Closing the Gap Between Event and Action

Real-Time AI Orchestration for the Modern Manufacturing Plant

Executive Summary

To realize the vision of Industry 4.0, manufacturing leaders have spent the better part of a decade investing heavily — upgrading ERP, MES, MRP, and operational systems, and building data lakes, digital twins, and a growing list of AI pilots. Yet the value creation so often promised has remained out of reach. With the emergence of low-code and no-code platforms and flexible integration APIs across systems, the instrumentation for these solutions is finally in place. But real challenges remain. Capturing data and optimizing processes directly on the production floor is still difficult. And closing the gap between an event and a coordinated response — then measuring its impact on the business — is still counted in hours, not seconds.

So, what's missing? Orchestration across both static and real-time systems: an event-driven, edge-capable fabric that connects existing systems, AI models, equipment, sensors, and people into a single live execution context.

That capability has a name: Operational AI. It is the AI that runs within the plant's live operational loop, reacting to events as they happen and driving coordinated action across systems, machines, and people in real time. It is not the same thing as the AI capabilities most factories have already purchased. The intelligence embedded inside ERP, MES, or PLM is analytical and assistive: it reasons over the data inside its own platform and surfaces answers and recommendations for a person to act on. Operational AI is executional and cross-system: it senses across every platform at once, decides based on the combined signal, and routes the action back to the system that owns it — within the window in which that action still matters.

What current factory applications lack: cross-system orchestration from event to action. Operational AI closes that gap with coordinated, real-time execution.

VANTIQ

Operational AI complements — rather than replaces — the AI embedded in ERP, MES, PLM, and other enterprise systems.

What this paper highlights:

- Operational AI is not the AI already inside your systems. The intelligence embedded in ERP, MES, and PLM (Product Lifecycle Management) analyzes and advises within one platform; whereas Operational AI senses, decides, and acts across all of these systems in real time.
- Data-at-rest architectures cannot drive real-time action. By the time data lands and a model reads it, the production line has moved on.
- Closed systems AI cannot orchestrate across ERP, MES, PLM, SCM (Supply Chain Management), CX, and robotics. Each platform's intelligence stops at its own data boundary.
- An event-driven, edge-first orchestration layer is what converts existing investments — rather than replaces them — into autonomous operations.
- Orchestration is broader than AI orchestration. The layer coordinates systems, machines, and people as well as the GenAI and agents embedded among them. AI is a participant in the loop, not the loop itself.

1. The Paradox of the AI-Ready Factory

Manufacturers across the globe are under sustained pressure to lift throughput, cut unplanned downtime, improve quality, and squeeze out the overhead that quietly erodes margin. Most CIOs reading this have already approved the spending that should, on paper, close that gap. Their MES (Manufacturing Execution

Systems) are modern. The data lake is populated. A digital twin exists. AI chatbots have been layered on top. And yet the same pattern repeats across every operational discipline. Let's explain with concrete examples.

- Production line orchestration depends on shift supervisors verbally coordinating MES, APS, and robotics decisions that should be moving in lockstep. When the schedule changes in APS, the others don't update at the same time: the MES dispatch list, the robot work queue, and the line-side material call each refresh on their own schedule, so they quickly fall out of sync. This costs throughput, idle robot time, and downstream rework.

- Predictive maintenance models flag a bearing trending toward failure, but the work order, the spare-part lookup, the line rescheduling, and the operator notification still pass through separate systems and separate people.

- Quality automation catches the defect, but the corrective action loop back to the upstream station is manual and lagging. By the time a quality engineer reviews the inspection results, traces the cause back to the offending station, and adjusts the process parameter, dozens of units have already been built with the same defect — sorted, reworked, or scrapped after the fact because the upstream station never got a real-time signal to correct itself.

- Worker safety alerts trigger on a single camera feed with no situational context drawn from MES state, equipment status, or the presence of other workers nearby.

The cost associated with not addressing these issues is significant. For example, when the manufacturing line is idle, the average manufacturing operation loses roughly $260,000 per hour, rising to over $2.3M/hour in the automotive industry.

VANTIQ

Operative Gap

Most factories cannot respond to events in real time. Systems such as ERP and MES look backward; they record what happened. The factory floor needs a platform that looks forward, reacts in milliseconds, and orchestrates action across systems, equipment, sensors, and AI agents as events unfold.

In every one of these cases, the signal arrives; the coordinated response does not. The contradiction is architectural, not a tooling problem. Data exists. Context does not flow. Each vendor's AI is optimized first and most deeply for the data within its own platform: ERP-native AI reasons best within ERP, MES-native AI over MES. Cross-vendor reach exists through connectors, but it is shallow, batch-paced, and built one integration at a time. The horizontal stack (PLM, ERP, CX, SCM) and the vertical stack (MES, APS, line, robotics) still operate as silos in any practical real-time sense.

AI bolted onto one of these systems becomes an expert with tunnel vision: it knows everything inside its own platform and very little, in real time, about what is happening in the other systems. The MES-aware AI system sees the ERP work order only after a batch sync. The ERP-aware AI learns about the line slowdown only when a report rolls up. No single horizontal or vertical system has a live, cross-system picture, so no single system can drive a coordinated real-time response.

2. Why Today's System Architectures Fall Short

Before introducing a new layer, it is worth being precise about why the architectures already in production — or already in proposals on a COO's and CTO's desk — fall short of real-time operational requirements. Five patterns dominate the landscape. Each is solving a real problem. None of them is solving this one.

Pattern 1 — Data lake plus chatbot

Backward-looking by construction. By the time the line data lands in the lake and a model reads it, the line has already moved on. Cloud round-trips add latency that the production line cannot absorb, and streaming every sensor reading upstream is operationally prohibitive. The architecture optimizes for analyzing what happened, not for acting on what is happening.

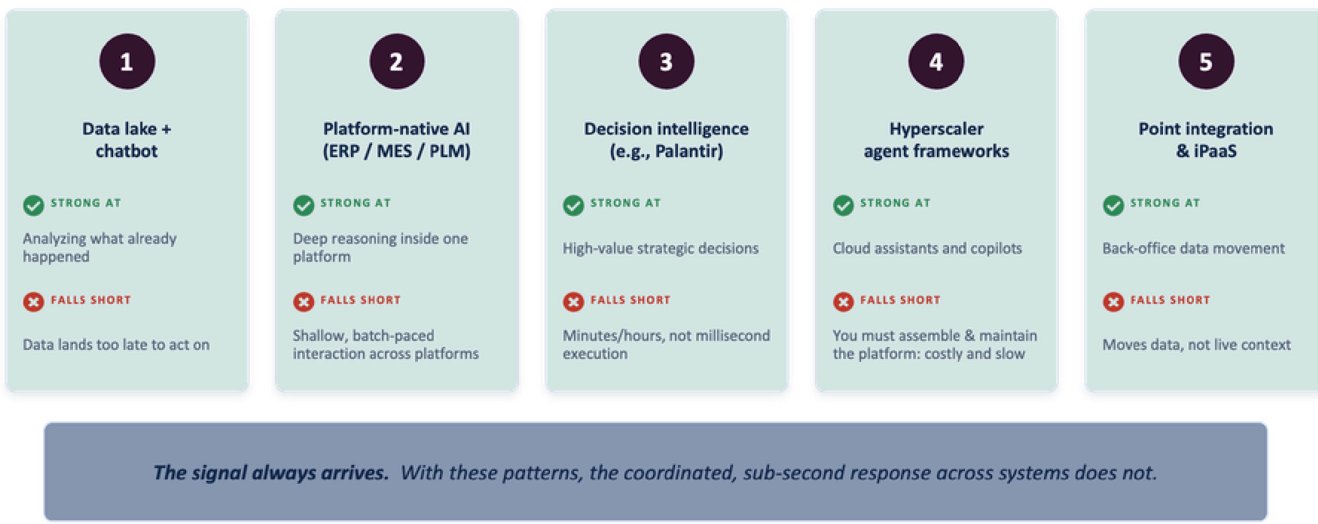

Figure 1: Five patterns dominate manufacturing IT. Each solves a real problem; none drives real-time, cross-system action

Pattern 2 — Vendor-native AI (ERP, MES, or PLM with "AI inside")

Optimized first and deepest for the data inside a single platform. Cross-vendor reach exists through connectors, but it is shallow, batch-paced, and built one integration at a time — adequate for periodic synchronization, not for real-time orchestration across PLM, ERP, MES, SCM, and robotics. The result is closed intelligence rather than orchestrated intelligence: a multiplying set of AI assistants, each fluent in its own data and only loosely aware of the others, none of which positioned to drive a coordinated response as events unfold.

VANTIQ

Platforms such as Palantir are built to inform high-value decisions, not to execute the high-frequency, real-time decisions that manufacturing requires.

Pattern 3 — Enterprise decision-intelligence platforms (e.g., Palantir Foundry / AIP)

These platforms are genuinely strong at retrospective analysis and AI-assisted strategic decisions, such as rerouting a global supply chain, redeploying a fleet, and adjusting an enterprise risk position.However, its architectural limits matter for real-time manufacturing. They are built to inform high-value decisions, not to execute high-frequency ones. For example, Palantir Foundry is data-driven: events flow through the Ontology — ingest, Ontology update, AIP condition, write-back — and latency accumulates at every hop. The platform is engineered for complex AI-assisted decisions with a human in the loop, measured in minutes or hours. Synchronizing thousands of sensors, coordinating hundreds of machines mid-production, or routing AMR fleets in real time — grid-protection, line-control, and robotics-coordination decisions on a millisecond scale — is not what the platform was engineered to do.

Pattern 4 — Build-it-yourself on a Hyperscaler agent framework

Hyperscaler products such as AWS Bedrock AgentCore, Google ADK, and Azure AI Foundry Agent Service are strong for cloud-resident assistants, copilots, and task automation. But as the operational execution layer for a plant, they carry key architectural gaps. The customer must first assemble the platform before building the application. This involves stitching agent runtimes, serverless functions, managed databases, identity, observability, and message buses into a working whole — and much of that wiring has to be redone for the next use case, because each one likely has its own data sources, latency targets, and security model.

The frameworks are cloud-native in architecture, so plants with strict network segmentation, intermittent WAN connectivity, or air-gapped requirements are at odds with the operating model. OT protocol support (OPC UA, MQTT, Modbus, DNP3, SCADA, PLC) is custom code that the customer must write and maintain. The underlying services also change according to the hyperscaler's release schedule — not the customer's — and not every service is available in every region, leaving the customer's team to absorb both version changes and geographic gaps.

Hyperscaler agent platforms sell you the blocks. You're stuck building the solution, wiring every connection, writing the OT drivers, and maintaining it all — forever, on their release schedule, not yours.

Pattern 5 — Point integration and iPaaS

These solutions move data, not context. Batch sync and request/response APIs operate on minute and hour timescales, scoped to specific vendor pairs. They are useful for back-office integrations; structurally unable to drive sub-second coordinated action on the production line.

2.1 What the factory floor needs for operational AI

Measured against what operational AI requires, none of the five patterns fits — though not all for the same reason. Patterns 1, 2, and 5 are data-at-rest problems that pretend to be real-time. Pattern 3 is a decision-intelligence platform pressed into an execution role for which it was not architected. Pattern 4 is a toolkit the customer must assemble into a platform — and re-assemble every time the cloud provider ships a change. None of them is wrong; each is solving a different problem.

What manufacturing needs is a new pattern that can work seamlessly with the existing systems. This new pattern/layer that sits between its events and its action needs to meet these seven requirements:

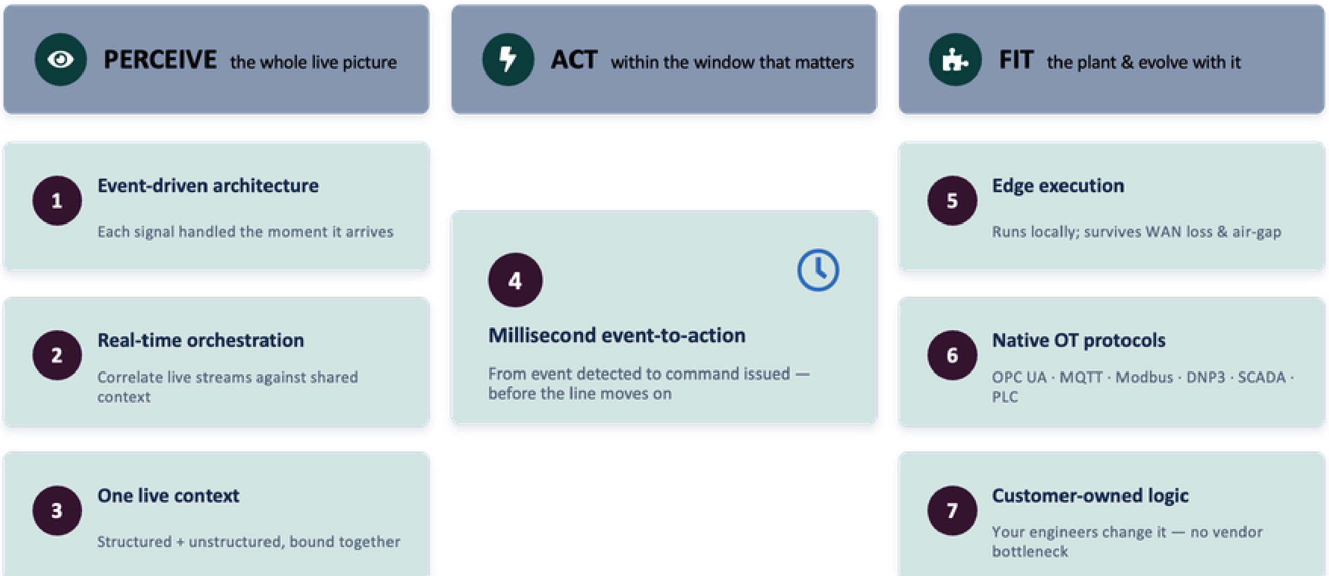

Figure 2: Three demands on the orchestration layer resolve into seven concrete requirements

It must perceive the whole live picture, binding signals from every system into a single real-time context rather than reasoning within a single platform. It must act within the window that matters, carrying a decision from event to command fast enough and close enough to the line that the response still changes the outcome. And it must fit the plant and evolve with it, speaking the floor's native protocols and staying in the hands of the customer's own engineers. Those three demands resolve into seven concrete requirements.

- **Asynchronous, event-driven architecture.** Plant signals do not arrive on schedule. A sensor fault, a quality flag, a safety alert, and a schedule change all surface independently and in parallel. A request/response or batch architecture forces those signals through a queue and a polling cycle; an event-driven architecture lets each one trigger its own handler the moment it arrives. Without this, the platform processes one event at a time and falls behind the line.

- **Real-time orchestration across events.** Individual events rarely tell the whole story. A vibration spike matters when paired with a maintenance record and a production-schedule change; a fall alert matters when paired with MES state and zone occupancy. The platform must correlate multiple live event streams against a shared context and determine the response to the combined signal in milliseconds. Without this, every event is addressed in isolation, and the coordinated response never forms.

- **Millisecond reaction times across the full event-to-action path.** Fast ingest is not the same as fast action. The clock that matters runs from event detected to command issued back into the MES, the robot, or the operator. Anywhere on that path that adds a batch window, a sync interval, or a human hand-off, the line has already moved on by the time the response lands.

- **Correlation across structured and unstructured sources in one live context.** Production reality spans MES, equipment telemetry, vision, CCTV, AMR fleets, and operator records — and no single system holds the full picture, by design. Each is the source of record for its own slice; the orchestration layer pulls those slices together in real time, decides on the combined signal, and routes the action back into whichever system owns the response — a work order into the MES, a setpoint change into the line, a dispatch into the Autonomous Mobile Robot (AMR) fleet. Without this, every system keeps producing locally correct decisions on locally incomplete information.

Operational AI must correlate live signals across plant systems, determine the right response in milliseconds, and continue executing locally even when cloud connectivity is unavailable.

VANTIQ

Vantiq is the real-time execution and orchestration layer for Operational AI, connecting systems, sensors, robots, AI agents, and people across the entire factory floor in one operational fabric.

- **Edge execution that survives WAN loss and respects plant network segmentation.** The line cannot stop when the cloud connection does. Critical safety and control decisions must continue to run locally on the plant network, even when the WAN is intermittent or the facility is air-gapped by design. A cloud-only architecture is not a deployment preference; it is incompatible with how regulated and high-availability plants are run.

- **Native OT protocol support as a platform capability.** OPC UA, MQTT, Modbus, DNP3, SCADA, and PLC are the ways the plant talks. If the platform does not speak their language natively, every sensor type, every line, and every firmware update becomes a custom integration that the customer writes, tests, and retests. The integration cost compounds with every site and every retrofit.

- **Customer ownership of application logic**. A manufacturing plant evolves constantly — new sensors, new lines, new rules, new sites. If every meaningful change routes through the vendor's engineering team on the vendor's timeline, the platform becomes the bottleneck the digital transformation was supposed to remove. The customer's own engineers must be able to modify the system without requesting permission.

3. Vantiq: A Real-Time Orchestration Layer for Operational AI

Vantiq is the real-time execution and orchestration layer for Operational AI. It does not replace the MES, ERP, PLM, SCM, or robotics stack. It sits across them as a vendor-agnostic real-time orchestration fabric, processing events as they occur, threading distributed signals into a live business context, and orchestrating coordinated action across AI agents, equipment, and human operators.

In short, Vantiq orchestrates the full floor — legacy and modern systems, equipment, robots and AMRs, and the people who run them — with GenAI and agents deeply embedded as participants, not as a separate AI layer.

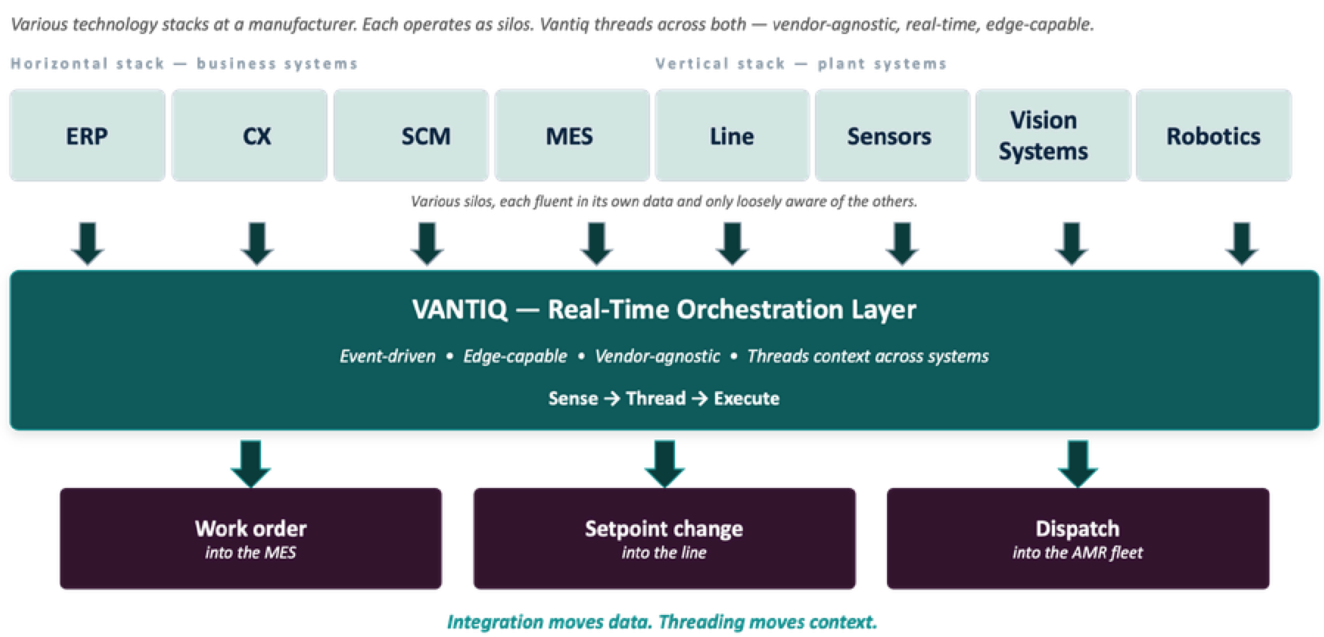

Figure 3: The Silo Problem — and the Threading Layer That Closes It

3.1 Four architectural pillars

Event-Driven Architecture (EDA). Events are processed as they happen rather than waiting for data to land in a lake. The execution model is continuous and asynchronous; every event is handled independently and in parallel across all concurrent sources. There is no batch window between the line and the response.

Edge-first design. Full application logic — event handling, state, decision rules, multi-system orchestration, closed-loop action — runs at the edge node, autonomously, without requiring cloud connectivity.

Operational AI is the real-time orchestration layer that coordinates AI, enterprise systems, robotics, and people — at the edge, and automates actions on events without replacing existing investments.

Only what matters is sent upstream. This is the architectural answer for plants with strict network segmentation, intermittent WAN connectivity, or air-gapped requirements.

True multi-agent (A2A) orchestration. Teams of specialized AI agents — an inspector who watches the line, a diagnostic agent who investigates root causes, a kaizen agent who recommends improvements, an orchestration agent who coordinates the response — work alongside equipment, robotics, and human operators. This is open agent-to-agent coordination across systems, not a single LLM with tool calls trapped inside one vendor's ecosystem.

These software agents coordinate directly with Physical AI on the floor — robots, cobots, AMRs, and embodied vision systems — so perception and action converge into a single loop. Individual Physical AI is only as useful as the coordination around it; Vantiq is the real-time layer that turns isolated robots and vision models into a coordinated whole working alongside systems and people.

Designed to extend, not replace. Vantiq is built to sit atop the MES, ERP, PLM, SCM, CX, robotics, and AMR fleets already in production. No rip-and-replace, no parallel data platform, no Ontology lock-in. The investment a manufacturer has already made becomes the substrate the orchestration layer threads together.

3.2 The Sense → Thread → Execute loop

In operational terms, every Vantiq-orchestrated response follows the same loop (Figure 2). The vocabulary varies by the use case, but the architecture does not.

Sense. Events arrive in real time from MES, equipment telemetry, vision and CCTV, IoT sensors, AMR fleets, ERP, and CX. Filtering and prioritization happen in line; each event is enriched with the initial context needed for downstream reasoning.

Thread. Distributed events are bound together into a single live business context. The platform then maps a step-by-step response — using deterministic rules and state machines when the path is well-defined and invoking LLM-based reasoning when the situation requires interpreting unstructured context, weighing trade-offs, or planning amid ambiguity. When LLM reasoning is in the loop, parallel multi-hop retrieval pulls related items, dependencies, and tacit knowledge from the knowledge graph, RAG, core systems, and technical documentation through a single unified interface, and conclusions are grounded in evidence to mitigate hallucination and long-context information loss.

Execute. AI agents decide and issue commands back into the systems that own the action — MES, robotics, AMRs, work-order systems, or the human on the floor. Impact scope and recommended countermeasures are surfaced; insights are proactively pushed to stakeholders and downstream systems.

Integration moves data. Threading moves context.

Traditional integration ships data between systems within minutes or hours and is scoped to specific vendor pairs.

Vantiq's threading layer ships live context — the state, the history, the related events, the operator's last action — across systems in milliseconds, vendor-agnostic by design.

Sense every signal. Analyze with full context. Execute with confidence — from raw event to grounded action, without the guesswork in between.

VANTIQ

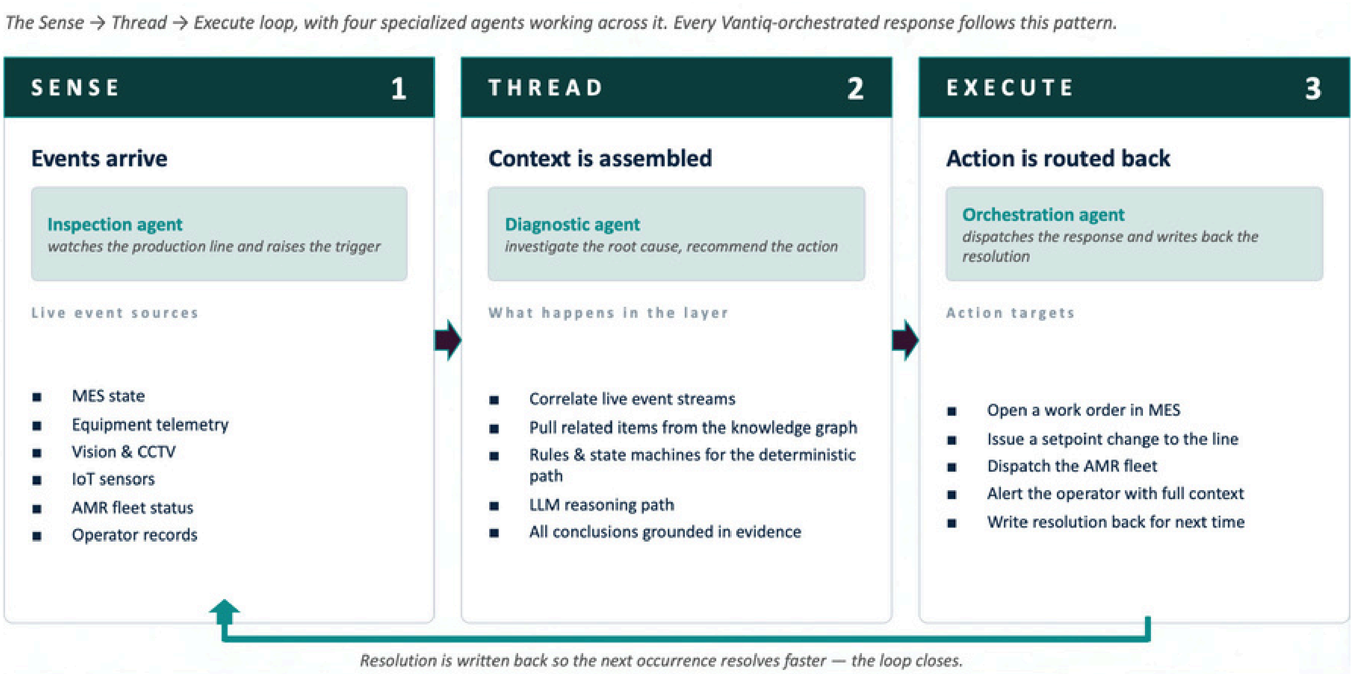

Figure 4: How an incident or anomaly gets resolved in real-time

4. Real-Time Monitoring Use Cases: From Sensor to Action

The architectural argument lands when it touches operations. Each of the following use cases follows the same loop — Sense, Thread, and Execute — and each replaces a workflow that today depends on dashboards, manual investigation, or shift-handoff knowledge transfer. The set is illustrative rather than exhaustive.

Real-time anomaly detection for quality control

Trigger: Sensor and camera signals on the line indicating temperature deviation, abnormal sound, or vibration outside tolerance.

Orchestration: AI inspects, diagnoses, and recommends; an AI agent pushes corrective action to the line, opens a work order, or alerts the operator with full context.

Outcome: Earlier defect detection, lower scrap and rework, fewer unplanned line stoppages, and a continuously improving quality model.

Predictive maintenance, CBM, and AMR coordination

Trigger: Equipment telemetry trending toward failure; Autonomous Mobile Robot (AMR) fleet state and production schedule diverging from plan.

Orchestration: Failure is anticipated before it occurs; Vantiq dispatches maintenance, reroutes production, opens work orders, and coordinates AMRs, robots, and humans against a shared real-time context.

Outcome: Higher MTBF, lower unplanned downtime, deferred capex, and a concrete step on the path toward dark-factory operations.

Cold-chain and process-condition monitoring

Trigger: Temperature, humidity, or pressure deviation in regulated storage or in-transit refrigerated logistics.

Orchestration: IoT stream is checked against thresholds and context; an alert is pushed to the responsible operator and automatically tagged with vehicle, zone, and storage-location metadata; the workflow is logged for audit.

Outcome: Reduced loss, faster anomaly response, stronger audit and compliance posture, and a defensible trail for regulators and customers.

Real-time inventory, demand forecasting, and auto-replenishment

Trigger: POS, e-commerce, logistics, and MES stock signals integrated as a continuous stream.

Orchestration: AI demand forecast is threaded against live inventory; replenishment orders are issued automatically based on warehouse and line-side stock; master data is kept current across systems.

Outcome: Fewer stockouts and less obsolete inventory — directly relevant to the kind of multi-tens-of-millions-of-dollars obsolete-inventory exposure typical at LED, semiconductor, and consumer-electronics manufacturers with very short product lifecycles.

Plant safety monitoring (fall and abnormal behavior detection)

Trigger: Vision and pose-estimation AI running on CCTV feeds from the line, the warehouse, and high-risk zones.

Orchestration: Edge inference flags unsafe posture, falls, or abnormal behavior; the responsible manager is alerted with camera and location context; the incident is logged for safety review.

Outcome: Lower incident rate, faster on-the-floor response, and measurable safety culture improvement that satisfies both insurers and regulators.

Connective tissue: real-time data plus generative AI on the floor

Across every use case above, Vantiq connects live operational signals to generative AI grounded in plant knowledge — explaining in natural language what is happening, what to do next, and why. This is what shifts the floor from intuition-based to data-driven and gives a frontline operator the same context a senior engineer would have.

Four AI agents — Inspector, Diagnostic, Kaizen, Orchestration — now close the loop that a modern MES and a well-stocked data lake never could.

5. Proof Point: LS Thira-UTech

LS Thira-Utech provides a smart factory and smart logistics solution — Intelligent Factory Control (IFC) — for semiconductor and secondary battery manufacturers.

Their customers had taken the standard playbook as far as it could go: a modern MES, a populated data lake, and AI chatbots reading from the lake. The chatbots could describe what had happened on the line; they could not act on what was happening now. The two-hour root-cause window stayed in place because every incident or anomaly still required a human to read the chatbot's summary, walk the line, check the systems, and decide what to do.

What they did

LS Thira-UTech changed their architecture rather than the tools. They embedded Vantiq at the core of their IFC platform as the real-time orchestration layer, sitting across MES, equipment telemetry, log files, CCTV and vision systems, operator records, and the AMR fleet. Events from all six sources flow through Vantiq in real time and are correlated against shared production context.

Four specialized AI agents now run on that fabric, each with a defined job in the response loop. The Inspector agent watches the line for anomalies and raises the trigger. The Diagnostic agent investigates the root cause by integrating correlated event streams — telemetry, MES state, recent operator actions, and historical incident records — and returns a grounded explanation. The Kaizen agent recommends a corrective action, drawing on prior resolutions for similar anomalies.

The Orchestration agent coordinates the response by dispatching maintenance, opening a work order in the MES, alerting the operator, and writing the resolution back to ensure the next occurrence is resolved faster. Together, the four agents close the loop that the chatbot-on-a-lake architecture could only narrate.

What it delivered

- **Six weeks from concept to working prototype,** against multi-quarter cycles that traditional MES development and data-lake projects require.
- **Tens of thousands of events processed in real time at the edge,** with MES, equipment, CCTV, vision, AMR, and operations data integrated into a single event-driven fabric.
- **Targeted 2× reduction in anomaly recognition and analysis time,** measured against the one-to-two-hour manual root-cause baseline.
- **A platform foundation that extends into predictive maintenance, defect detection, and autonomous robotics** without re-architecting each time a new capability is added.

6. Vantiq and the Shift to Industry 5.0

Industry 4.0 optimized the factory for automation. Industry 5.0 — increasingly referred to as Manufacturing 5.0 — reframes the goal around three pillars: human-centricity, resilience, and sustainability. The vision is not the "lights-out" plant run entirely by machines, but the collaborative factory, where human judgment and machine precision operate in continuous synergy.

"Only Vantiq could orchestrate this level of manufacturing integration."

— Youngmin Lee, Director of Manufacturing Intelligence, LS Thira-UTech

Cobots take on physically punishing and repetitive tasks; people focus on adaptation, problem-solving, and judgment that machines cannot replicate. This collaborative model runs on Physical AI — the robots, cobots, AMRs, and perception systems that sense and act in the physical world — and on the orchestration layer that keeps them coordinated with each other, with plant systems, and with people in real time. The hardest part of this transition is rarely the hardware. It is connecting humans, machines, and AI into a single coordinated loop and capturing the tribal knowledge that walks out the door as a generation of experienced operators retires.

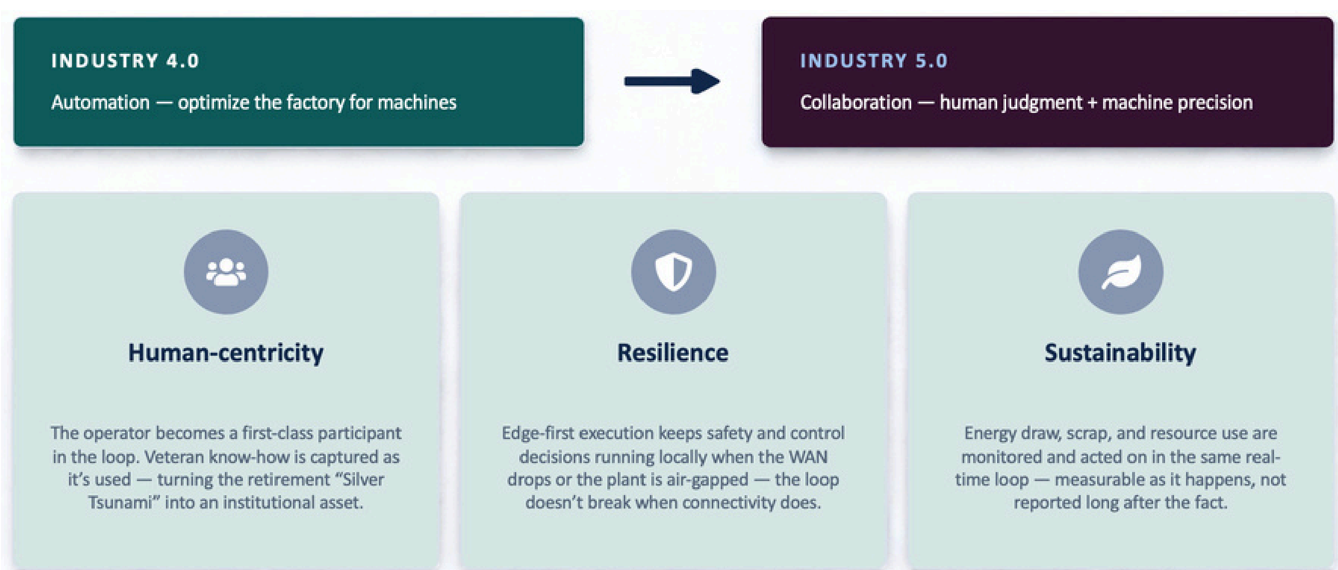

Figure 5: Industry 5.0 reframes the goal around three pillars — each an orchestration problem Vantiq is built to solve

This is an orchestration problem, and it is the problem Vantiq is built to solve. The same real-time layer that threads MES, equipment telemetry, vision, and AMR signals together also treats the human operator as a first-class participant in the loop, not a downstream recipient of a report.

When an agent diagnoses an anomaly, Vantiq pushes the explanation, the recommended action, and the reasoning to the operator in natural language, giving a frontline worker the same context a senior engineer would have. And because every resolution is written back into the shared context, the expertise of veteran operators is captured as it is used and made available to the next person on the line, turning the "Silver Tsunami" of retirements from a knowledge-loss risk into an institutional asset.

Edge-first execution directly delivers the resilience pillar: critical safety and control decisions continue running locally when the WAN drops or the plant is air-gapped, so the collaborative loop does not break when connectivity does. And the same real-time visibility that lifts throughput and quality makes sustainability measurable — energy draw, scrap, and resource use can be monitored and acted on in the same loop as everything else, rather than reported after the fact. Industry 5.0 describes where manufacturing is going. Real-time orchestration is how a plant gets there without ripping out the systems — or sidelining the people.

7. From Automation to Autonomy: A Practical Path Forward

7.1 Three decisions that flow from the argument

- **Treat real-time orchestration as a layer, not as a feature of any one application.** It is the connective tissue between systems the manufacturer already owns.

Industry 5.0 describes where manufacturing is going. Real-time orchestration is how a plant gets there — without ripping out the systems, or sidelining the people.

The only architecture decision that matters: build the orchestration layer once, and never re-platform again.

- **Choose an execution layer that is vendor-agnostic across ERP, MES, PLM, SCM, CX, and robotics,** so that AI-model changes and system upgrades do not trigger re-platforming. Closed intelligence will keep multiplying; only an open orchestration layer keeps the IT landscape open to the future.

- **Prove value in weeks on a single high-pain use case,** then extend the same fabric across the plant and across sites. Anomaly response, condition-based maintenance, and AMR coordination are all natural starting points.

7.2 A six-week-to-value adoption pattern

- **Weeks 1–2.** Select one high-pain trigger — for example, anomaly recognition on a critical line — and the systems it touches. Define the manual baseline (mean time to recognize, mean time to resolve, scrap impact).

- **Weeks 3–4.** Thread events from those systems into Vantiq. Stand up the first specialized agents — typically an Inspector and a Diagnostic — to replace the first manual investigation step.

- **Weeks 5–6.** Close the loop. Add Orchestration and Kaizen agents that write back into the MES, dispatch workflows, and alert humans with context. Measure against the manual baseline.

- **Beyond week 6.** Extend the same fabric to additional lines, additional sites, and adjacent use cases — predictive maintenance, defect detection, AMR coordination, multi-site knowledge sharing. Each new use case lands on the same orchestration layer rather than restarting from infrastructure.

Conclusion

Data-centric architectures look backward. Event-driven orchestration looks forward, and it acts. The mandate facing manufacturing IT leaders is no longer "add AI to the factory." It is to give AI, software and physical AI alike, a real-time matrix across the factory — a matrix that respects the systems already in place, runs at the edge where the line data actually lives, speaks OT protocols natively, and remains in the customer's own hands as the operating environment evolves.

That is the missing link. That is what Vantiq closes.

Next step

A 60-minute architecture workshop maps your highest-value real-time use case to a six-week proof of value, with a measurable comparison against today's manual baseline. Contact your Vantiq representative to schedule.