Executive whitepaper series on enterprise physical AI autonomy

When smart systems disagree

The hidden breakdown of decision intelligence in modern manufacturing

Table of contents

Authors

Vijay Narayan

EVP, Global Head for Physical AI and Head of Manufacturing, Logistics, Energy & Utilities

https://www.linkedin.com/in/vijaynarayan1/

Dr. Samih Fadli

CTO for Physical AI and Chief AI Officer of Manufacturing, Logistics, Energy & Utilities

https://www.linkedin.com/in/samihfadli/

Executive summary

As agentic AI accelerates execution, decisions without shared context can institutionalize fragmentation, producing sequences of micro-decisions that degrade quality, stability, margins and service levels. Integration and data platforms improve visibility but do not resolve the problem, because data connectivity does not create shared meaning, governed trade-offs or a persistent record of enterprise decisions.

The consequences are significant. Fragmented decision intelligence increases yield loss, lengthens decision latency, drives inconsistent execution across plants and limits institutional learning. Accountability weakens as outcomes become difficult to trace across systems.

Maturity does not require replacing specialized platforms. What is required is a way to unify decision context across systems so enterprise-level trade-offs can be understood, resolved and remembered. Manufacturers that establish this capability will shorten decision cycles, compound learning and strengthen governance.

The limits of distributed intelligence

Manufacturing leaders have invested heavily in artificial intelligence across planning, production, supply chain, quality, maintenance and customer operations. The result is a dramatic increase in local decision capability. Advanced planning and scheduling platforms optimize production sequences against throughput, cost and due date adherence. Procurement and supplier risk platforms model continuity with precision that was not commercially available a decade ago. Quality management systems classify defects at a resolution that often exceed human inspection. Enterprise asset management platforms anticipate equipment failure hours before it announces itself in yield. Commercial forecasting systems anticipate demand volatility at a cadence that supports real-time replanning. Each system appears intelligent within its own domain and contributes to measurable value when evaluated in isolation.

The enterprise-level question is different.

Has the manufacturing organization itself become more intelligent, or has it assembled a portfolio of disconnected decision engines that now produce competing recommendations at industrial speed?

In many enterprises, the answer is uncomfortable.

Intelligence has been deployed widely, but it has not been unified. The enterprise now has more decisions being generated by more systems, but fewer mechanisms for reconciling those decisions into a coherent operational posture. The enterprise has deployed intelligence everywhere, yet it cannot think as one system.

This is the hidden breakdown of decision intelligence in modern manufacturing. It does not usually appear as a dramatic failure. It appears as repeated schedule changes, unresolved trade-offs, late escalations, inconsistent decisions between plants and recurring debates about which system has the more accurate view of reality. The enterprise has automated parts of the decision process, but the burden of coordination still falls on people. Human teams become the unofficial arbitration layer between systems that do not share a common enterprise context at a moment when the volume and velocity of decisions is outpacing what informal arbitration can reliably absorb.

The manufacturing decision landscape has become more intelligent, but less coherent

The modern manufacturing enterprise is a dense network of decision systems. Enterprise resource planning platforms govern financial commitments and material flows, while manufacturing execution systems coordinate production events. Advanced planning and scheduling platforms allocate capacity, quality management systems enforce specifications and traceability, and product lifecycle management systems carry engineering change history and configuration control. Enterprise asset management systems prioritize maintenance, warehouse and transportation platforms coordinate inventory movement, and commercial forecasting systems anticipate customer demand. Within this landscape, agentic workflows are increasingly layered across platforms, generating recommendations, triggering actions and accelerating execution.

The issue is not that these systems are weak. It is that they are strong within boundaries that do not match the real structure of manufacturing decisions. A manufacturing decision almost never belongs to just one domain.

A material substitution affects supplier risk, cost, line performance, defect probability, warranty exposure, customer commitments and regulatory obligations. Maintenance intervention affects equipment uptime, labor availability, production scheduling, inventory buffers and delivery reliability. A quality alert affects throughput, supplier scorecards, work orders, customer allocation and root-cause analysis. Yet each decision engine typically evaluates the situation from inside its own operational frame, identifiers, event definitions, and sense of priority.

This creates a structural problem. Local intelligence can optimize an objective while degrading the performance of the system. A planning system may protect throughput by pulling production forward, while a procurement system protects cost by favoring a lower-cost supplier. A quality system, in turn, may tighten inspection thresholds, even as a maintenance platform recommends downtime to preserve asset life.
Each recommendation may be valid. Together, they can be incompatible.

The enterprise is not suffering from a lack of intelligence but from the absence of a shared decision architecture.

Scenario one

The material substitution that looked rational everywhere but failed system-wide

Consider a manufacturer facing a disruption in a critical component. The supply chain platform identifies an alternate supplier with acceptable lead time and price. Procurement confirms contract viability. The planning system updates the schedule to preserve customer delivery. The manufacturing execution environment accepts the substitution and routes the change to the line. In terms of continuity, the enterprise appears to have responded effectively.

The problem emerges in the interactions that no single system owns. The alternate material has a slightly different tolerance profile. The quality system begins flagging a defect pattern that is statistically meaningful but operationally ambiguous because its inspection parameters were calibrated to the original supplier's material characteristics. The production system interprets the issue as a line adjustment problem. The supplier management system treats it as a supplier qualification matter. The planning system continues optimizing against throughput because demand pressure remains high. Maintenance, unaware that the rebalance has elevated the displaced line's utilization, schedules preventive work that reduces capacity. Within 72 hours, a secondary defect pattern emerges on the displaced line. The defect is detected. The root cause is not identified because it is the accumulated consequence of four independent system responses to a single upstream event and no system in the enterprise reconstructs the full causal chain.

The failure is that no system has enough cross-domain context to understand the situation as one enterprise event.

What began as a supply continuity decision becomes a quality, production stability and customer service issue. By the time leadership sees the pattern, the organization has already absorbed cost, delay and operational distraction. The enterprise absorbs the impact, attributes it to supplier quality, and rarely learns what actually happened. The most expensive part of the problem is not the original disruption. It is the enterprise's inability to reason across its own decision systems in time.

Scenario two

Maintenance, production and quality each optimize correctly and still collide

A second pattern appears at the intersection of asset health and production pressure. An enterprise asset management platform identifies a high-risk failure on a critical piece of equipment and recommends intervention within a defined window. The production planning system, under pressure to meet a service-level commitment, defers the maintenance window because the line is needed for a high-priority order. The quality system detects a subtle increase in variation but does not yet classify it as a formal defect trend, because its threshold logic was designed for more pronounced signals. Operations leaders are left to reconcile three decision streams, each based on legitimate evidence but lacking a shared interpretation of what the combined picture means.

The coordination challenge is made worse by time. Manufacturing decisions are sequenced. A delay in maintenance changes the risk profile of the next production run. A production reroute changes labor allocation and inspection workload. A quality threshold adjustment changes yield and throughput. These decisions unfold as a chain, but the systems generating them often reason as if each decision were independent. In reality, the manufacturing enterprise is a coupled system, and a decision in one area changes the meaning of signals elsewhere. The shift leader who needs to decide how the next four hours of production should proceed reconstructs the answer manually, by opening four screens, reading four timestamps, and applying judgment the systems cannot apply because they do not share a common operational memory. The reconstruction takes fifteen minutes. It happens three times a day, in every plant, and it is the most practiced skill in the operations function.

This is where agentic AI can amplify both value and risk. When autonomous or semi-autonomous workflows accelerate decisions inside each platform, the enterprise gains speed. But speed without coordination can institutionalize inconsistency. Faster maintenance recommendation, production rescheduling, and faster quality escalation do not automatically create a better enterprise decision. Without shared enterprise context, speed can simply move fragmentation closer to real time.

The consequences are operational, financial, and organizational

Fragmentation does not stay contained within systems. It shows up in operations, slows decisions, creates inconsistency and erodes the enterprise's ability to learn.

Operational drag

Manufacturing organizations often measure improvement at the domain level, yet losses from fragmentation appear across handoffs. Plans are revised repeatedly. Expedites increase. Inventory buffers grow because teams do not fully trust the next decision signal. Quality investigations take longer because root-cause evidence is distributed across systems that do not share a common reasoning history. The enterprise pays for intelligence in each domain, then pays again for humans to reconcile the conflicts between them. In a typical discrete manufacturer, the cumulative effect is yield erosion in the range of one to three percentage points absorbed as routine operational variance and translated into tens of millions of dollars of recoverable margin that is never recovered.

Decision latency

In manufacturing, latency does not only mean slow analytics. It means the time between signal, interpretation, arbitration and action. When multiple AI-enabled systems generate recommendations without a shared decision frame, escalation becomes the default governance mechanism. Teams wait for meetings, approvals and manual interpretation. The organization may have advanced models and still depend on informal human mediation when volatility matters most, which is the moment when the reconciliation tax is most expensive.

Inconsistent execution

Similar events may produce different decisions across plants, regions or business units because each location has a different combination of systems, data maturity, operating norms, and escalation practices. This limits scale. A manufacturer can standardize a process document more easily than it can standardize decision intelligence across fragmented platforms. As AI adoption expands, this inconsistency becomes more difficult to manage because the reasoning behind decisions is increasingly embedded inside systems rather than visible in a single institutional record.

Loss of institutional learning

Each platform learns from its own transactions, outcomes and feedback loops. But the enterprise does not necessarily learn from the full decision episode. The knowledge that a material substitution affected a defect profile, changed machine behavior, altered throughput and triggered customer exposure should become part of the manufacturer's institutional intelligence. In fragmented architectures, that learning is distributed, partial and often temporary. The reasoning behind last quarter's quality release is not available.. The supplier qualification judgment made two years ago is not shared with the engineer facing a structurally similar judgment today. The organization repeats patterns it should have learned from, not because of negligence, but because the enterprise lacks a persistent cross-domain memory of its own decisions.

In an industry where the senior operational workforce is entering a generational retirement, this is not a theoretical concern but a balance sheet event in slow motion.

Why integration alone does not solve the problem

Many manufacturers respond to this challenge by expanding integration. They connect systems through APIs, increase data sharing, centralize information into data platforms and build dashboards that show more of the operating environment. These are necessary foundations. They improve visibility, reduce some manual work, and create access to broader datasets. But they do not by themselves create enterprise decision intelligence.

Integration moves data but it does not guarantee shared meaning. A quality event, a production deviation, a supplier risk flag and a maintenance alert may all be visible in the same environment and still lack a common interpretation. Data can be connected without the enterprise agreeing on what the event means — which trade-offs matter, who has authority and how the decision should be remembered.

Connectivity is not coherence or primarily a data problem. Manufacturers that have consolidated operational data into unified analytical platforms report that the analytical coherence of their data has improved substantially, while the decisional coherence of their enterprise has not. Analytics describe what happened. They do not arbitrate what should happen next. A unified data platform allows the enterprise to look back at its decisions with greater clarity. It does not allow the enterprise to make its next decision as one coordinated organism.

The deeper limitation is architectural. Manufacturing enterprises need a way to coordinate reasoning across systems, not merely exchange information between them. They need decisions to be interpreted in relation to enterprise objectives, operational constraints and prior outcomes. They need a governed approach to resolving conflicts among systems that optimize different objectives. And they need the learning from each decision episode to persist beyond the platform that generated the initial recommendation.

Why agentic AI raises the stakes

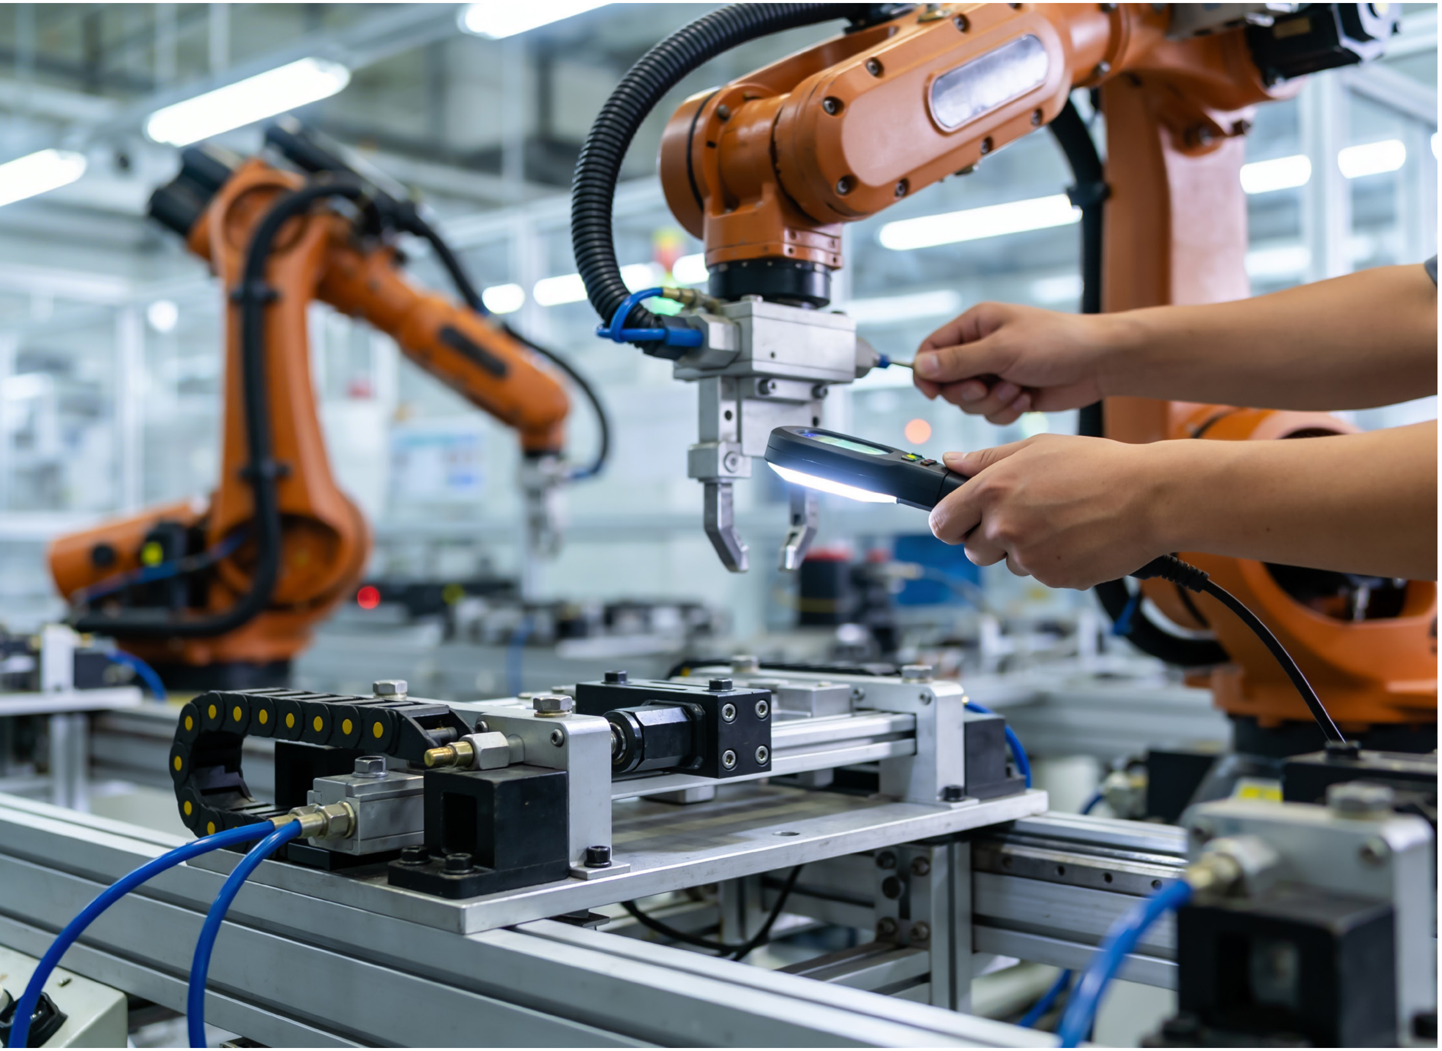

The move from analytics to agentic execution changes the risk profile for manufacturers. Traditional analytics presented insights for human interpretation. Agentic workflows can recommend, sequence and trigger actions across planning, procurement, production, quality and service operations. This creates enormous potential, because manufacturing decisions can be accelerated and executed closer to the moment of need. It also increases the cost of fragmentation. If each agentic workflow operates within a narrow domain boundary, the enterprise may automate the very contradictions that human teams previously slowed down long enough to examine.

This matters because manufacturing is a trade-off business. Improving yield may reduce throughput, protecting service levels may increase quality exposure, and lowering input costs may create process variation. Human leaders have historically managed these trade-offs through escalation, experience and operating cadence. Agentic AI challenges that model because decisions can move faster than the organization can convene, interpret and reconcile. The governance problem is therefore not only whether an AI action is permissible. It is whether the action is coherent in relation to the broader enterprise state.

The risk is not a single rogue decision. The more likely risk is a series of locally rational micro-decisions that collectively erode performance. A procurement agent protects continuity through substitutions. A scheduling agent rebalances the line to sustain output. A quality agent tightens inspection to ensure conformance. A logistics agent expedites shipments to meet customer commitments. Each decision is defensible. Together, they can erode margins, destabilize schedules, drive avoidable rework and pull management into constant decision reconciliation instead of strategy.

The executive failure mode: Fragmented accountability

Fragmented intelligence also creates an accountability problem for executives. When outcomes deteriorate, it becomes difficult to determine whether the issue came from the data, the model, the process, the operating constraint or the interaction among several systems. Leaders may hear different explanations from different functions. Supply chain points to supplier variability. Operations point to schedule pressure, quality to process drift, maintenance to asset condition, and finance to cost discipline. Each function may be describing a real part of the problem, yet none owns the complete decision episode.

This ambiguity weakens management discipline. Performance reviews become debates about whose system was right rather than what the enterprise learned. Continuous improvement efforts focus on localized remediation because the full causal chain is hard to reconstruct. Technology teams are asked to integrate more data, while business teams continue to operate without a shared decision memory. The organization may improve individual dashboards and still fail to build a common understanding of how decisions actually propagate across the manufacturing network.

This ambiguity carries into the regulatory and audit functions. Every regulated manufacturer, from aerospace and medical devices to automotive and food processing, operates under frameworks that require the enterprise to reconstruct the reasoning behind operational decisions. In a fragmented architecture, the decision trail exists in fragments across the systems that participated in the decision without unified record of how those fragments were reconciled. Compliance teams describe the resulting reconstruction work as the single largest source of unbudgeted effort in the regulatory function. Regulators are beginning to ask not whether individual systems are controlled, but whether the enterprise can demonstrate coherent governance across them. Few manufacturers can answer that question today.

For senior leaders, the implication is clear — agentic AI cannot be governed as a collection of isolated productivity tools. It must be treated as an emerging decision fabric across the manufacturing enterprise. That fabric needs shared context, consistent policy interpretation and a persistent record of decisions and outcomes. Without those capabilities, the enterprise may increase automation while weakening accountability.

What maturity begins to look like

A mature manufacturing enterprise does not expect every system to become a universal decision engine. Specialized platforms will remain essential. Planning systems should focus on planning. Quality systems should focus on quality. Maintenance systems should focus on asset reliability. Maturity means ensuring these systems do not reason in isolation when their decisions interact. The enterprise must preserve specialization while creating coherence across domains.

In practical terms, this means that critical decisions should be interpreted against a shared enterprise context. A material substitution should not only be evaluated by price and lead time, but by its relationship to quality history, process capability, equipment behavior, customer commitments and prior outcomes. A maintenance deferral should not only be evaluated by immediate production demand, but by its effect on defect probability, reliability risk, schedule stability and service exposure. The enterprise should be able to understand the decision episode as one event, even when multiple systems contribute to it and retain that understanding long enough for the next similar event to benefit from it.

This is the direction of travel for manufacturing AI. The next advantage will not come from deploying one more isolated intelligent tool. It will come from turning distributed intelligence into enterprise intelligence. Manufacturers that establish this capability will create a compounding advantage: every decision will make the next decision better, not only within a system, but across the operating enterprise.

The way forward

Looking ahead, manufacturing leaders will need to move beyond the assumption that more AI systems will automatically produce a more intelligent enterprise. The next stage of maturity will depend on the ability to unify decision context across the manufacturing value chain. This does not require replacing existing platforms. It requires a new architectural posture in which existing systems can contribute to a shared enterprise understanding of decisions, trade-offs, and outcomes.

The organizations that make this shift will be able to reason across planning, procurement, production, quality, maintenance, logistics and customer commitments with greater consistency. They will shorten decision cycles because fewer conflicts will require manual arbitration. They will improve resilience because operational lessons will compound across plants and events. They will also strengthen governance because decisions will be traceable through a shared interpretation of why actions were recommended, accepted, modified or rejected.

The defining challenge is no longer whether manufacturing enterprises can deploy intelligent systems. They already can. The challenge is whether they can make those systems contribute to a coherent enterprise intelligence capability. Manufacturers that solve this problem will not simply automate more decisions. They will develop the capacity to think, decide and learn as an integrated operating system. Those that do not will continue to optimize locally while remaining constrained by fragmentation at the enterprise level.

cognizant

Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

World Headquarters

300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Tel: +1 201 801 0233

European Headquarters

280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office

Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters

1 Fusionopolis Link, Level 5
NEXUS@One-North, North Tower
Singapore 138542
Tel: +65 6812 4000

WF4823550