



KAI WAEHNER LANDSCAPE Q4 2026

INDUSTRIAL IOT LANDSCAPE

Vendors, Open Standards, and the Unified Namespace Shift

Published: October 2026

INSIDE THIS REPORT

Contents

Introduction	03	07 Where Enterprise IT Connects	23
01 From Industry 4.0 to Industrial IoT: Why the Architecture Won	04	08 Five Trends Shaping Industrial IoT	26
02 What Industrial IoT Actually Covers	07	Virtual PLCs Move Control to Software	27
03 How to Read This Landscape	09	Digital Twins Close the Loop	27
04 The Four Categories from Edge to Cloud	11	Generative and Agentic AI Reach the Plant	27
Industrial Connectivity	12	OT Metadata Joins Enterprise Governance	28
SCADA and Edge Platforms	13	The Categories Are Converging	28
Unified Namespace (UNS) and Data Infrastructure	14	09 Six Questions Before Any Vendor Shortlist	29
Applications and Industrial AI	15	A Note from the Author	32
05 Open Standards Decide Who Stays Replaceable	16	About the Author	33
OPC UA or MQTT?	17		
Standards Beyond the Protocols	18		
Regulation and Licensing Decide Portability Too	18		
Zenoh and NATS: Protocols to Watch	19		
06 The Industrial IoT Stack and the Unified Namespace	20		
What Is a Unified Namespace?	21		

Industrial IoT (IIoT) connects machines, sensors, controllers, and plants to software, so that production data becomes accessible, contextualized, and actionable from the edge to the cloud. **The center of gravity in this market has moved back to the plant.** The hyperscalers retreated from the IoT application layer, and OT-native platforms, open standards, and open source filled the gap.

Three forces are converging at once. The hyperscalers dismantled large parts of their IoT portfolios and kept only infrastructure primitives. Private equity is assembling orphaned industrial software into new independent platforms, with Velotix as the most visible example. And AI has reached the plant floor, from copilots that generate SCADA configurations to digital twins that execute closed-loop decisions.

This landscape maps the industrial IoT market across four categories from edge to cloud, covers the open standards that decide whether today's decisions remain reversible tomorrow, explains why the Unified Namespace (UNS) has become the dominant integration pattern on the plant side and how it connects into the broader IT/OT enterprise architecture, and gives architects a practical framework for choosing their industrial IoT stack.

This is not a ranking. No vendor pays to appear here. The analysis is based on experience advising enterprises across the Global 2000, combined with ongoing research into product developments and adoption patterns. It is an independent practitioner perspective, not a formal research methodology.

Industrial IoT Landscape Q4 2026



KAI WAEHNER




Figure 01. The Industrial IoT Landscape Q4 2026. Four categories from edge to cloud. Bubble size reflects market footprint; filled bubbles are platform leaders, outlined bubbles are specialists and emerging players, dashed outlines mark open source projects.

01

CHAPTER 01

From Industry 4.0 to Industrial IoT: Why the Architecture Won

Industry 4.0 promised a revolution. The market answered with architecture: hyperscalers retreated to primitives, and OT-native platforms, open standards, and open source filled the application layer.

Industry 4.0 was announced as a revolution more than a decade ago. The vision was right. The early execution was not. Lighthouse factories produced impressive demos while most plants stayed stuck in pilot purgatory: point-to-point integrations, proprietary protocols, and cloud platforms that treated the factory as just another data source.

The market answered with architecture instead of hype. The clearest evidence is the hyperscaler retreat. Google Cloud IoT Core was shut down in 2023. AWS discontinued IoT Things Graph, IoT Analytics, IoT Events, and Fleet Hub, and retires Greengrass V1 in October 2026. Microsoft ends standard support for Azure IoT Edge 1.5 LTS in November 2026 and retires Azure IoT Central on March 31, 2027, pointing customers to Azure IoT Operations instead. The pattern is consistent. Hyperscalers keep the primitives: ingestion, storage, compute. The application layer moved back to vendors who understand the plant.


The second signal is consolidation, running in two directions at once. Private equity is assembling independents: in March 2026, TPG launched Velotic, combining GE Vernova's former Proficy business with PTC's former Kepware and ThingWorx businesses, led by CEO Brian Shepherd with James Heppelmann as Executive Chairman and more than 300 million dollars in revenue, and Cumulocity completed a management buyout from Software AG in January 2025 and operates as an independent IIoT platform from Düsseldorf again. And the giants are buying context: in October 2026, Schneider Electric agreed to acquire PTC for 22.6 billion dollars, its largest deal ever, adding the design and product lifecycle layer to a portfolio that already spans AVEVA for operations and, since June 2026, Cognite for industrial data and AI; closing is expected in the third quarter of 2027. The timing carries a telling detail: PTC's industrial IoT platform assets had already left for Velotic months before the big deal. Industrial software is being reassembled on both sides of the size spectrum.

The market behind these moves is large. The IoT-focused research firm IoT Analytics puts the enterprise IoT market at 324 billion dollars in 2025, growing 13 percent year over year, with the industrial IoT platform segment around ten billion dollars and growing toward thirty billion within the decade. The platform layer is small compared to hardware and connectivity, but it is where the architectural decisions are made. Those decisions are the subject of this landscape.

The Hyperscaler Retreat



IoT application services retired or scheduled for retirement between 2023 and 2027, while infrastructure primitives remain



Hyperscalers keep the primitives for ingestion, storage, and compute. The application layer moved back to the plant.

Retirement dates from vendor announcements. Analysis: Kai Waehner - kai-waehner.de

Figure 02. The hyperscaler retreat. IoT application services retired between 2023 and 2027 above the timeline, replacements launched below it.

02

CHAPTER 02

What Industrial IoT Actually Covers

Four categories from edge to cloud define this landscape, and three adjacent areas stay deliberately outside its scope.

This landscape organizes the market into four categories. They mirror the four columns of Figure 1 and read from edge to cloud.

Industrial connectivity provides access to machine data: protocol servers, edge DataOps, contextualization, and the emerging OT data catalogs. Buyers here are solving the first-mile problem: getting clean, modeled data out of heterogeneous equipment.

SCADA and edge platforms operate the plant: visualization, supervisory control, HMI, and increasingly virtual PLCs. Buyers here run production and carry uptime responsibility.

Unified Namespace (UNS) and data infrastructure move and distribute events: MQTT brokers, Sparkplug B, and the remaining hyperscaler services. Buyers here design the data backbone between plant and enterprise.

Applications and industrial AI turn data into decisions: IIoT platforms, digital twins, analytics, and copilots. Buyers here own improvement programs, from OEE to predictive maintenance.

Four adjacent areas are deliberately excluded. Consumer IoT is a different market with different economics. Enterprise IT connects to industrial IoT but is not part of it. That includes data streaming with Apache Kafka and Apache Flink as well as the broader integration and analytics stack; Chapter 07 maps the boundary, and the [Data Integration Landscape](#) covers that side in full. And OT cybersecurity, the market of vendors such as Claroty, Nozomi Networks, and Dragos, is a discipline of its own under growing regulatory pressure from NIS2 and the EU Cyber Resilience Act, and deserves its own analysis rather than a token column here. So is network connectivity: 5G, private cellular, TSN, and LPWAN are the transport beneath this landscape, not part of it.

03

CHAPTER 03

How to Read This Landscape

The chart organizes the market into four columns from edge to cloud and three tiers per column. Bubble size reflects market footprint, dashed outlines mark open source projects. The columns are categories, not scores.

The chart is organized as four columns and three tiers. The columns are categories, not scores, and they read from left to right as edge to cloud. Within each column, platform leaders sit at the top, specialists in the middle, and emerging players at the bottom. Bubble size reflects relative market footprint in that category. A dashed outline marks open source projects, because open source is a structural feature of this market, not a footnote.

Vendors within each cell are representative examples, listed without ranking. Many companies could appear in several columns. Each product is placed where its center of gravity sits, and some vendors therefore show up in more than one column, once per major product line. The landscape identifies the role a technology plays in the architecture, not which vendor is best. The right choice depends on your equipment fleet, your deployment constraints, and your team, which is what the decision framework in Chapter 09 is for.

04

CHAPTER 04

The Four Categories from Edge to Cloud

Vendor profiles across industrial connectivity, SCADA and edge platforms, the Unified Namespace, and industrial applications and AI.

Industrial Connectivity

The first mile of industrial IoT is protocol access and contextualization. This zone decides how hard everything downstream becomes. It exists because there are two ways to a unified API: standardize the wire, which is what OPC UA does, or unify in the client layer with drivers per device, which is what this zone does. Both are legitimate, and the right mix differs per fleet. Industrial DataOps, the discipline behind the second path, is a practice rather than a standard: borrowed in name from DevOps, it treats plant data as a product, with modeling, versioning, and governance applied at the edge.

Kepware (by Velotiv). Kepware remains one of the most widely deployed industrial protocol servers, with hundreds of drivers spanning PLCs, DCS, and legacy equipment across manufacturing, oil and gas, and utilities. Now part of Velotiv after the TPG transaction, it sits in a portfolio alongside Proficy and ThingWorx, with a roadmap pointed at tighter integration across connectivity, SCADA, and IIoT applications. The Spring 2026 release modernized deployment, management, and scalability for fleet-wide rollouts.

The open question is portfolio gravity. Kepware grew as a neutral connectivity layer that fed every SCADA and IIoT platform equally, and its value depends on staying that way. Best suited for heterogeneous brownfield environments where driver coverage decides project timelines; buyers should watch whether integration into the Velotiv stack changes the neutrality that made Kepware the default.

HighByte. HighByte Intelligence Hub defined the industrial DataOps category: modeling, contextualizing, and routing OT data at the edge before it hits brokers or clouds. It is the tool where OT engineers and data teams meet, and a common companion to a UNS rollout, turning raw tag data into modeled, self-describing payloads. HighByte shipped the first MCP Server in industrial DataOps in version 4.2, and version 4.3 added high availability, deeper Microsoft Fabric and Databricks integration, and MCP governance with turnkey data tools for OPC UA and MQTT connections, so AI agents can discover and read plant data through natural language. IDC named HighByte a Leader in its 2026 Industrial DataOps MarketScape. HighByte is a data layer, not a visualization or control product. Best suited for organizations building a contextualized namespace across multiple sites who want the modeling discipline without replacing existing SCADA or brokers.

Litmus. Litmus Edge combines broad driver coverage with edge analytics and fleet management in one platform, with Git-based version control for distributed edge deployments through Edge Manager and a visual UNS builder in Litmus Unify. The Litmus Data Catalog, in private preview since April 2026, extends DataOps into OT metadata management, with planned integrations into Atlan, Collibra, Purview, Dataplex, and Unity Catalog. That direction matters: it is the first serious attempt to give enterprise data governance visibility into plant-floor assets, covered as a trend in Chapter 08. Best suited for enterprises that want connectivity, edge processing, and fleet operations from one vendor and are building toward governed industrial data.

Cirrus Link. Cirrus Link authored the MQTT Sparkplug specification, now an OASIS standard, and builds the Sparkplug modules used inside Ignition deployments worldwide. It is a small company with outsized architectural influence. Where Sparkplug B is the payload standard, Cirrus Link is usually near the design. Best suited for Ignition-centric UNS architectures and any team that wants Sparkplug expertise from the source.

Specialists in this zone include **Softing**, with deep OPC UA tooling and gateways trusted in machine-builder ecosystems, and two open source projects that changed the economics of connectivity. **Apache PLC4X** is a universal protocol adapter that gives applications unified access to PLCs across many native protocols with one API. **Eclipse Milo** is the reference OPC UA implementation for the JVM, embedded in many commercial products.

Worth watching: **FlowFuse** commercializes Node-RED for industrial use, turning the most widely deployed low-code integration tool on the plant floor into a managed, governed platform. Node-RED is often the unofficial connectivity layer in brownfield plants, and FlowFuse gives that reality an enterprise wrapper.

SCADA and Edge Platforms

This zone operates production. It is covered deliberately briefly here: supervisory control is a deep market of its own and a candidate for a dedicated deep dive in a future edition of this landscape series. The profiles below cover the representative leaders in the industrial IoT context.

Siemens. Siemens spans this zone like no other vendor: WinCC and WinCC OA for SCADA, Industrial Edge as the application platform, and the S7-1500V virtual controller for software-defined automation. The virtual PLC direction is strategic, proven publicly in Audi's EC4P program, and pulls control workloads onto standard IT infrastructure; the shift has its own analysis as Trend 1 in Chapter 08. The tradeoff is ecosystem depth: the more Siemens layers a plant adopts, the more the architecture follows Siemens conventions. Best suited for enterprises standardizing on Siemens automation that want one accountable vendor from controller to edge, with the virtual PLC path as a real differentiator; keep the data layer open to preserve options.

Ignition (Inductive Automation). Ignition is the best-known platform of the modern SCADA generation: web-native, cross-platform, unlimited licensing, native MQTT and Sparkplug B support, and a large global integrator community. Many Unified Namespace reference architectures in manufacturing are built on it, combining SCADA, HMI, historian, and Sparkplug publishing in one platform. The 8.3 long-term-support release moved it further toward the data layer: the Event Streams module maps event-driven data between sources such as tag changes and Kafka topics and handlers such as scripts and databases, a Kafka connector works as both source and target, and a rebuilt historian suite on QuestDB replaced the previous internal store. A SCADA platform speaking Kafka natively is the IT/OT boundary of Chapter 07 turning into product. Its commercial model disrupted the per-tag incumbents and is discussed in Chapter 05 as a procurement lesson. Best suited for organizations building UNS-centric architectures with strong integrator support, from single plants to multi-site standards.

AVEVA (by Schneider Electric). AVEVA carries the broadest industrial software portfolio in this zone, from System Platform and Plant SCADA to the PI historian universe covered in the applications column. Strength lies in large, multi-site standardizations, deep process-industry credibility, and the installed base. Complexity and licensing structure are the recurring counterarguments. Best suited for process industries with existing AVEVA or PI investments where portfolio depth outweighs platform simplicity.

Rockwell Automation. FactoryTalk covers SCADA, MES, and analytics deeply integrated with Rockwell's controller base. In Rockwell-standardized plants it is the path of least resistance, with tight controller integration that no third party matches. Outside them, its pull weakens, which is the platform-native pattern familiar from every other landscape in this series. Best suited for Allen-Bradley environments in North American discrete manufacturing.

Specialists include **COPA-DATA** with zenon, strong in energy and pharmaceutical compliance scenarios, **Beckhoff** with TwinCAT, PC-based control that anticipated the virtual PLC shift by two decades, and **CODESYS**, the vendor-neutral runtime whose Virtual Control products bring IEC 61131 logic onto containers. Worth watching: **N3uron**, a lean edge platform combining DataOps and web SCADA, gaining traction as a pragmatic UNS building block.

Unified Namespace (UNS) and Data Infrastructure

This zone is the backbone between plant and enterprise. The architectural pattern that defines it, the Unified Namespace, is covered in Chapter 06.

HiveMQ. HiveMQ built its business on enterprise MQTT at scale: clustering, observability, and Sparkplug awareness, proven in connected-car and manufacturing deployments, and it remains the broker most often named in UNS reference architectures. The substantial product expansion is HiveMQ Pulse, launched in 2025: model-driven governance for MQTT data, with a centralized semantic graph that defines expected topics, payload structures, and types, and distributed agents that validate payloads and compute KPIs at the edge. In September 2026 the company bundled broker, Edge, Pulse, and a new AI layer under a single HiveMQ Platform brand, positioning toward industrial AI. The packaging is new; the architecture underneath is the Pulse work. With Pulse, HiveMQ now competes in the contextualization layer alongside the DataOps vendors, part of the convergence trend in Chapter 08. Best suited for enterprises running the UNS as production infrastructure with commercial support, compliance requirements, and multi-site topologies.

EMQX. EMQX pairs an open source core with a commercial platform and pushes aggressive scale numbers, spanning MQTT brokers, the NanoMQ edge gateway, and cloud services. For organizations that want open source foundations with a commercial upgrade path, it is the natural counterpart to HiveMQ in evaluations. Best suited for teams that value the open core model and want flexibility from embedded edge to large clusters.

Amazon Web Services. AWS IoT Core remains the primitive for device connectivity into the AWS cloud after the portfolio pruning of recent years. It is solid infrastructure for cloud-side ingestion at scale. It is no longer an application strategy, and the discontinued services around it are the strongest argument in this paper for keeping the plant-side architecture vendor-neutral. Best suited as the cloud landing zone for AWS-committed enterprises, behind a neutral UNS rather than instead of one.

Microsoft. Azure IoT Operations is Microsoft's Arc-enabled, Kubernetes-based answer to the same lesson: an edge-to-cloud data plane with MQTT at the center, replacing the retired IoT Central and aging IoT Edge lines. It is the most UNS-shaped offering any hyperscaler has produced, and it is young. Best suited for Azure-committed enterprises willing to adopt a Kubernetes operating model at the edge, with the migration history priced in.

Specialists include **Solace**, whose PubSub+ event mesh handles multi-protocol, fine-grained routing in large industrial and financial environments and powers SAP's Advanced Event Mesh. **Eclipse Mosquitto** is the lightweight open source broker running in countless production sites and embedded devices. **NATS (Synadia)** is the CNCF-governed messaging system spreading from cloud-native into industrial edge scenarios; MachineMetrics, an IIoT analytics provider, runs it as the transport layer for high-frequency machine data across thousands of connected machines. Worth watching: **United Manufacturing Hub**, the open source, Kubernetes-based UNS stack, and **Zenoh (ZettaScale)**, the protocol arriving from robotics and automotive, profiled as a standard in Chapter 05.

Applications and Industrial AI

The top of the stack turns contextualized data into decisions. This is where IIoT platforms, digital twins, analytics, and copilots live, and where the business case for everything below gets proven, from OEE gains to predictive maintenance.

Siemens. Insights Hub and the Xcelerator portfolio anchor Siemens on the application side, with the digital twin business, including the NVIDIA Omniverse partnership, as the most differentiated asset. The PepsiCo results in Chapter 08 show what the closed-loop twin story looks like in production. Best suited for enterprises with Siemens automation below and simulation ambitions above, where the twin portfolio compounds.

ThingWorx (by Velotix). ThingWorx pioneered the IIoT application platform category and now restarts under Velotix ownership alongside Kepware and Proficy. The combination is coherent: connectivity, operations, and applications under one roof, focused purely on industrial software. Execution risk is the counterweight, as three product lines from three corporate parents become one company. The first combined signals are visible: ThingWorx 10.2 shipped in September 2026 with an AI assistant, an MCP server, agent services, and MQTT v5 support, while the Proficy release cadence continued on schedule. Best suited for existing ThingWorx customers, who gain a focused owner, and for buyers who want an integrated industrial stack from an independent vendor; watch the first combined roadmap before committing greenfield.

AVEVA. The PI System is widely regarded as the most deployed industrial data historian in the world, and CONNECT extends it into a cloud data platform. For many enterprises, PI data is the richest industrial dataset they own. AVEVA's challenge is turning that installed base into a modern, open data platform faster than UNS-native architectures route around it. The parent is raising the stakes: with the agreed PTC acquisition and the Cognite deal, Schneider is assembling design, industrial data, and operations under one roof, which puts AVEVA's data layer at the heart of a digital thread strategy once the deals close. Best suited for process industries where PI is already the system of record and the question is how to open it up, not whether to keep it.

Cumulocity. Independent again after the January 2025 management buyout, Cumulocity is one of the few pure-play IIoT platforms left in Europe, with device management, application enablement, and edge deployments validated at large scale, and a strong white-label footprint among machine builders and telcos. Gartner named Cumulocity a Leader in its 2025 Magic Quadrant for Global Industrial IoT Platforms, the fourth consecutive year counting the Software AG era. Independence restored focus; the question is reach against the ecosystems of the giants. Best suited for OEMs and operators that want a neutral, European, full-stack IIoT platform without hyperscaler or automation-vendor gravity.

Specialists include **Tulip**, the frontline operations platform that turned no-code apps into a manufacturing category, **XMPro**, focused on intelligent digital twins and agentic decisioning in asset-heavy industries, and **Eclipse Ditto**, the open source digital twin framework used inside several commercial platforms. Worth watching: **Tier0 by FREEZONEX**, an industrial data platform from Singapore and Hangzhou that unifies signals into a real-time namespace with an AI app builder for UNS-based MES-class applications, deployed at SUPCON for SMT factory digitalization, with an open source edge built on Node-RED, a semantic MQTT broker, and TimescaleDB. And **EnvisionSCADA by Omnicron**, applying generative AI to SCADA application engineering.

05

CHAPTER 05

Open Standards Decide Who Stays Replaceable

OPC UA and MQTT are complementary layers, Sparkplug B makes MQTT payloads interoperable across vendors, and performance, implementation quality, and licensing decide whether a standard actually serves you.

Open standards create portability at every layer of the industrial IoT stack. They determine whether hardware, brokers, and platforms can be exchanged, or whether today's convenience becomes tomorrow's lock-in. They are not free of trade-offs: performance, implementation quality, and licensing differ widely, which this chapter covers alongside the standards themselves.


OPC UA is the semantic standard of the plant. It normalizes how machines describe themselves: data types, hierarchies, information models. Vendor support is near-universal, from PLCs to SCADA to clouds. The semantics live in the companion specifications, and there are many. The VDMA counts more than 400 OPC UA companion specifications across all industries, and even in machinery fewer than two thirds are maintained in a current version. The practical entry point is OPC UA for Machinery, which covers identification, machine status, job management, and energy monitoring and applies to almost any machine without extensions. The portability check is information model fidelity: whether a vendor exposes real companion specification models, ideally certified, or a flattened proprietary view with its own variables. An OPC UA server with custom variables is connected, not interoperable.

MQTT is the distribution standard. It moves events at scale with minimal overhead and powers the Unified Namespace pattern. **Sparkplug B is an open OASIS specification that defines topic structure, payload encoding, and state management for MQTT in industrial environments, making MQTT data interoperable across vendors instead of merely transportable.** The portability check is Sparkplug conformance and clean topic architecture.

OPC UA or MQTT?

OPC UA and MQTT are complementary standards that operate at different layers: OPC UA normalizes device-level semantics at the machine, while MQTT distributes contextualized events across the plant and to the enterprise. The practical pattern of 2026 is OPC UA at the equipment interface, contextualization in the DataOps layer, and MQTT with Sparkplug B as the distribution fabric. Any vendor forcing a choice between the two is describing its product boundary, not the architecture. In practice, many plants deliberately run only one of them: OPC UA end to end where the automation stack supports it, or MQTT-first where simplicity wins. Combining both adds moving parts and pays off when contextualized distribution at scale is the goal, the pattern described in [OPC UA, MQTT, and Apache Kafka: The Trinity of Data Streaming in Industrial IoT](#). The convergence now runs at the specification level too: OPC UA Pub/Sub over MQTT carries OPC UA semantics on MQTT infrastructure, so the two standards stopped being rivals where they meet.

OPC UA and MQTT: Complementary Layers



A forced choice between OPC UA and MQTT describes a product boundary, not the architecture.

Analysis: Kai Waehner · kai-waehner.de

Figure 03. OPC UA and MQTT as complementary layers: device semantics at the machine, contextualization in the DataOps layer, distribution through MQTT and Sparkplug B.

Standards Beyond the Protocols

ISA-95 contributes the naming hierarchy that gives UNS topic trees their structure: enterprise, site, area, line, cell. **IEC 61499**, with Eclipse 4diac as its open source implementation, standardizes distributed control logic and matters as virtual PLCs make control portable. A newer front is interoperability above the data layer: Margo, a Linux Foundation initiative founded by ABB, Microsoft, Rockwell, Schneider Electric, Siemens and Capgemini and now carried by some forty vendors, standardizes how edge applications, devices, and orchestration software interoperate across multi-vendor environments. Preview releases have shipped since January 2026 with general availability targeted for late 2026; it is the standard to watch for the edge application layer.

Regulation and Licensing Decide Portability Too

Since September 12, 2026, the access-by-design obligation of the EU Data Act applies to connected products placed on the EU market. Product data and the metadata needed to interpret it must be accessible to the user by default, in a structured, commonly used, machine-readable format. For machine builders, a companion specification is the most direct way to meet the metadata requirement without inventing a format. For operators, open data access on new equipment moves from a procurement wish to a legal right.

Regulation reaches beyond data access. The Cyber Resilience Act makes machine builders document which software and firmware runs in every product they ship, and that inventory comes from PLC projects and machine networks. Digital product passports, starting with batteries in 2027, require product and process data straight from production. This is not a European peculiarity. China steers sustainability through state policy and green finance, requires sustainability reports from its largest listed companies, and already tracks every

electric vehicle battery with a digital identity on a national platform. The architectures differ: Europe keeps passport data with the manufacturer, China centralizes it on a state platform. Either way, plant data becomes regulated output.

Lock-in does not only come from protocols. **The EU Data Act**, applicable since September 2025, gives users of connected products a legal right to access and share the data those products generate, which turns data portability from an architecture preference into a compliance requirement in Europe. Licensing is the other lever.

Per-tag licensing punishes exactly the data growth that industrial IoT creates. Ignition's unlimited-tag model disrupted the category for that reason, and unlimited or capacity-based licensing has become a procurement requirement in modern evaluations. Test both before signing: whether the data leaves in open formats, and whether the license survives scaling. Standards carry their own bill: on Siemens controllers, OPC UA is a per-CPU runtime license with hard node limits per controller class, and a standards-based read can cost up to two orders of magnitude more in latency and controller load than the same controller's native protocol on the same wire. How a client asks often matters more than which protocol it speaks: in [cross-vendor benchmarks](#) spanning Siemens, Beckhoff, Schneider Electric, Rockwell, and Mitsubishi controllers, the same data asked for well versus badly differed by one to two orders of magnitude, a larger gap than between any two protocols. Batched reads and held-open connections are worth more than most protocol migrations. [Standards buy portability, semantics, and auditable security, not performance](#). Budget and measure them like any other architecture component.

Zenoh and NATS: Protocols to Watch

Two protocols are entering industrial IoT from adjacent worlds. **Zenoh**, driven by ZettaScale under Eclipse governance, became the first non-DDS middleware option for ROS 2, is used in autonomous vehicle programs including the Indy Autonomous Challenge and V2X projects, and carries an ITU recommendation for intelligent transport systems. Its strengths are very low latency, peer-to-peer topologies, and constrained networks.

NATS, CNCF-governed with Synadia behind it, brings lightweight cloud-native messaging to the edge, with production IIoT deployments like MachineMetrics. Neither displaces MQTT in the factory today. MQTT plus Sparkplug B has the broker ecosystem, the tooling, and the installed base. But robots, AMRs, and software-defined vehicles, the carriers of what analysts now call physical AI, are entering plants, and they bring their protocols with them. McKinsey expects physical AI to arrive first in manufacturing and logistics, where the economics are clearest, which makes this protocol boundary a near-term planning item rather than a research topic. Watch the boundary.

06

CHAPTER 06

The Industrial IoT Stack and the Unified Namespace

How the categories compose into one architecture, and why the UNS replaced point-to-point integration as the dominant pattern.

What Is a Unified Namespace?

A Unified Namespace (UNS) is an event-driven architecture pattern in which all plant data is published into a centrally structured, hierarchical namespace, typically an MQTT broker organized along ISA-95, so that every application consumes current, contextualized data from one place instead of maintaining point-to-point connections. Inside the plant it is the dominant pattern; across the enterprise it is one building block next to data streaming, APIs, and governance, which Chapter 07 maps.

One clarification matters: raw sensor values dumped onto a broker are not a UNS. Without contextualization in the DataOps layer, a broker is a message bus, not a namespace. The namespace earns the name when data arrives modeled, named, and self-describing. The maturity step behind that clarification is data contracts: without machine-readable schemas, a namespace cannot guarantee structure, validate payloads, or evolve safely across teams, sites, and vendors.

The Industrial IoT Stack




Figure 04. The Industrial IoT Stack. Five layers from physical operations to intelligence and agents, with the Unified Namespace as the central distribution layer.

Figure 4 shows how the four landscape categories compose into one architecture. Physical operations at the bottom: PLCs and their virtual successors, sensors, robots. Connectivity and DataOps above it: OPC UA, PLC4X, contextualization, OT data catalogs. The Unified Namespace in the center. Applications and digital twins above it: SCADA, MES, historians, simulation. Intelligence and agents at the top: copilots, analytics, agentic AI.




Figure 05. Point-to-point OT integration versus the Unified Namespace: coupled, stale, and brittle versus current, contextualized, and decoupled.

The contrast in Figure 5 is the architectural argument of this paper. Point-to-point OT integration couples every application to every source. Each new consumer multiplies connections, each change breaks something downstream, and every system holds a different version of the plant's state. The UNS inverts this. **Producers publish once. Consumers subscribe independently.** Adding the tenth application costs the same as adding the second. This is the same architectural shift that event streaming brought to enterprise IT, applied to the plant, which is exactly why the two worlds connect so cleanly at the boundary described next.

07


CHAPTER 07

Where Enterprise IT Connects

Each OT layer has a natural counterpart on the IT side. Four exchanges cross the IT/OT boundary, physical operations deliberately has no direct one, and data streaming with Apache Kafka sits adjacent to this landscape rather than inside it.

Industrial IoT does not end at the plant gate, and it does not extend into the enterprise data stack either. The two architectures meet at a boundary, and each OT layer has a natural counterpart on the IT side.

Mapping the Stack to Enterprise IT



Physical operations reach enterprise IT only through the layers above. No direct integration by design.

Enterprise IT categories are covered in the Data Streaming and Data Integration Landscapes. Analysis: Kai Waehner · kai-waehner.de

Figure 06. Mapping the Industrial IoT Stack to enterprise IT. Four exchanges across the IT/OT boundary; physical operations connects only through the layers above.

The intelligence layer exchanges models and training data with the lakehouse and analytics platforms. The applications layer exchanges orders and transactions with business systems, the classic MES and ERP integration. The Unified Namespace exchanges real-time events with data streaming, where Apache Kafka and Apache Flink carry plant events into the enterprise; this is the main data highway, and the reason Kafka appears in this paper as an adjacent technology rather than a landscape entry. The boundary is turning into product on both sides: most modern platforms now integrate with Kafka natively, Ignition 8.3 for instance publishes and consumes Kafka topics through its Event Streams module, and the analytical platforms accept the Kafka protocol for ingestion. On the IT side of that exchange the pattern has a name: the UNS structures live OT data, and [data products](#) govern and distribute it across the enterprise with schemas, ownership, and access control. The combination underneath is not new: OPC UA, MQTT, and Kafka as the trinity of industrial data streaming was the argument in 2022, and it is the default pattern of 2026. And the connectivity and DataOps layer exchanges metadata and lineage with enterprise data catalogs and governance platforms.

One nuance on that last exchange: enterprise catalogs govern the namespace, but the metadata that feeds them is produced in the DataOps layer. No MQTT broker integrates with Collibra. The asset models, tags, and lineage come from the contextualization tools, which is why the OT data catalog trend in Chapter 08 starts there. The same logic explains another missing arrow: direct OT-to-lakehouse ingestion, promoted by middleware and lakehouse vendors alike, recreates reverse ETL at plant scale; curated data reaches the lakehouse through the namespace and the streaming layer, not around them.

A second nuance on the word "act". When data platforms say they act on operational data, they mean dashboards, workflow triggers, and recommendations. Closing the loop into control logic remains the job of the automation stack: the PLCs, virtual PLCs, and SCADA systems with their safety and engineering discipline.

The separator is determinism and safety certification, not speed: an orchestration or analytics engine retries and escalates when something runs late, while a safety PLC has no such option. What decides and coordinates above the data movement is workflow orchestration, a layer deliberately outside this landscape and covered in [Unified Orchestration in Manufacturing](#). Physical operations has no arrow in Figure 6 by design. PLCs reach enterprise IT only through the layers above them. Any architecture that connects controllers directly to enterprise applications has skipped the layers that make the connection safe, governed, and maintainable.

The enterprise side of this boundary is covered in the [Data Integration Landscape](#), whose Industrial IoT chapter describes the same boundary from the other direction: the IIoT layer captures and bridges, the integration layer carries and connects. The event streaming layer itself is analyzed in the [Data Streaming Landscape](#).

08

CHAPTER 08

Five Trends Shaping Industrial IoT

Virtual PLCs, closed-loop twins, generative and agentic AI, OT metadata governance, and the convergence of the categories.

Virtual PLCs Move Control to Software

A virtual PLC runs deterministic control logic as software on standard IT hardware instead of dedicated controller hardware. Software-defined automation moves control from the cabinet to industrial PCs and servers. Audi's EC4P program is the flagship: thousands of traditional PLCs, IPCs, and HMIs replaced by virtual counterparts on rack servers, built on Siemens virtual controllers with standard networking and virtualization underneath. Analysts project the virtual controller market to reach several billion dollars by 2035, and ABI Research published its first Software-Defined Automation competitive ranking in September 2026, naming Schneider Electric, Siemens, Bosch Rexroth, SUPCON, Honeywell, and CODESYS as leaders. Safety has caught up as well: CODESYS Virtual Safe Control is the first virtual safety controller certified to IEC 61508 SIL3 with FSoE. The consequence for this landscape: control becomes deployable, versionable, and orchestrable like software, and the boundary between the automation vendor and the IT platform starts to blur.

Digital Twins Close the Loop

Digital twins are moving from visualization to execution. A public flagship example: PepsiCo, working with Siemens Digital Twin Composer and NVIDIA Omniverse, recreates lines with physics-level accuracy and reports identifying up to 90 percent of potential issues before physical changes, a 20 percent throughput increase, and 10 to 15 percent capital expenditure reduction. The what-if layer is becoming an expectation. The next step, twins that trigger actions back through the orchestration and control layers, is where the top of the stack meets the bottom.

The interoperability standard behind industrial twins is the Asset Administration Shell (IEC 63278), driven by the IDTA, used in data spaces such as Catena-X, and piloted in Manufacturing-X projects such as Factory-X. Twins that expose an AAS remain exchangeable across vendors. The semantics underneath are converging as well: the VDMA is moving its companion specification semantics into the IEC Common Data Dictionary, a technology-agnostic dictionary that both OPC UA and AAS models can reference. Product passports show where this lands in practice: AAS-based demonstrators generate the passport from the production event itself, triggered by an MQTT message or a REST call when a unit passes its end-of-line test, so the passport becomes part of the production process rather than documentation after the fact.

Generative and Agentic AI Reach the Plant

Industrial software vendors have ended the free copilot phase and are monetizing generative AI, from configuration generation to natural-language analytics. The agentic step is next: Gartner projects that by 2030, semiautonomous AI agents will orchestrate 10 percent of key production, quality, and maintenance use cases, up from around 2 percent today, with humans keeping final approval. The plumbing for this is being built now: HighByte ships MCP tools that let agents discover and read OPC UA and MQTT data through natural language, and MCP servers are spreading across industrial vendors. The market has moved past pilots. IoT Analytics declared the pilot era over after Hannover Messe 2026, and Gartner renamed its Magic Quadrant for Industrial IoT Platforms into one for Industrial AIoT Platforms in September 2026. Not every plant decision is an LLM call, though. Deterministic decision models, the category TypeSafe AI calls System One models with Jev as its example, execute explainable microsecond decisions at the edge where LLM latency and non-determinism are disqualifying. They complement LLM-based agents rather than replace them. Both kinds

of agent need exactly what this landscape describes: a contextualized namespace to read from and governed pathways to act through. An AI agent without both is fast, capable, and ungoverned, and plants without that foundation will run agent pilots that never leave the lab. The trust and governance side of this shift is covered in the [Trusted Agentic AI Landscape](#).

OT Metadata Joins Enterprise Governance

An OT data catalog inventories and governs industrial assets, tags, and lineage at the plant level and connects that metadata to enterprise data governance platforms. It is the newest category in this landscape. Litmus is first to market with planned integrations into Atlan, Collibra, Purview, Dataplex, and Unity Catalog, and HighByte's modeling capabilities point the same direction. The significance is organizational as much as technical: for the first time, the plant's assets, tags, and lineage become visible to the same governance programs that manage enterprise data. Lineage is the sharpest difference between the two worlds: in OT, transformations happen at the edge before data ever reaches a warehouse, so the trail must start at the asset, not at ingestion. Regulation adds a second driver. Machine-readable security advisories in the CSAF format only help when the plant knows which device runs which firmware in which cabinet. Asset and firmware inventory becomes the shared basis for CRA conformity, NIS2 evidence, and obsolescence management. Industrial data stops being a silo with its own rules.

The Categories Are Converging

Every column of this landscape is growing toward its neighbors, and the evidence is shipped product, not positioning. HiveMQ Pulse moved a broker vendor into contextualization with a semantic graph and distributed validation. The Litmus Data Catalog moved a DataOps vendor into governance. Ignition's Event Streams moved a SCADA platform into native event distribution with Kafka. HighByte's MCP services moved contextualization into the agent interface. The categorical boundaries that structure this landscape are real today and softer every quarter, which is exactly why vendors are placed by center of gravity rather than by claimed scope. For buyers, the practical consequence: evaluate what a platform does well in production, not the number of boxes its diagram covers, and expect overlap in every shortlist. Independent analysis points the same way: IoT Analytics observes that value increasingly concentrates in platforms that orchestrate data and workflows across sites.

09

CHAPTER 09

Six Questions Before Any Vendor Shortlist

Six questions determine which parts of this landscape matter for your situation. Apply them before any vendor shortlist.

Six questions determine which parts of this landscape matter for your situation. Apply them before any vendor shortlist.

Q1. Where does your data need to be actionable first?

If the primary consumers are operators and control loops, start in the SCADA and edge platform column and build up. If the primary consumers are enterprise analytics and AI, start with connectivity and the UNS and let applications follow. Most failed IIoT programs started at the top of the stack without the bottom.

Q2. One vendor stack or best-of-breed through a UNS?

A single-vendor stack is legitimate when the fleet, the team, and the roadmap already belong to that ecosystem. It buys speed and accountability at the price of following the vendor's conventions. A best-of-breed architecture connected through a UNS keeps every layer replaceable. It requires more architectural ownership. A Unified Namespace, together with open interfaces toward the enterprise, is what makes the best-of-breed path viable at all.

Q3. Which standards are non-negotiable for your equipment fleet?

Inventory before shortlisting: which protocols does the installed base speak, which information models exist, which OPC UA companion specifications apply and whether devices are certified against them, and is Sparkplug B conformance required. A platform that covers 95 percent of the fleet natively beats a more elegant one that needs custom drivers for your most critical line.

Q4. What are your deployment constraints?

Air-gapped sites, intermittent connectivity, data residency, and edge compute footprints eliminate more vendors than feature lists do. Verify edge, on-premises, hybrid, and cloud options early. The hyperscaler retirement wave is a reminder that deployment strategy should not depend on one provider's portfolio decisions. Virtual PLCs change this calculus as well: once control runs as software, deployment constraints become infrastructure decisions rather than hardware facts. Two checks apply on every branch, covered in depth in [Edge to Cloud and Back](#): the edge dials out and never accepts inbound connections into the site, and local operation must survive offline with a store-and-forward buffer sized deliberately rather than discovered in the first outage.

Q5. Does the licensing model survive scaling?

Industrial IoT multiplies tags, topics, and consumers by design. Per-tag and per-connection licensing turns architectural success into a budget problem, and standards can carry runtime licenses of their own, such as OPC UA per CPU on Siemens controllers. Model the license at ten times today's data points before signing, and scale it geographically too: multi-region rollouts meet different legal regimes and price lists, and open source and open standards are the levers that keep those negotiations balanced.

Q6. Is closed-loop automation or agentic AI a near-term priority?

If yes, data freshness and governance become selection criteria now, not later. Agents and closed-loop twins need current, contextualized data and auditable pathways to act. That pulls the UNS, the DataOps layer, and the enterprise boundary from Chapter 07 into the first project, not the roadmap. Even if the answer is no, plan the direction: most plants are brownfield, and retrofit, modernizing existing equipment with new connectivity and software while keeping the installed base, is how this foundation gets built step by step.

A NOTE FROM THE AUTHOR

The architecture is the strategy. Industrial connectivity, the Unified Namespace, operational platforms, and industrial AI are layers of one converged architecture, and decisions in each layer constrain what is possible in the others. A plant with excellent SCADA but point-to-point integration will starve its AI initiatives. A namespace without contextualization is a message bus, not a foundation. An AI agent without governed pathways is fast, capable, and ungoverned.

The vendor landscape reflects a market rebuilding itself: hyperscalers retreating to primitives, private equity assembling new platforms, open source carrying more of the foundation every year, and standards deciding who stays replaceable.

The plant floor is becoming software. Data must flow from edge to cloud in real time. Everything else follows from that.

ABOUT THE AUTHOR

Kai Waehner

ADVISORY FIELD CTO

Kai Waehner is an Advisory Field CTO who has spent over 20 years helping enterprises make their most consequential data, process, and AI architecture decisions. He works with Fortune 500 and Global 2000 companies across Europe, North America, the Middle East, Asia, and Australia, and follows a deliberately vendor-neutral approach in his advisory work that prioritizes the right architectural choice over the easiest sell.

His effectiveness as an advisor comes from range. He moves fluidly between a strategic conversation with a CIO and a deep architecture review with an engineering team, and across more than a dozen industries, from financial services and manufacturing to telecom, retail, and the public sector. That range is grounded in 100+ speaking engagements, from technical conferences like AWS re:Invent and QCon to CIO and CTO executive summits.

Kai is known for his independent technology landscapes for data streaming, data integration, process intelligence, SCADA, and trusted agentic AI. He also writes the blog at kai-waehner.de, covering industry use cases, technical best practices, and emerging topics for enterprise architects, CTOs, CDOs, and data engineers. Kai is available for advisory engagements, workshops, and keynotes worldwide, as well as media collaborations.

ABOUT THIS LANDSCAPE

This landscape is an independent analyst perspective, not a quantitative ranking. Vendor selection, tier placement, and market footprint reflect the author's assessment based on public information, vendor announcements, and two decades of enterprise architecture work with the platforms covered. No vendor paid for inclusion or placement, and no vendor reviewed or approved its section before publication. Revenue and customer figures are drawn from public vendor statements where available. The author runs an advisory practice through Kai Waehner GmbH and has held roles at Talend, TIBCO, and Confluent. He also serves as Global Field CTO at Kestra, a workflow orchestration vendor outside this landscape's scope. This landscape was produced through Kai Waehner GmbH, independently of any vendor engagement.

WEBSITE

kai-waehner.de

LINKEDIN

linkedin.com/in/kaiwaehner

NEWSLETTER

[Subscribe for data streaming, industrial IoT, and agentic AI insights](#)

KEEP READING · STAY INFORMED

The plant floor is becoming **software**. Data must flow from edge to cloud in real time. Everything else follows from that.

BLOG & NEWSLETTER

Regular updates from the industry

Subscribe for the latest thinking on enterprise architecture, data streaming, industrial IoT, and trusted agentic AI.

kai-waehner.de/news

LANDSCAPES

All Kai Waehner Landscapes

Data streaming, data integration, process intelligence, trusted agentic AI, and industrial IoT, each as an independent vendor map.

[Explore the landscapes](#)

CONNECT

LinkedIn & X

Follow along for ongoing analysis of industrial IoT, data streaming, and the enterprise software landscape.

[LinkedIn](#) | [X \(Twitter\)](#)

MORE READING

The Trinity of Modern Data Architecture

How data integration, process intelligence, and agentic AI fit together.

[Read more](#)