IDC

SPOTLIGHT

Sponsored by: Wipro

By leveraging Edge AI, supported by clean data and IT/OT convergence, organizations can enable real-time, machine-level intelligence — unlocking autonomous operations that are scalable, resilient, and fit for security-critical industrial environments.

Automation to Autonomy at Edge: Connecting Data, AI, and IT/OT Convergence

May 2026

Written by: Mukesh Dialani, Research Vice President, IDC US

Synopsis

Industrial enterprises are approaching an inflection point. After years of investing in automation, many are now facing a harder question: how do they move from systems that simply execute instructions to systems that can reason, adapt, and act with minimal human intervention? This is the promise of autonomous operations, and Edge AI is the enabling layer that makes it achievable at scale.

The foundation of this shift is technological, but it requires a fundamental rethinking of how data is collected, structured, managed and trusted across the enterprise. Organizations operating in energy, manufacturing, utilities, and industrial infrastructure segments are finding that the volume of real-time data generated at the operational edge far exceeds what can be practically transmitted to centralized cloud environments for processing. Decisions that once took minutes or hours now need to happen in milliseconds — at the site, at the machine, now.

Yet, for these organizations, the path from automation to autonomy is blocked by three challenges that most organizations have not yet resolved: fragmented and untrustworthy data, siloed IT and OT environments that lack a common operating model, and the absence of scalable edge infrastructure capable of running AI inference reliably in constrained, often security-critical environments.

Addressing these challenges is essential for industrial enterprises seeking to unlock the full potential of Edge AI and transition toward autonomous operations. Clean, trustworthy data foundations, converged IT/OT architectures, and purpose-built edge platforms are no longer optional components of digital transformation strategies — they are prerequisites for enabling intelligent, real-time decision-making at the operational edge and for sustaining competitive advantage in the next phase of industrial evolution.

AT A GLANCE

KEY STATS

According to IDC research:

- The global product engineering and OT services market is projected to reach US$344 billion by 2029 with a 5-year CAGR of 10%.
- In an IDC Services Path Survey in October 2025, 38% of respondents indicated they were currently outsourcing their Edge needs to engineering services providers. Another 47% indicated that they were planning to do so in the next 2 years.

CHALLENGE

» Challenges faced by customers include fragmented and untrustworthy operational data, organizational resistance to IT/OT collaboration, lack of clear strategy to scale beyond, pilots, security risks and building internal capability to govern and sustain autonomous systems.

WHAT'S IMPORTANT

- Customer needs are evolving: Mature and seasoned buyers of Edge AI services expect providers to know more about their industry's challenges and provide related guidance, execution road maps, and a clear ROI commitment.

Figure 1: ***Customer Challenges that Increased Edge Project Costs***

Source: IDC's *Edge View 2025, n = 800*

Situation Overview

- **The data volume problem has changed the economics of centralization:** Modern industrial environments generate data at a pace and volume that challenges traditional architectures. A smart manufacturing facility may operate thousands of sensors across production lines, HVAC systems, energy meters, and safety monitors — each producing continuous streams of telemetry data. Aggregating all this data into a central cloud or data center for processing is not only expensive and time consuming, but also operationally inadequate. Data fragmentation across OT systems, CMDB gaps, and inconsistent device identity also prevents unified operational visibility. Lack of clean, labeled, and contextualized data undermines AI model performance in production environments. Edge computing resolves this by enabling local inference and decision-making. AI models deployed at the edge can detect anomalies, trigger automated responses, and update operational parameters without waiting for a trip to and from the cloud. The economic and operational case is clear: if organizations successfully deploy edge AI infrastructure, they will see meaningful reductions in unplanned downtime, faster incident response, and lower bandwidth costs.

- **The move from automation to autonomous operations:** Traditional automation is deterministic — it executes predefined rules reliably, but it cannot adapt when conditions change. Autonomous operations represent the next layer of sophistication: these are systems that can observe their environment, interpret what they see, and take contextually appropriate action — even in situations that were not explicitly anticipated during design.

  This shift requires building AI models that are trained on high-quality operational data, deployed in environments where they can act on real-time inputs, and are governed by feedback loops that allow ongoing refinement. None of this is possible without a stable and rigorous data foundation. Clean, structured, and well-contextualized data is the prerequisite for AI that can be trusted to operate with significantly reduced human oversight.

- **IT/OT convergence is no longer optional:** For most industrial enterprises, operational technology (OT), the systems that run physical infrastructure — has historically been kept separate from information technology (IT). OT systems prioritize stability and uptime; IT systems prioritize agility and connectivity. But as digital intelligence moves closer to the operational layer, the boundary between these worlds is blurring and collapsing. Security and risk concerns also make IT/OT convergence a governance challenge as much as a technical one

  Regulatory frameworks such as IEC 62443 were developed to formalize requirements for secure, integrated IT/OT environments. Beyond compliance, the business case for convergence is compelling: organizations that have unified their IT and OT layers are better positioned to achieve continuous visibility across operations, enable faster and more confident decision-making, reduce the security exposure that comes from managing two disconnected technology stacks and duplication in investments.

The challenge is that convergence introduces real complexity. OT systems often run on proprietary protocols, have long operational lifespans, and cannot tolerate the kind of downtime that IT system updates may require. Bridging these environments demands a careful, architecture-first approach which prioritizes interoperability, security by design, and operational continuity.

Wipro's Intelligent Edge Platform: uEdgeAI powered by Dell

Wipro's **OTNxT** platform under **Wipro Intelligence** provides an integrated edge-to-enterprise architecture that supports data acquisition, processing, and analytics at the edge and provides a unified AI infused orchestration layer to visualize and manage seamlessly. A key element of the platform is the use of a knowledge graph with AI models, Wipro's OTNxT TM ensures that outputs are context-aware, explainable, and governed, which is essential for adoption in regulated and safety-critical environments to connect and contextualize data across systems. This enables better visibility and supports more adaptive operations. The platform uses distributed architecture with processing capabilities at the edge. Edge nodes run analytics locally and connect to centralized systems for visualization, orchestration and AI model management. A unified data model helps ensure that data is clean and structured, improving the quality of insights.

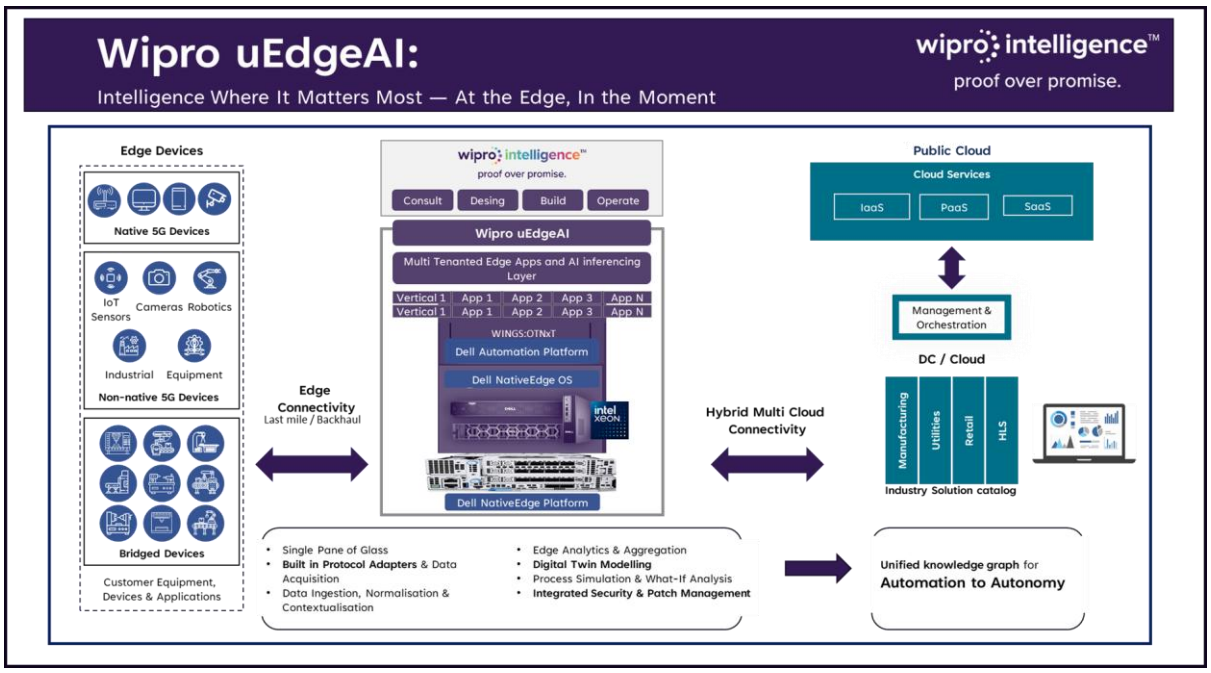

Wipro's uEdgeAI, developed in collaboration with Dell, leverages Dell NativeEdge endpoints and the Dell Automation platform, integrated natively with Wipro's OTNxT platform, to deliver a scalable autonomy stack blueprint. This blueprint includes standardized OT/IT/IoT connectors for data collection, a unified data model for clean telemetry, an event-and-rules layer for real-time decision-making, and an orchestration layer for lifecycle management. Partnering with Dell, it integrates with industrial-grade edge infrastructure, allowing customers to deploy consistent and supportable edge nodes and clusters across multiple sites. Additionally, Wipro's OTNxT incorporates defense-in-depth security controls throughout the edge-to-enterprise stack, adhering to industrial security standards like IEC 62443. In collaboration with

Dell, Wipro's uEdgeAI is designed to scale from CPU-only nodes to GPU-accelerated edge clusters, enabling tailored deployments based on specific use cases across industry verticals.

As a case study with one of world's largest beverage manufacturing company the challenge wasn't a lack of tools—it was a lack of operational truth: incomplete CMDB, inconsistent device identity, OEM-driven silos, and reactive OT operations that couldn't scale across sites. Wipro's OTNxT changed the physics of that environment by turning IT-of-OT telemetry into a governed control plane. It discovered and normalized assets into a unified ontology, enriched ServiceNow CMDB with the missing structure and relationships, and correlated events into workflow-ready signals so response became consistent, not heroic. The engagement narrative shows CMDB fulfilment moving from 45% to 75%, and with that foundation in place, customer gained the leverage to standardize lifecycle operations resulting in 30-35% reduction in MTTR thus minimizing production downtime and services disruptions which reduced operational cost by 15 -20%, strengthen cyber governance, and run OT services with a true "single pane of glass" across Purdue L2/L3/L3.5.

In utilities, Wipro's OTNxT supports real-time analysis of operational data to improve efficiency and compliance. In oil and gas, edge AI is used for safety monitoring and automated response. In manufacturing, the platform enables coordinated operations across production and logistics systems, supporting more efficient workflows and reduced downtime.

The platform includes capabilities for asset discovery, network mapping, and vulnerability analysis. This helps organizations understand their operational environments and prioritize security actions. By focusing on critical risks, organizations can improve security posture by reducing threat radius without increasing operational complexity.

Wipro collaborates with a wide network of partners to advance its uEdgeAI offerings. Its alliance with Dell Technologies facilitates seamless integration of Wipro intelligence platforms into industrial edge infrastructure, enabling zero-touch deployment and lifecycle automation leveraging Dell NativeEdge. Additionally, through strategic partnerships with technology leaders like Intel and advanced connectivity providers, Wipro enables enterprises to build and operationalize production-ready Edge AI solutions for an AI-driven future.

IDC Perspective

Edge AI has moved from a niche investment to a mainstream strategic priority for industrial enterprises. The organizations that are advancing further and fastest are not those with the most sophisticated AI models — they are the ones that have solved the upstream data problem. Model quality is ultimately bounded by data quality, and data quality in industrial environments is a function of how well IT and OT systems have historically been integrated and how rigorously data governance has been applied across the operational layer.

Wipro's approach to this market is notable for its grounding in the practical realities of industrial deployment. Rather than leading with AI capability alone, the company's OTNxT platform addresses the foundational layers — asset discovery, data normalization, knowledge graph construction, and event correlation — that make AI-driven autonomy operationally viable. This sequencing matters: organizations that attempt to deploy AI before resolving these foundational issues consistently encounter model reliability problems that erode confidence in technology.

The company's uEdgeAI blueprint, developed in partnership with Dell, reflects an important market insight: scalable edge AI requires repeatable deployment patterns. One-off, custom deployments are expensive and operationally fragile. Standardized connectors, unified data models, and centralized policy orchestration allow organizations to move from single-site pilots to scalable multi-site fleet deployments without a proportional increase in cost of deployment risk.

IT/OT convergence will remain a multi-year journey for most organizations, but the roadmap is clear. Enterprises that strategically invest now in clean data infrastructure, secure integration architecture, and edge-ready AI platforms will have an advantage in operational efficiency, regulatory readiness, and the speed at which they can deploy and refine autonomous capabilities to transform their customer's business. Those that delay are likely to be left behind as early movers develop operational AI systems that continuously learn and improve.).

Challenges Faced by Customers

- **Fragmented and untrustworthy operational data:** Data collected across PLCs, SCADA platforms, historian databases, and IoT sensors are rarely integrated, consistently labeled, or governed with the rigor that AI requires. Incomplete CMDB records, inconsistent device identity, and raw telemetry without contextual metadata imply that organizations often discover the seriousness of their data problem only when their first AI model fails to perform reliably in production.
- **Organizational resistance to IT/OT convergence:** IT and OT functions have historically operated with separate leadership, budgets, and tools and technologies and a degree of mutual suspicion about each other's risk tolerance. Convergence requires more than a technology decision; it demands executive sponsorship and a shared governance model to overcome the organizational inertia that consistently slows or stops integration programs.
- **Difficulty scaling beyond the pilot:** Many organizations have demonstrated edge AI in a controlled single-site pilot, only to find that replicating it across the relevant operations infrastructure is far more complex than expected. Variability in OT infrastructure age, network quality, and local operating practices demands standardization and orchestration capabilities that were not built into the original pilot design.
- **Cybersecurity risk in converged environments:** Most OT infrastructure is designed for isolated environments with limited native security withing IP networks, and this introduces vulnerabilities with consequences that go well beyond typical IT incidents. A compromised industrial control system can cause physical damage, safety failures, or extended operational shutdown, making security architecture a business-critical investment rather than a compliance checkbox.
- **Building internal capability to govern autonomous systems:** When AI systems make critical real-time decisions on equipment operations or safety responses, organizations need talent who understand how those models work and can recognize when they are failing. Most industrial enterprises do not yet have this capability at scale, and building it through training, model monitoring workflows, and escalation protocols must happen before autonomous systems are trusted in production.
- **Balancing innovation speed with operational risk tolerance:** Industrial environments have very low tolerance for failure, which creates structural tension with the rapid iteration that AI improvement requires. Organizations must build governance frameworks including simulation environments, human-in-the-loop approval thresholds, and rollback mechanisms that allow AI capabilities to be tested and refined without exposing critical systems to a scenario that can jeopardize or stall operations thereby severely impacting operations performance and business metrics.

Guidance for Technology and Services Buyers

The key task for CIOs, CTOs, heads of OT/IT infrastructure, and enterprise architects is to build the foundational systems that enable and accelerate the move from automated to autonomous operations possible without increasing operational risk. Scaling edge deployments beyond pilots require orchestration and lifecycle management capabilities most organizations do not yet have. This makes the case for partnering with engineering services firms that have the required lifecycle capabilities, domain expertise and the right infrastructure and technology partnerships in place.

The following priorities are critical and reflect what organizations must work towards. Rather than try and reinvent the wheel, IDC recommends that organizations work with services partners to accelerate their edge AI transformation initiatives.

- **Resolve the data foundation before scaling AI:** AI models are only as reliable as the data they are trained and operated on. Before committing to scaled edge AI deployments, technology leaders should audit the quality, completeness, and consistency of operational data across their environments. Common problems include incomplete CMDB records, inconsistent device identity across sites, and raw telemetry that lacks the contextual metadata needed for meaningful analysis. Investment in data normalization, asset discovery, and knowledge graph infrastructure should be treated as a prerequisite. Organizations that have structured their operational data with a unified ontology are able to move faster in deploying AI with greater confidence in the outputs.

- **Design edge architecture for flexibility and lifecycle management:** Edge deployments that are architected as rigid, site-specific solutions become expensive to maintain and difficult to scale. Technology buyers should focus on architecture that separates the intelligence layer (models, rules, event logic) from the infrastructure layer (compute, connectivity, storage), allowing each to evolve independently. The ability to push model updates, policy changes, and configuration modifications to edge nodes from a single control plane is not a nice-to-have; it is what separates pilot projects from scaled, production-grade deployments. Organizations must evaluate platform vendors on their lifecycle management capabilities, and not just their inference performance.

- **Treat security as an architectural requirement, not an add-on:** IT/OT convergence dramatically expands the attack surface of industrial environments. Technology leaders must ensure that security is embedded into the architecture from the outset covering network segmentation, identity and access management for edge nodes, encrypted data transmission, and vulnerability management across both IT and OT assets.

  Alignment with frameworks such as **IEC 62443** provides a useful structure for security governance in IT/OT environments. Organizations should also invest in continuous monitoring capabilities that can detect and respond to anomalies across the converged environment without requiring operational shutdown.

- **Select partners that reduce integration complexity:** The edge AI ecosystem includes a wide range of vendors across computer hardware, connectivity, software platforms, and systems integration. The organizations that deploy most effectively are those that have structured their vendor relationships to minimize integration complexity by choosing platforms with pre-built connectors for common OT protocols and established partnerships with hardware providers.

  Standardized deployment blueprints packaged configurations of hardware, software, and connectivity that can be replicated across sites significantly reduce the cost and risk of scaling. When evaluating technology partners,

organization must prioritize those that can demonstrate repeatability across multiple deployment environments, not just success in a single proof of concept.

Guidance for Business Buyers

For business leaders such as COOs, head of operations, plant managers, and business unit executives, the opportunity with edge AI is tangible and near-term. But realizing it requires active leadership vision and engagement on questions that go beyond technology choice. The organizations that succeed treat edge AI as an operating model change, with technology as the enabler.

- **Define business and technology outcomes first, then work backward to requirements:** The most common mistake customers make in envisioning and implementing edge AI programs is starting with the technology and working forward to the business case. This approach almost always produces pilots that are technically successful but operationally disconnected and irrelevant to areas where business needs to improve. Business leaders should begin by identifying top two or three operational outcomes that would most meaningfully change performance by reduced unplanned downtime, faster regulatory reporting, lower energy consumption per unit of output, or faster incident closure. From there, they can work backwards to determine what data is needed, what AI capabilities are required, and what infrastructure investments are justified. This approach ensures that the program is anchored to business value and makes the return on investment measurable from the outset.

- **Establish governance for autonomous decision-making:** As systems become more autonomous, the governance question, that is, who is accountable for decisions made by AI becomes more consequential. Business leaders should define, in advance, which categories of decisions can be delegated to autonomous systems, which require human review, and which must always remain under direct human control. In regulated industries, autonomous systems that take actions affecting infrastructure or worker safety, environmental compliance, or financial reporting may create legal and regulatory exposure if governance frameworks are not in place. Stakeholders must ensure clear ownership, audit trails for AI-driven decisions, and exception-handling protocols before autonomous capabilities are deployed in production.

- **Protect and treat operational data as a strategic asset:** The operational data generated at the edge of equipment telemetry, process parameters, energy usage patterns is among the most competitive, critical and sensitive data an industrial enterprise possesses. Business leaders must insist on data architecture decisions that keep raw operational data and proprietary AI models within controlled environments, even as processed insights are shared more broadly across the organization.

  When working with technology partners and systems integrators, ensure that data residency, access controls, and intellectual property protections are clearly defined in contractual terms. The goal is to enable the insights that drive autonomous operations while retaining full control over the underlying data that makes those insights possible.

- **Build for scale:** Many organizations have demonstrated edge AI in controlled pilot environments, only to find that scaling across dozens or hundreds of sites introduces costs and additional complexity that were not anticipated. Business leaders should evaluate edge AI programs on their path to scale from the beginning — not just on their performance in the initial deployment.

Customers must ask if this architecture be replicated at other sites without custom engineering at each location. What does the operating model look like when this is running across 50 sites instead of 5? What skills and support structures does the organization need to sustain autonomous operations over time? Programs that cannot answer these questions credibly should be redesigned before they move beyond pilot.

Conclusion

The move from automation to autonomous operations is not a distant aspiration. It is underway, and organizations investing in foundational capabilities that include clean data infrastructure, secure IT/OT convergence, and scalable edge AI architecture are building a structural advantage that will compound and give them a competitive edge over time.

However, this is easier said than done. Data fragmentation, integration complexity, and the organizational challenge when governing AI-driven decisions are real barriers. But they are solvable, and the organizations that have solved them are already seeing the results that include fewer unplanned operational failures, faster and more consistent incident response, lower compliance exposure, and operations that can adapt to changing conditions without waiting for human intervention.

Edge AI succeeds when it is treated as an operating model change rather than a technology project. That means starting with the data foundation, designing for scalability from the beginning, embedding security into the architecture, and defining the governance model for autonomous decision-making before systems are deployed in production. Services and Technology partners that can support this full-stack approach from data normalization and IT/OT integration through to AI inference and lifecycle management at the edge will be the ones that help organizations cross the line from pilot to scaled, sustained value.

The competitive advantage of autonomous operations is not just efficiency. It is the capacity to sense changes in the operational environment and respond faster and more intelligently than a competitor relying on human-in-the-loop processes. Building that resilience requires investment today in the infrastructure, data, and governance capabilities that make autonomous operations reliable at scale.

About the Analyst

***Mukesh Dialani*, Research Vice President**

Mukesh Dialani is a Research Vice President for IDC's Worldwide Digital Engineering and Operational Technology Services research. He is responsible for executing field research and custom research projects across the entire lifecycle of hardware and software products. Based on this background that included working with end customers in the engineering services domain, his core focus also includes operational technology and emerging technology areas related to Industrial IoT, Computer Vision, Robotics, AR/VR and digital transformation pertaining to engineering services.

MESSAGE FROM THE SPONSOR

Wipro's OTNxT platform under **Wipro Intelligence** enables enterprises to transition from fragmented OT environments to intelligent, AI-driven operations. By combining edge-to-enterprise data integration, real-time analytics, and a unified orchestration layer, it delivers context-aware, governed insights critical for regulated and mission-critical industries.

In partnership with Dell, **Wipro's uEdgeAI** provides a secure, scalable autonomy stack with standardized connectivity, real-time decision-making, and lifecycle orchestration across distributed edge environments. Proven in large-scale deployments, it drives faster incident resolution, reduced operational costs, and enhanced cyber resilience, enabling organizations to **achieve single-pane visibility and future-ready, autonomous operations at scale**.

Learn more about Wipro's IoT and Edge offerings at https://www.wipro.com/infrastructure/iot-and-edge/

Learn more about Dell NativeEdge solutions at : https://www.dell.com/en-in/shop/storage-servers-and-networking-for-business/sf/nativeedge

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Asia/Pacific Headquarters
168 Robinson Road
Capital Tower, Level 20
Singapore 068912
65.6226.0330
Twitter: @IDC
blogs.idc.com
www.idc.com