

WHITEPAPER

Industrial AI Needs a Foundation Most Operations Still Don't Have

Everyone is building on top of it. Almost nobody has built it.



For operations leaders in asset-intensive industries.

Executive Summary

The technology is not failing.

The foundation is missing.

Billions are being invested in industrial AI. Most of it is underperforming. RAND Corporation documented that the majority of enterprise AI projects in 2025 failed to deliver business value. Manufacturing is no exception.

The most valuable knowledge in any operation does not live in systems. It lives in people. And almost none of the AI currently being deployed can reach it.

Sensor networks capture process conditions. Historians record what the plant did. MES, DCS, and SCADA systems preserve enormous volumes of operational activity. But the layer that explains why decisions were made, what the experienced operator noticed, what the night shift discovered, and what actually worked when similar situations happened before is still missing from most industrial AI environments

A model can identify that a parameter has moved outside its normal operating range. It cannot understand that operators have already seen the same behaviour every winter for the last four years and that the issue was eventually traced to a seasonal cooling instability that never appeared in the root cause report.

This paper is about that gap. And it explains why 2026 is becoming a turning point.

80.3%

Enterprise AI projects failed to deliver business value in 2025

RAND Corporation, 2025

76.4%

Manufacturing AI projects failed to deliver promised business value in 2025

RAND Corporation, 2025

Contents

Executive Summary	<u>2</u>
1. The AI Pressure Is Real and It Is Arriving Fast	<u>4</u>
2. What AI Actually Needs to Work in Operations and Where the Data Really Lives	<u>6</u>
3. The Tribal Knowledge Problem at Organisational Scale	<u>8</u>
4. Why This Moment Specifically	<u>10</u>
5. From Shift Intelligence to Organisational Memory	<u>12</u>
6. The Compounding Return	<u>15</u>
7. Five Things to Do Before Your Next AI Investment	<u>17</u>
8. The Reframe	<u>19</u>
Take the AI Readiness Assessment	<u>21</u>



The AI Pressure Is Real and It Is Arriving Fast

The conversation has changed. Twelve months ago, AI in industrial operations was a pilot programme, a proof of concept, a line item in a digital transformation roadmap that everyone had agreed was important and nobody had fully funded.

Today, boards are asking for timelines, budgets are moving, and operations directors are fielding requests not to “explore AI” but to “deploy AI” and deliver measurable results.

The potential is real. Predictive maintenance, AI-assisted root cause analysis, intelligent shift briefings, anomaly detection, and operational copilots already exist. Real examples are already working on the floor across asset-intensive industries. The question is no longer whether the technology works. It is whether the foundation underneath it does.

But something is happening alongside the ambition. Picture an AI copilot answering a field operator’s question about a recurring pressure deviation with complete confidence, based on sensor history alone. The answer is plausible. It is also missing three months of shift observations that would have changed it entirely. Nobody catches this immediately. A maintenance call is made. The deviation recurs two weeks later.

This is not a technology failure. It is a foundation failure. And it is already happening in organisations that moved fast on deployment without asking what the AI was being given to reason on.

Teams are adopting AI tools without shared standards, structured operational inputs, or a connected knowledge environment. The governance conversation is well established. What receives less attention is whether the people responsible for validating AI outputs have the structured knowledge to do so meaningfully. By the time the gap surfaces, the decision has been made, production affected, or costs committed.

60%

of AI projects without AI-ready data will be abandoned.

Gartner



This is what is starting to be called AI theater.

Impressive demonstrations. Plausible outputs. Underwhelming operational impact. The AI was given sensor data and a well-configured model. It was not given the operational knowledge to reason about what any of it meant in the context of that specific plant, that specific asset, and the people who work with it every day. The data confirms what operations teams already know.

The evidence is consistent across multiple independent sources.

RAND Corporation documented that 80.3% of enterprise AI projects in 2025 failed to deliver their promised business value, with this number being 76.4% for the manufacturing industry. Gartner found one in five AI projects in operations collapses entirely, and 57% of I&O managers have at least one failure behind them. In manufacturing the failure rate is 76.4%, with OT/IT integration and data quality as the dominant causes. Gartner estimates 60% of AI projects without AI-ready data will be abandoned.

The problem is not the technology.

It is what the technology is being asked to reason on.

What AI Actually Needs to Work in Operations and Where the Data Really Lives

Ask any data scientist what AI needs to function well, and you will hear the same answer: clean, connected, structured data. That is correct, and the old principle holds as firmly as ever: garbage in, garbage out. But in asset-intensive operations, even that framing is incomplete, because it assumes the most important data exists somewhere and simply needs cleaning. In operations, much of it was never captured at all.

Operations generate two fundamentally different types of data. The first is what sensors produce: temperatures, pressures, flow rates, machine states, and event logs. This infrastructure has been built over decades and, in most mature operations, works well.

The second is what people produce. The operator who notices a pump sounds unstable before the vibration trend changes. The supervisor who recognises that a recurring issue only appears during a specific operating sequence. The process engineer whose previous intervention solved the symptom but not the cause.




Image 1 - What the Gap Looks Like in Practice.

This is also why human-in-the-loop oversight is more demanding in practice than it appears. The principle is straightforward: a qualified person should be able to review and override AI outputs before they drive consequential decisions. Most operations leaders agree. But the bottleneck people cite is rarely the human; it's the absence of structured operational knowledge that prevents a human from making a fast, confident, well-informed judgment.

The engineer is not a bottleneck

They can instantly see the operational history behind an AI recommendation is the quality layer that makes AI outputs trustworthy.

The goal isn't fewer humans in the loop

It's giving people what they need to review, override, and trust AI outputs before they drive consequential decisions.

Regulators across industries have already started arriving at similar conclusions.

In pharmaceutical manufacturing, Annex 11 has long required documented human oversight of computerised systems. In process industries and chemicals, safety frameworks increasingly expect human judgment to remain present within operational decision chains. In oil and gas, the conversation has shifted from whether humans should remain involved to whether organisations have the operational knowledge architecture necessary to support meaningful human oversight. And across all asset-intensive sectors, Article 14 of the EU AI Act will require documented human oversight of high-risk AI systems when full enforcement begins in August 2026.

Gartner

The Tribal Knowledge Problem at Organisational Scale

When companies build internal AI tools, such as operational copilots, deviation explainers, shift briefing assistants and anomaly triage systems, they hit the same wall within months. The model has no operational context. It can process a sensor feed and detect values outside normal range. It cannot reason about what experienced teams have already learned unless that knowledge is structured.

This is not an AI model problem. It is an organisational knowledge problem. Knowledge that should move naturally across shifts, teams, sites, and years dissipates instead. It lives in verbal conversations, PDF reports nobody revisits, disconnected spreadsheets, and fragmented handovers.

72%

of manufacturing employers report difficulty filling roles globally, with Manufacturing and Production skills among the hardest to find.

(ManpowerGroup Talent Shortage Survey, 2026).

The challenge is no longer retaining people. It is retaining operational intelligence. Three situations illustrate what this costs in practice.

1. A reactor unit shows a pressure deviation every time a specific raw material batch arrives from a particular supplier. Each shift, logs it as a process adjustment and closes it. Nobody connects the dots because each shift sees one isolated event. Six months later, a process engineer investigates why yield is consistently below design. The answer is sitting in eighteen months of shift records that were never analysed together.

2. A supervisor who managed the same unit for eight years retires. Three months later, the team is dealing with a seasonal heat exchanger behaviour that he handled instinctively every year. Nobody thought to ask him before he left.
3. A CI manager develops a solution to a recurring equipment reliability issue, documents it locally, and moves on. A year later, a colleague on a sister site starts the same investigation from zero. The knowledge never travelled.

These are not unusual situations. They are the predictable cost of operational knowledge having no structured home. When that knowledge is captured and connected, the organisation gains something more valuable than documentation. It gains continuity. And AI systems gain something to reason on beyond process signals.




Image 2 - Connected Operational Knowledge: Before and After Structuring It.

Many organisations have already invested in operations digitalisation. Shift communication tools, frontline platforms, and operational excellence programmes have created real value.

What they have not collectively solved is whether the human operational knowledge generated every day is being structured in a way that makes AI reasoning meaningful and improvement genuinely transferable. That gap is not a criticism of existing investments. It is the next layer those investments are waiting for.

Why This Moment Specifically

The operational knowledge gap is not new. What has changed is the cost of leaving them unresolved. Three forces are converging in 2026 that turn operational knowledge from a best practice into a strategic requirement.

1. AI investment raises the cost of a weak foundation

For years, fragmented operational knowledge was an efficiency problem. Significant, but manageable. Today it is an AI problem, and AI problems are expensive differently. Companies are committing material budgets to AI platforms on the assumption they have the data foundation to support them. Many do not.

When experienced operators find AI outputs shallow or disconnected from operational reality, trust erodes fast. The organisation draws the wrong conclusion: AI is not ready for operations. The accurate conclusion and the expensive one to ignore is that the foundation was not ready for AI. The consequence is not just wasted investment. It is institutional fatigue with AI at exactly the moment when genuine operational value is becoming achievable.

2. The workforce pipeline is narrowing from both ends

Retirement is half the story. The other half is attraction. Younger generations are showing markedly lower interest in industrial floor roles. Operations are competing for talent against technology and professional services and increasingly losing at the recruitment stage.

ManpowerGroup's 2026 survey of 39,000 employers across 41 countries found 72% of manufacturing employers report difficulty filling roles, with Manufacturing and Production skills among the hardest to find globally. The people leaving take knowledge tied to specific assets, specific sites, and years of pattern recognition that no onboarding programme replaces.

The people not arriving mean that knowledge cannot be rebuilt from the next generation either. The window to capture what experienced operators and engineers still carry is closing from both ends simultaneously, not gradually.

3. Regulatory, safety, and sustainability pressure is demanding more from operational data

The third force is external accountability, arriving from multiple directions simultaneously.

The EU AI Act's Article 14 enforcement from August 2026 is not a future consideration. For asset-intensive industries deploying AI in safety-critical contexts, it is a current compliance requirement. And meaningful human oversight is only possible when the humans in the loop have the structured knowledge to exercise it.

Sustainability commitments create a parallel obligation. Emissions reporting and ESG frameworks increasingly require evidence of operational decisions, not just process outputs. Sensor data records the outcome. Operational knowledge records the decision. Regulators are beginning to ask for both.

Compliance frameworks are converging on the same requirement from three directions at once: AI governance, environmental accountability, and operational safety. All three require the same thing: structured, traceable, connected operational knowledge.

From Shift Intelligence to Organisational Memory

The shift is where operational knowledge is generated: where the operator notices something before the sensor fires, where the supervisor makes the call that affects the next twelve hours, where the decision is either learned from or lost. Starting there is the correct entry point.

But the shift is not the scope. The organisations that extract the most value are those that connect what happens in shifts to engineering analysis, CI programmes, asset reliability, cross-site learning, and the AI systems that support decisions across the organisation. When every shift generates structured, connected records consistently across roles, value compounds upward.

ROLE	WHAT CHANGES
<u>Field Operators</u>	Observations logged once, reaching the right people automatically. Issues that disappeared into verbal handovers are visible from the first report.
<u>Panel Operators</u>	Each shift starts with a structured picture of what the unit did. No more reconstructing history from memory or copy-pasting between systems.
<u>Shift Supervisors</u>	Shifts close with confidence. Actions are tracked and the incoming team has everything they need without a fifteen-minute verbal briefing.
<u>Operations Managers</u>	Patterns across shifts and teams become visible. Systemic issues surface before they become incidents.
<u>Process Engineers</u>	Investigations start with a structured operational history already in place. Deviation patterns and previous interventions are searchable, not scattered.
<u>CI and OpEx Managers</u>	Improvements travel with evidence, not just with the person who developed them. What worked on one line is transferable because it is documented.

This is not a reporting environment. It is operational intelligence infrastructure. Traditional reporting moves information upward for management visibility. This connects what happened, what was observed, what decisions were made, what actions followed, and how those outcomes connect across shifts, teams, and time.




Image 3 - Cyzag Shift Handover View: every shift creates a structured knowledge record, accessible to everyone who needs it, in the role context that makes it actionable.

The same infrastructure that moves knowledge upward distributes it outward. An operator reviews an asset's operational history before starting a task without finding the right colleague to brief them.

A process engineer investigates a deviation without reconstructing events from scratch. A new team member reaches competence faster because the organisation's accumulated knowledge is accessible rather than locked in individual memory. Accountability follows naturally, not as a management imposition, but as a by-product of how the work gets done.

This is not a reporting environment. It is operational intelligence infrastructure. Traditional reporting moves information upward for management visibility. This connects what happened, what was observed, what decisions were made, what actions followed, and how those outcomes connect across shifts, teams, and time.

SECTION 5




Image 4 - Cyzag Operations Manager Dashboard: patterns that were previously invisible because knowledge lived in people rather than systems become actionable at every level of the organisation.

The Compounding Return

Most operational technology investments ask you to wait. Wait for implementation to complete, for adoption to stabilise, for data to accumulate. The commercial case sits on a horizon that keeps moving.

The return on operational knowledge infrastructure works differently.

The first layer begins immediately. From the first structured shift, incoming teams have the context they need. Observations no longer disappear into verbal exchanges. Actions become traceable. Supervisors spend less time reconstructing what happened and more time managing what is happening.

None of this requires heavy IT, months of configuration, or a change management programme that outlasts the enthusiasm that started it. When knowledge capture fits how shifts already work, adoption follows because the tool makes the job easier.

These are first-layer results, visible before they show up in financial reporting. That visibility is what sustains the adoption that makes everything else possible.

The 1st layer

Nouryon saw a fourfold increase in operator engagement in just two and a half months.

Stolt Tankers doubled the frequency of its PDCA cycles and scaled continuous improvement practice across a fleet of 104 vessels.

A leading specialty chemicals manufacturer saved two hours per shift from the first weeks of deployment

These are first-layer results, visible to the people experiencing them before they show up in financial reporting. That visibility is what sustains the adoption that makes everything else possible.

The 2nd layer

Nobian prevented €980k per year in raw material losses at a single site by connecting shift observations to process decisions in a structured way, with combined impact exceeding €1M across seven sites.

Nouryon achieved €400,000 in average savings per plant per year, scaling to 23 global plants within six months.

That is what the second layer looks like when the first layer has been running long enough to compound.

Gartner

Most investment conversations focus on ROI, which is the return on what you spend. Fewer ask about the cost of inaction. In asset-intensive operations, the COI of not building the knowledge foundation is not zero. It is the recurring losses that never get traced to their cause, the improvements that never transfer between sites, and the AI investments that underperform because the foundation was never in place. The longer the foundation is absent, the higher that cost compounds.

The second layer develops progressively and is where bottom-line impact becomes direct and measurable. Recurring production losses become identifiable and solvable. Energy and material inefficiencies become visible in operational context. Deviation investigations accelerate. CI programmes scale because improvements travel with documented reasoning rather than with specific individuals.

The two layers are not separate investments. The first funds and enables the second. You do not have to wait for transformation to be complete before value begins.




Image 5- The Two-Layer Return on Operational Knowledge Infrastructure.

Five Things to Do Before Your Next AI Investment

The data on AI project failure is consistent, and the causes are well understood. For asset-intensive organisations, five actions separate those that get real operational value from AI from those that fund another cycle of AI theater.

1. Conduct an honest operational knowledge audit before committing to AI development

Most organisations can describe their sensor infrastructure precisely. Very few can describe the state of their human operational knowledge layer with the same precision. Before approving the next AI initiative, ask: where does our operational knowledge actually live? How much is structured, searchable, and accessible? How much exists only in the heads of people who may not be there next year? The answer determines how much of your AI investment will reason on a complete picture versus an incomplete one.

2. Budget for the knowledge foundation, not just the model

Organisations that skip the data foundation pay significantly more in remediation costs later. For asset-intensive operations, this means allocating real resources to structuring operational knowledge before or alongside AI deployment, not as an afterthought. Platforms like Cyzag are built specifically for this layer. Regardless of the tool, the work needs to happen and needs to be budgeted as a first-class investment.

3. Define what meaningful human oversight actually requires in your operation

Organisations that skip the data foundation pay significantly more in remediation costs later. For asset-intensive operations, this means allocating real resources to structuring operational knowledge before or alongside AI deployment, not as an afterthought. Platforms like Cyzag are built specifically for this layer. Regardless of the tool, the work needs to happen and needs to be budgeted as a first-class investment.

4. Set realistic timelines and measure adoption alongside outcomes

AI initiatives fail partly because expectations are set at the technology level rather than the operational level. A model can be technically live in weeks. The knowledge foundation that makes it trustworthy takes longer. Plan for both timescales. Measure adoption by operators and shift supervisors. That is the leading indicator of whether the second layer of return will materialise.

5. Build an institutional knowledge infrastructure, not just an AI capability

The organisations seeing the strongest long-term results are not necessarily those with the most sophisticated models. They are the ones that have built knowledge infrastructure that compounds over time: operational learning that accumulates shift by shift, site by site, and becomes more valuable the longer it runs. External expertise can accelerate but cannot replace institutional knowledge. The goal is an operation that gets smarter with every shift.



The Reframe

There is a particular pressure on operations leaders right now. The pressure of feeling behind: behind on AI, behind on digital transformation, behind on the capabilities industry publications describe as already deployed at leading organisations.

That pressure is understandable. But the framing underneath it is wrong.

Every capability the next wave of industrial AI promises assumes the same thing. It assumes your organisation has clean, connected, structured operational data reflecting not just what your process did, but what your people decided, why they decided it, and what they learned. It assumes the foundation described in this paper.

Without that foundation, AI in operations produces what sensor-only data produces: pattern recognition without context, recommendations without history, automation without judgment. In asset-intensive operations, technically impressive and operationally trustworthy are not the same thing.

The organisations building this foundation now are not behind. They are building the prerequisite layer that industrial AI depends on to function meaningfully.

You are not behind. You are building what everyone else assumed you already had.

The urgency here is not theoretical. The workforce carrying decades of operational knowledge is still in your organisation, but that window is narrowing from both ends. The AI investments you are making can still be built on a foundation that makes them work rather than retrofitted onto one that does not. The regulatory deadlines are fixed and do not move to accommodate implementation timelines.

Three forces, each with its own clock. The organisations that act now are working with all three. Those who wait will be working against them.

The question is not whether to build the foundation. Every serious AI initiative in asset-intensive operations will eventually require it. The question is whether you build it now, intentionally, or later, under pressure, after the cost of not having it has already been paid.

Building the foundation starts here

The knowledge that drives operational performance has always lived closest to the work. In the experience of operators, the judgment of shift leaders, the patterns that experienced teams recognise before any sensor does. Cyzag is the AI Digital Operations Platform that captures that layer in real time across every shift, every site, and every team.

What starts on the floor as structured handovers scales up to site-level performance management, cross-site learning, and an AI that gets smarter the more your operation runs. No multi-year implementation. No large IT programmes. Teams contributing value from the first week.

Our customers include Perstorp, Nobian, Nouryon, and Stolt Tankers. We work across process manufacturing, chemicals, energy, liquid terminals, and oil and gas.

The best-run operations do not happen by chance. They happen by design.

cyzag.com

info@cyzag.com

**Fewer
Incidents**

Lower Costs

**More reliable
performance**

**Shift after shift,
site after site.**

Where do you stand?

Find out if your operational knowledge foundation is strong enough to drive real AI outcomes.



10 Questions

Quick assessment to benchmark your readiness



Instant Results

See your AI readiness score and key gaps



Actionable Insights

Get tailored recommendations to build an AI-ready foundation

Take the Free Assessment

