informed network redesign that found $50 million in savings on a single product line, not by adding capacity but by rationalizing where capacity sat.

Every capital decision is either an opportunity to improve AI readiness (adding sensors, connecting systems, updating the plant model) or a missed chance that creates the next generation of brownfield constraints (Figure 4).

Figure 4

Application	What AI enables	Evidence
Layout & line design	Digital twin tests throughput and space utilization before metal moves.	**A Global Industrial Technology Company:** +20% productivity, +30% volume flex, +40% space efficiency.
CapEx scenario modelling	AI compares investment options on cost, risk and flexibility before capital is committed.	**A leading Global Healthcare Company:** all CapEx planning runs on an AI-based environment with virtual commissioning.
Debottlenecking	AI pinpoints where capacity is being lost and simulates interventions before teams act.	**North American Climate & Energy Solutions Provider:** from 52→27 plants; $50M saved on one product line.
Brownfield readiness assessment	Sensor and data audit sequences retrofit investment correctly.	**Global Mobility Technology Company:** mandatory data maturity audit before any AI deployment.
Virtual commissioning	New lines validated in digital twin before physical connection.	**A major European automaker:** line ramp-up from 18 months to 3 months.

The five key dimensions

Taken together, the five key dimensions we identified in our interviews define the path toward systemic AI. For each, the functions with high decision frequency and repeatable signals offer the highest-value places to start. Think about where AI can make an impact in minutes, rather than months. Start with one value stream and 2–3 KPIs, build the minimum reusable foundation, prove a closed loop with explicit escalation, then pass one gate before expanding to a second site without rebuilding integrations.

Five dimensions support systemic AI

Dimension 01
Integrate planning, production, quality and logistics so decisions flow end to end

Dimension 02
Build shared data, platforms and guardrails so you don't rebuild for every plant

Dimension 03
Redesign decision rights and operating rhythms so AI is part of daily work

Dimension 04
Connect physical and agentic AI to create the closed loop

Dimension 05
Design for humans in the lead to define accountability as autonomy scales

Dimension 01

Integrate planning, production, quality and logistics so decisions flow end to end

At Dow, one of the world's largest materials science companies, Accenture supported a shift from siloed, function-led operations to a more connected system, where data flows across production, maintenance and operations rather than sitting within them. By digitizing core processes and establishing a more consistent data foundation, the organization enabled more coordinated, real-time decision-making across sites. The result was not just better visibility, but a step change in how decisions are made, moving from local optimization to system-level coordination.[3]

For most manufacturers, integration and data work will not be a clean sequence. Leaders will often have to advance them in parallel. The data foundation does not have to be finished to begin integrating AI systems, but it has to be good enough to begin their learning (depending on the criticality of the system).

More mature systems are now moving to coordinate across time horizons: signals from demand inform supply in near real time; production schedules adjust based on quality drift and maintenance insights; inventory policies respond dynamically to changing constraints. Manufacturers leaving planning, production, quality and logistics as separate systems will continue to fall behind as the system absorbs friction between them.

The same logic applies upstream. When AI informs capital planning—site selection modelling, line configuration, commissioning simulation—manufacturers compress time to value on CapEx investments and design long-term flexibility into the asset base, not just operational efficiency.

Three enablers distinguish organizations that achieve integration from those that stall at dashboards.

First, **data standardization and interoperability are non-negotiable**. Machine data, site-specific taxonomies and fragmented system landscapes must be translated into shared formats and common semantic models. This is often the heaviest lift, as shop floor data is notoriously complex and fragmented. Many organizations will have to adopt their data gradually. As a vice president of industry automation explained, “For scaling up, the problem is the data standardization. That's really the struggle, the big stuff."

Second, **integrate IT and OT**. Enterprise, operational and engineering data must cooperate and converge. This is what moves AI beyond descriptive insight toward predictive and agentic decisions that span functions.

Third, **embed governance into architecture**. The best models combine centralized standards with decentralized execution: central teams define the guardrails, standards and protocol, while plants and functional leaders retain ownership of outcomes. This balance preserves speed while ensuring reliability and reuse.

Dimension 02

Build shared data, platforms and guardrails so you don't rebuild for every plant

One of the most consistent findings in our research is that manufacturers succeeding with AI today did not wait for a perfect data foundation before deploying their first use cases. They built the "foundation alongside the fireworks," to borrow a phrase from a consumer goods executive. That meant starting with a handful of high-value, repeatable deployments to build the proof and the confidence that unlocks the next round of investment.

Across use cases, the goal is for seamless connection, from shop-floor signals to engineering records to process knowledge. As one executive put it: "The data by itself is meaningless. The combination of the data is where it makes a difference—how you connect the data and the metadata together. That is what makes for smart manufacturing."

Making that combination possible means tackling the hardest problem in manufacturing IT: connecting enterprise systems (planning, finance, procurement) with operational systems (machines, sensors, quality) and engineering systems (product specifications, change management, bills of materials). Accenture Research indicates that 55% of plant managers still rely on mostly manual processes for data discovery, while 38% still remain hesitant to apply generative AI in their facilities due to inconsistent quality.[4] The solution depends not on replacing these layers but creating a shared semantic model above them.

When common definitions exist across plant processes, AI can reason between them, and AI is increasingly what makes building and maintaining that semantic layer tractable in the first place. One global automation manufacturer does this through a hardware-agnostic platform layer that communicates with legacy equipment from across the vendor landscape simultaneously. The platform does not replace the underlying systems; it renders them interoperable, enabling agentic AI to traverse the full stack, from shop floor sensors to procurement records in ERP. That is the architectural move that transforms three disconnected system layers into a single, trusted data foundation for systemic AI.

Again: what works in one plant will not necessarily scale to many plants. Shared platforms are crucial.

When AI use cases are built on custom integrations, every new plant requires the same work again from scratch. When they are built on shared platforms with common data standards, success can be redeployed without rebuilding.

Ecosystem strategy changes the buy-versus-build calculus

The question is no longer whether to build or buy. It is what is worth building at all.

Foundation models, industrial AI platforms and maturing ecosystem tools have made building the base layer from scratch indefensible. What manufacturers should retain internally is narrower but more strategic: domain knowledge, process context, proprietary operational data and integration architecture. That is where competitive differentiation lives.

Critically, the need for partners has intensified. Some manufacturers may assume that generative AI and vibe coding had finally made it possible to build complex systems like MES and ERP without specialist help. The experiment largely failed. As one senior digital manufacturing executive put it: "Maybe there are some companies that will try to do it. Good luck. It's much more complex than people anticipate." Architecture, security and integration complexity do not disappear because a developer has a capable coding assistant. The companies that tried are coming back.

No single provider delivers the full stack. The right model acts as a system integrator: modularize the problem and source specialized partners for each module.

Three characteristics define credibility: interoperability across heterogeneous plant environments; trust around proprietary process data; and security standards that hold across every node in the partner network. Domain depth and IT/OT integration track record remain essential—because nothing from any vendor works out of the box without deep integration expertise on top of it.

And finally, across redeployments, it’s strong governance that makes transformation stick. One manufacturer's closed-loop operating model showed what guardrails look like in practice: digital twins serving as the bridge between AI insight and physical action. Agents captured detailed process data, generated insights, tested proposed changes in simulation and pushed approved changes back to machines almost instantly.

The governance decision at the center of this model must remain human-led. Agents can run simulations and propose changes, but humans validate before anything is applied to the physical plant. That human release gate is non-negotiable. The system may generate and test changes, but nothing reaches the production floor without explicit human approval.

"You definitely want to start with the digital twin and the simulation,” a head of automation engineering and digital at a global technology company told us. “When you have a simulation and digital twin first, the data is already there, and once the data is ready, it's really easy to do systematic AI."

Shared data, platforms and governance standards: these are what make systemic AI scalable.

Dimension 03

Redesign decision rights and operating rhythms so AI is part of daily work

Every quarter spent scaling AI on top of an operating model designed without AI in mind will meet with frustration. AI deployment too frequently stalls because the organization didn't design to run with it.

The fix is a fundamental redesign of how decisions get made, who owns them and how performance is managed. When AI coordinates decisions across planning, production, maintenance and supply, ownership and accountability must reflect that reality. When models improve continuously, governance must support monitoring, retraining and ongoing performance management as permanent disciplines.

Which is to say, operating models need to align with how workers experience AI.

Aligning on decision rights means defining three things before AI goes live: what decisions can be automated, what decisions require human validation and what conditions trigger escalation. These are operational specifications that determine whether autonomy scales or stalls.

Making AI part of daily work

If decision rights define the boundaries of AI, then operating rhythms determine whether AI is actually used. This is not a question of whether humans remain in the lead of operations (a point we address more fully in Dimension 5), but whether the structures, rhythms and incentives of the organization are built to make AI a reliable part of how its daily decisions get made.

We found that leading manufacturers are embedding AI into the daily cadence of work: shift handovers reference AI-generated insights, planning reviews are built around live model outputs and exception handling is structured around the escalation paths defined in governance. AI becomes part of how workers work rather than an option or a separate layer.

Consider a Chinese heavy machinery manufacturer's cloud lab. It centralizes all AI demand across business units, enforces a strict two-year ROI gate and prioritizes efforts by feasibility, technical maturity and commercial value.

> **The results compounded quickly: mandatory AI-assisted coding for all developers since early 2025, productivity doubled by mid-year and a 65% reduction in workload by year-end. The sequencing—governance before scale—is what made the difference.**

Manufacturers that make these shifts in governance and adoption unlock compounding value: each deployment strengthens the enterprise, and each refinement improves performance across sites. Those that don't will find themselves rebuilding governance, re-establishing ownership and renegotiating responsibilities with every new pilot.

Dimension 04

Connect physical and agentic AI to create the closed loop

Physical AI excels at execution. It detects, adjusts and moves. Agentic AI, meanwhile, excels at coordination. It sequences, prioritizes and decides across workflows. On their own, each delivers incremental improvement. Together, they create something more significant: a closed loop where the factory doesn't just respond to what's happening, but anticipates, adapts and continuously improves.

In our interviews, executives noted this as the point where everything changes: when agents coordinate the system in real time—adjusting lines, responding to demand changes, managing quality drift and timing maintenance so decisions actually change what happens next, not just what gets reported.

From coding robotics to teaching robotics

Today, most factories still depend on automated guided vehicles (AGVs) built 20 or 30 years ago. They are rules-based, follow strict paths and are unable to adapt. The shift now is to autonomous mobile robots (AMRs) with dynamic navigation and AI-driven fleet coordination, while the near future points to humanoids handling preparation area tasks such as kitting and component staging.

The difference isn't just hardware. It's how these systems learn. Physical AI learns from data instead of following hand-coded rules. It draws on operational data, synthetic simulation data and data from human-operated demonstrations. A global industrial technology company, for example, is training humanoids using decades of workstation simulation capability, with deployments to customer sites happening within weeks.

As one Accenture expert in physical AI put it,

> **“We are moving from coding robotics to teaching robotics. That is the nature of physical AI. Rather than programming a robot with explicit rules, physical AI models are trained on data and can generalize to novel situations."**

Early results are promising. One logistics robotics developer reported a 30% gain from AI-enabled AMR deployment vs static AGV due to the system’s ability to learn dynamically instead of relying on traditional rules-based programming. Meanwhile, executives said that early indications show the potential for AI-enabled human “senses” like vision and sound. At one Chinese manufacturer, cameras “see” workers place components into a box, and an AI model checks for accuracy, while at another manufacturer, AI listens and senses the physical act of connecting wires in real time. In both cases, error rates are greatly reduced while safety increases.

The self-reinforcing loop

What makes the convergence of physical and agentic AI genuinely distinctive is that it compounds over time. Physical AI generates the operational data that improves agentic models. Better agentic models direct better physical AI performance. The loop is self-reinforcing and continuously widens the gap between manufacturers who have built it and those who haven't.

For one manufacturer, an agentic system connected to ERP, manufacturing execution and maintenance systems delivered three simultaneous improvements:

30–40%	downtime reduction
25–30%	maintenance capacity increase
4–6%	improvement in overall equipment effectiveness (OEE)

Integrating physical and agentic AI also creates opportunities for reuse. When agentic logic is built on common platforms and shared standards, what works in one plant can be redeployed across others with greater ease.

Autonomy heightens need for security-by-design

Convergence of physical and agentic technologies, however, also expand vulnerabilities. Manufacturing environments impose a higher reliability standard than office workflows. As agentic logic reaches deeper into OT environments—connecting programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems and real-time control loops—the attack surface expands. A software error on a shop-floor control system can stop production, damage equipment or put people at risk. Systemic AI therefore requires governance that makes speed and safety compatible from the start.

Cybersecurity must be co-designed with the orchestration architecture: separating IT and OT networks, applying strict verification of every connection including internal ones, monitoring for anomalies in real time and maintaining incident-response plans that account for physical safety. AI systems must also carry explicit limits on their authority over production-critical decisions.

Manufacturers deploying AI without these controls in place are creating future IT problems, to say nothing of the serious operational and legal risk.

A director of strategy planning at a large manufacturing conglomerate said,

> **"AI should not be given the full authority to shut down production completely, or to stop the production line completely, or to stop a machine completely where it will cause a mass pile-up within the production floor."**

Dimension 05

Design for humans in the lead to define accountability as autonomy scales

Autonomy without accountability is senseless risk. The manufacturers making the most progress with AI are clear and deliberate about where human judgment remains necessary and what is the roadmap for achieving the broader objectives.

As AI and robotics mature, the nature of human work in manufacturing shifts. Factory managers are unambiguous about this issue: in Accenture's survey of 552 plant managers, 70% identified workforce transformation as the single most critical factor for AI success—above technology, governance or even data. Yet 49% identify training investment as a major hurdle and 46% report that workers fear their roles will become obsolete.[5] Roles will likely change. Not away from the floor, but toward a different kind of contribution: supervising systems, validating outputs and intervening when conditions fall outside defined boundaries. Agents coordinate workflows. Robots execute standardized tasks. People make the calls that matter most. Designing that workforce with intention is what makes the difference between scaling and stalling.

Robust testing and validation remain critical before deployment, especially for live operational environments. Systemic AI should complement and not bypass established safety controls and operate within the boundaries of existing industry regulations.

At the same time, the regulatory landscape for AI is evolving quickly. Manufacturers adopting these capabilities should stay alert to emerging requirements, ensuring systems are not only effective but also compliant and resilient over time.

Trust is the multiplier

Technology adoption in manufacturing has always been a composite problem shaped by behavior and technology as well as physical, temporal and spatial constraints. AI is now part of this dynamic because its systems are adaptive rather than fixed. Workers who do not understand or cannot explain how a model reaches its conclusions are unlikely to trust it when the stakes are high.

Gaps in trust are why adoption slows when systems feel opaque or imposed, while co-designing AI with operators builds long-term confidence and uptake. That risk is visible in the data: only 21% of employees feel they have a voice in how AI is introduced.[6] When people experience AI as something done to them, trust and adoption erodes. Adoption accelerates when the people closest to the work have shaped how the system behaves. Workers who can explain why a particular defect pattern appeared in a specific production run are the same ones who will ensure an AI system performs on the factory floor rather than just in a lab. Frontline expertise surfaces corner use cases that no model training anticipated.

AI also changes how skills are built. Embedding guidance, diagnostics and AI-generated work instructions directly into workflows shifts training from classroom events to continuous, on-the-job learning. And this kind of embedded learning is still rare.

only

19%

of employees in our research say they work in teams that experiment with AI and digital tools to develop, learn and improve together.[7] Closing that gap is part of designing a workforce that can run with systemic AI, not just adopt it.

For brownfield environments with mixed equipment and varying operator experience, this is one of the most practical near-term benefits of AI deployment: compressing the time it takes to certify operators and processes without sacrificing quality or safety.

And make no mistake, the path to systemic AI runs through the workforce, not around it. “The pilot is the human, and the AI is the co-pilot,” as a head of manufacturing excellence at a global consumer good company put it. There can be no organizational deployment without trust and buy-in from the people in charge of it.

What a global automotive supplier got right

A global automotive supplier identified planning as its highest near-term value area across all manufacturing AI applications. Planning runs primarily on data, making it less dependent on physical infrastructure maturity and faster to show results. But executives were explicit: AI systems must operate within clear safety, governance and accountability boundaries regardless of where they are deployed.

Even as AI supported planning, quality inspection and decision workflows, operational teams remained responsible for KPIs once solutions moved into production. Controls that keep humans in the decision path were embedded by design, particularly in environments with low tolerance for error. That approach enabled the company to move from experimentation toward systemic use while maintaining trust on the shop floor.

One automotive manager at this company put it bluntly:

> **"Planning has the most potential to have AI impact on a larger scale and to see actual benefits and fast results."**

Accountability mechanics: who owns outcomes when autonomy is live

Systemic AI only scales when accountability is operational. Four modes of control need to be defined, documented and made auditable before autonomy goes live:

Autonomous execution within guardrails. The system acts when risk is low and defined constraints are satisfied.

Human validation required. Actions are proposed; a supervisor approves before execution.

Escalation required. Predefined thresholds trigger escalation to a named role within a set response time.

Stop and rollback. Safety, quality or security anomalies trigger automatic return to manual control with full incident logging.

Ownership must be equally explicit: plant operations own KPIs, central teams own standards and guardrails and platform teams own performance monitoring and change control.

The performance standards that underpin this accountability are just as important. One manufacturer operates with three defined thresholds: prediction accuracy target of 98–98.5%, a false discovery rate below 1% and a model drift threshold that automatically triggers retraining. When something drifts outside those boundaries, a named owner responds. That is what operational accountability looks like in practice.

Compounding advantages

The manufacturers pulling ahead aren't waiting for better technology. They're building the operating foundations that make the technology they already have compound across every plant in the network. That is the competitive advantage systemic AI creates, and it widens with every deployment.

That's why the five dimensions outlined in this report are so important. They represent the architectural choices that determine whether AI delivers isolated wins or enterprise-wide growth: faster product launches, more resilient supply chains and an operation that gets smarter with every cycle. Each dimension removes a different bottleneck. Together, they redesign how manufacturing works across every stage of a plant's life.

Research Methodology

Accenture Research conducted 36 independent executive interviews with senior manufacturing and technology executives at leading global companies across Europe, North America and Asia-Pacific in March 2026. Participants spanned functions including operations, digital transformation, engineering, supply chain and AI strategy, and represented facilities and enterprise operations across Europe, North America and Asia-Pacific. Interviews were designed to surface real-world evidence of AI deployment at scale—barriers encountered, organizational changes required and performance and value outcomes. We use Generative AI in our research production process. Our research experts review and validate the generative AI outputs with traditional research methods where possible, applying Accenture's Responsible AI standards. Unless otherwise stated, all insights, findings and company examples in this report are derived from these interviews. These qualitative findings are supplemented by Accenture's proprietary quantitative research.

The first report, Making autonomous supply chains real (2025), surveyed 1,000 senior executives across 10 industry sectors to examine the shift toward autonomous supply chains, assessing 29 core supply chain activities for current and projected automation levels and machine delegation across roughly 168,000 data points.

The second report, Rethinking the course to manufacturing's future (2025), combined 15 executive interviews with heads of production at leading automotive, industrial and aerospace companies with a survey of 552 factory managers across Asia, Europe and the US, capturing a dual perspective on manufacturing's trajectory through 2040.

The third report, Talent Reinventors: Delivering value with and for people in the age of AI (2026), combined two global surveys of 1,320 C-suite executives and 4,560 employees across 20 industries and 12 countries with more than 55 interviews, a focus group of 75 workers and analysis of more than 2 million job postings and 3,000 earnings call transcripts to identify the talent practices driving measurable business performance.

References

1. Accenture, Making Autonomous Supply Chains Real, 2025
2. Accenture, Making Autonomous Supply Chains Real, 2025
3. Accenture, Factories of the Future, Dow Case Study
4. Accenture, Rethinking the Course to Manufacturing’s Future, 2025
5. Accenture, Rethinking the Course to Manufacturing’s Future, 2025
6. Accenture, Talent Reinventors, 2026
7. Accenture, Talent Reinventors, 2026

Acknowledgements

The Authors would like to thank, **Ute Wellenberg**, **Edward Blomquist**, **James Velasquez**, **Regina Maruca**, **Jed Bonniwell** and **Andrea Brady**

About Accenture

Accenture is a leading solutions and services company that helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed across the enterprise, bringing together the talent of our approximately 786,000 people, our proprietary assets and platforms, and deep ecosystem relationships. Our strategy is to be the reinvention partner of choice for our clients and to be the most client-focused, AI-enabled, great place to work in the world. Through our Reinvention Services we bring together our capabilities across strategy, consulting, technology, operations, Song and Industry X with our deep industry expertise to create and deliver solutions and services for our clients. Our purpose is to deliver on the promise of technology and human ingenuity, and we measure our success by the 360° value we create for all our stakeholders.

Visit us at accenture.com.

Accenture Research

Accenture Research creates thought leadership about the most pressing business issues organizations face. Combining innovative research techniques, such as data-science-led analysis, with a deep understanding of industry and technology, our team of 300+ researchers in 20 countries publish hundreds of reports, articles and points of view every year. Our thought-provoking research developed with world leading organizations helps our clients embrace change, create value and deliver on the power of technology and human ingenuity.

Visit us at accenture.com/research.