WNDRVR

Edge AI in Manufacturing

Turning Data into Real-Time Decisions

A Wind River ebook, produced in partnership with GDS Group

Contents.

Introduction

Manufacturers across every sector – from automotive and aerospace, to machine tooling, life sciences, consumer goods, power generation, water treatment, chemicals, oil and gas, food processing and more – are facing a new operational reality.

Plants are more instrumented than ever, generating torrents of data from PLCs, sensors, robots, cameras and quality systems. Yet the bulk of that data still goes unused or arrives too late to influence outcomes on the line. Traditional cloud-centric analytics struggle to respond with the speed, resiliency and performance outputs that industrial environments demand. And sovereignty and security are becoming increasingly critical concerns in a volatile world.

Edge AI is emerging as the architectural shift that closes this gap, enabling manufacturers to act on data where it is created while still benefiting from cloud-scale learning. By processing data locally, manufacturers can detect issues earlier, respond faster and keep critical operations running – even when connections to the cloud are limited or interrupted.

This eBook brings together insights from a number of roundtables hosted by Wind River, in partnership with GDS Group, looking at edge intelligence in the manufacturing space, and why hybrid architectures – balancing cloud and edge according to workload – are becoming essential.

It highlights the cultural, operational and IT challenges manufacturers face, and provides a pragmatic roadmap for implementing edge AI in complex industrial environments. Edge AI is coming: is your business ready?

Chapter 1

The Shift Toward Real-Time Intelligence

Manufacturing performance hinges on milliseconds.

Whether detecting defects on a camera feed or identifying a vibration anomaly before a bearing fails, decisions must be made faster than cloud latency allows. Moving all data off-site for analysis is not only slow, but costly and often unnecessary. The idea that "not all data deserves to travel" captures a growing recognition that edge processing is essential for responsiveness and for controlling the escalating costs of cloud storage and compute.

But latency is not the only driver. Many manufacturers operate in regulated, sovereign or air gapped environments where cloud use is either constrained or entirely off-limits. As one manufacturing leader explained, "Governments don't really take very well to us using the cloud, so all our data is kept locally" – a reality that is driving adoption of edge-first architectures that comply with regulatory or sovereignty concerns.

And of course, uptime and availability are also key drivers. For continuous process environments such as power, water, chemicals, oil and gas, and food – where shutdowns are enormously expensive and disruptive – keeping decision-making at the edge is critical to maintaining safe, stable operations. When connectivity drops – as it inevitably does in industrial sites – operations cannot pause while waiting for a server halfway across the world.

Edge AI enables a greater degree of autonomy, where inference models, control logic and fail-safes continue to run even when networks don't. This improves resilience, safety and uptime, especially in high-value processes.

"Not all data deserves to travel. Keeping data at the edge – near the camera, sensor or robot – lets you act quickly, contain costs and meet sovereignty and security needs."

Warren Bayek, VP Intelligent Edge, Wind River

Chapter 2

Operational and IT Challenges Manufacturers Are Facing

Almost every organization contributing to the roundtables described a landscape marked by fragmented systems, legacy equipment, cultural resistance and inconsistent data governance.

Even advanced firms noted that while they may be data-rich, they remain insight-poor, unable to fully exploit operational information due to silos, inconsistent interfaces or dependence on manual interpretation.

One leader in the aerospace sector explained that while enterprise-level strategy is sophisticated, "we're very poor at taking that data and getting meaningful outcomes from it at the grassroots shop-floor level" – especially when new toolsets cannot be used due to sovereignty or air-gap constraints. Others highlighted the burden of outdated systems, with one executive noting that decades of equipment deployed globally leads to unversioned PLC code, inconsistent configurations, and field escalations that drain engineering bandwidth because "nobody really knows what's going on until they go on site and open the hood."

Across industries, culture was consistently cited as a barrier. Many organizations have veteran workforces steeped in manual processes and tacit knowledge. As one food manufacturing executive reflected: "The biggest challenge is inherent culture. Digital transformation represents a material threat to some individuals, whose expertise was built in a pre-digital era."

The buy versus build debate also emerged as a recurring challenge. Few organizations have the scale to build proprietary AI infrastructure, especially given skill shortages. One leader put it simply: "You can't hire 100 specialists just for digitalization and AI, so deciding what to make ourselves versus what we buy is a constant challenge." Evidence across the discussions suggests a growing trend: buy the foundational platform, build the differentiating intelligence.

"Culture is harder than tech. You need wins along the way, clear governance and a unified vocabulary – otherwise terms such as 'unstructured,' 'AI,' and even 'edge' mean different things to different teams."

Warren Bayek, VP Intelligent Edge, Wind River

Chapter 3

The Emergence of Hybrid Edge-Cloud Architectures

In response, manufacturers are increasingly adopting a hybrid strategy, where AI models, operational logic and deterministic control loops run on edge devices; and filtered data, events and telemetry flow upstream to the cloud for cross-site analytics, model retraining and governance.

This pattern marries the strengths of both layers: real-time, sovereign processing at the edge, and scalable learning in the cloud.

In Wind River's **Think-Sense-Act-Optimize** framework, this balance is explicit.

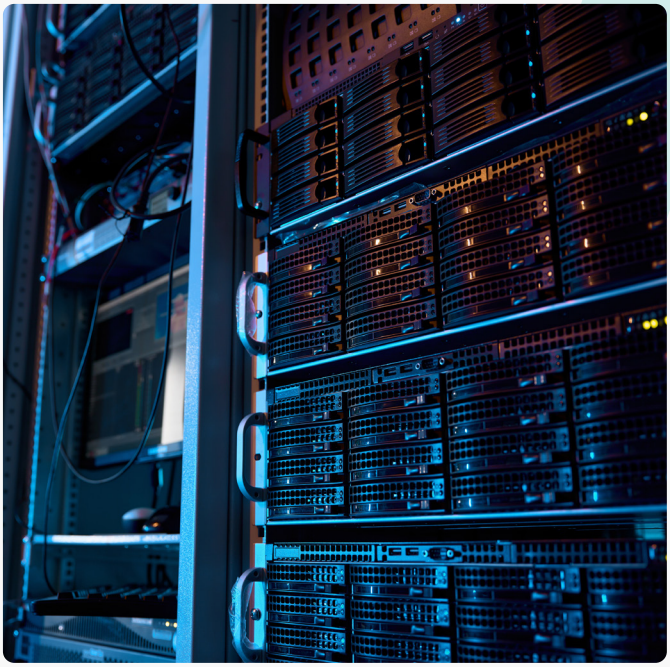

Plants **sense** through a low-latency edge software layer, **think** via cloud-native analytics pipelines, **act** at the machine through low-latency edge execution, and **optimize** by applying enterprise-level insights back to distributed sites through centralized orchestration. This model supports digital twins, closed-loop optimization and continuous learning – all while maintaining determinism and safety.

In practice, hybrid architectures also help address the fragmentation many manufacturers are struggling with. Several participants explained that consolidating multiple small edge devices into a single on-site edge server created an easier to maintain backbone for analytics and AI. One paper manufacturer noted that after early experiments with many small edge devices, the organization moved to "a very sophisticated and high performing server structure" on-site, because it was simpler to manage, scalable and future proofed for AI workloads.

Chapter 4

Real-World Use Cases Where Edge AI Delivers

Edge AI is not a theoretical exercise. Manufacturers are already finding value in several high impact workloads:

01

Vision-based quality inspection

Cameras feeding inference models at the edge can identify defects instantly, reject parts autonomously, and reduce scrap dramatically. Such an approach prevents defective parts from moving downstream and reduces both rework and customer returns. Doing this in the cloud would introduce latency that renders such automation impossible.

02

Predictive maintenance

Edge-based vibration and temperature models detect early anomalies and alert maintenance before catastrophic failures occur. Several participants highlighted the operational and financial upside of shifting from reactive to predictive maintenance, with edge compute enabling real-time execution even during connectivity outages.

03

Continuous process control

In power generation or water treatment plants, edge-based AI models can monitor boilers, turbines or pumps in real-time and trigger alarms or controlled slowdowns before a trip forces a full shutdown – helping avoid extremely costly restarts. Keeping decision-making at the edge is critical to maintaining safe and stable operations.

04

Line and cell optimization

Models running line-side can monitor production rates, cycle times, and energy consumption to identify bottlenecks and inefficiencies in real-time. The system recommends or automatically applies adjustments to parameters, helping increase throughput and reduce energy use without compromising quality.

05

Safety and compliance monitoring

Edge AI can analyze video, sensor or LIDAR to detect unsafe conditions, missing protective equipment or deviations from standard operating procedures, without streaming sensitive imagery off-site. Operators and supervisors are alerted instantly, helping organizations comply with privacy rules while improving workplace safety.

Benefits of Edge AI in Plants

Edge AI brings analytics and decision-making closer to the assets that matter most:

- Real-time detection of quality issues before defective products move downstream.
- Earlier identification of equipment problems, enabling predictive maintenance and fewer unplanned stoppages.
- Local optimization of processes and energy use, improving throughput and reducing waste.
- Greater resilience when cloud connectivity is limited, with critical decisions made directly at the edge.

Chapter 5

Data Strategy: Getting the Foundations Right

A recurring theme across the discussions is that data governance must mature before AI can scale.

Many organizations highlighted years – sometimes decades – of underinvestment in structured data management. As one transformation leader put it starkly: "You can't monetize what you don't control. Before we can even talk about AI at the edge, we're putting solid data foundations and governance in place, with quick wins that prove value along the way."

As such, gaining a better understanding of the data at your disposal was a key challenge. Executives argued that unstructured data now holds more untapped value than ever before – especially for agentic AI models that can interpret logs, images or audio without rigid schemas. Others cautioned that meaningful AI still depends on thoughtfully tokenized, contextualized inputs. The concept of garbage-in-garbage-out is a clear and present danger in the manufacturing space; models still require well-prepared input to produce reliable outcomes.

A practical data strategy for edge AI should therefore include:

- **Clear workload triage** that identifies what needs to stay local and what can go in the cloud.
- **Consistent governance** that travels with the workload.
- **A normalized edge data layer** that standardizes, cleans and structures raw data directly at the edge, close to the sensors, machines or devices that generate it.
- **A digital thread** linking design, manufacturing and service, in order to unify teams and processes.

Chapter 6

Culture, Skills and the Operating Model for Edge

Successful edge AI adoption depends as much on people as on architecture.

Manufacturers repeatedly noted that while technology challenges can be solved, the real transformation lies in aligning OT and IT teams around shared outcomes.

Several roundtable participants emphasized the need for long-term partnerships, especially in regulated sectors where equipment lifecycles typically run to decades. One medical device leader explained that outsourcing AI infrastructure expertise is often necessary, but that "the challenge is finding partners with a long-term perspective" because short-term contracts do not align with product lifecycles that stretch across many years.

Organizations also stressed the importance of creating "wins along the way" to build trust. As one executive explained, value must be demonstrated through measurable improvements in KPIs such as first-pass yield or unplanned downtime – otherwise AI remains a pilot rather than a production tool. Without clear outcomes, organizations fall into the pilot purgatory that several participants warned about.

Chapter 7

A Roadmap for Manufacturers: Crawl, Walk, Run

A phased approach repeatedly emerged as the safest and most reliable way to scale edge AI.

During the **Crawl phase,** organizations select a single high-value use case – such as vision inspection or predictive maintenance – and run it on a pilot line using production-ready tooling. Edge runtimes, model containers and governance baselines are introduced early, ensuring that the pilot proves not just technical feasibility but operational safety and lifecycle management.

In the **Walk phase**, manufacturers template the deployment, extending successful pilots to more assets, lines or cells. At this stage, organizations begin centralizing policy management, security baselines and model orchestration, so that updates and governance remain consistent as the footprint grows. Edge infrastructure becomes standardized rather than bespoke.

In the **Run phase,** manufacturers expand across sites and geographies, building a closed-loop system where edge telemetry feeds enterprise analytics, which retrain the models, and which are then redeployed safely back to distributed production environments. This is where digital twins, fleet analytics and end-to-end digital thread capabilities start to deliver value across engineering, production, quality and service.

As one expert put it: "Operationalizing AI – not doing AI for its own sake – is what really creates the value."

Conclusion: Edge as a Strategic Capability

Edge AI is transforming from a niche technology into a core capability for future manufacturing.

It provides the real-time responsiveness that physical processes demand, supports data sovereignty and safety, reduces operational costs, and enables a consistent platform for distributed plants. Perhaps most importantly, it allows organizations to unify OT and IT around shared, measurable outcomes – driving both resilience and competitiveness.

Ultimately, the manufacturers contributing to these conversations shared a common insight: technology alone is not the limiting factor. The real work lies in establishing governance, building a unified data strategy, fostering cross-functional collaboration, and reshaping your organizational culture.

When these elements align, manufacturers can turn edge AI into a strategic advantage – one that continuously improves quality, throughput, reliability and cost across every plant.

"Technology isn't the limiter anymore. The real work is governing data, architecting for reliability, and building a culture that uses AI in trustworthy, outcome driven ways."

Warren Bayek, VP Intelligent Edge, Wind River

Next Steps: From Ideas to Action

To move from ideas to a concrete plan, consider engaging with Wind River for:

- An Edge AI Workshop to assess your current environment, prioritize use cases, and map them to the Think/Sense/Act/Optimize framework.
- A Pilot Design Session to define a scoped edge AI project with clear success criteria, architecture blueprints, and a roadmap for deployment.

“Consolidate your workloads onto a unified platform – VMs, containers, edge devices and private/hybrid cloud. You then get a single pane for security, scale and consistent reliability across your estates.”

Warren Bayek, VP Intelligent Edge, Wind River

From Think, Sense, Act to Optimize on the Line

Wind River enables manufacturers to move from disconnected experiments to a closed-loop Think, Sense, Act, Optimize system running directly on the line. By combining intelligent software at the edge with centralized orchestration, plants can continuously learn from their operations and turn those insights into real-time decisions.

Think

Model, predict and decide

Sense

capture the state of every asset

Act

Execute decisions where milliseconds matter

Optimize

Close the loop across sites and fleets

WNDRVR

A Wind River ebook, produced in partnership with GDS Group