



WHITE PAPER

The Industrial Edge Data Plane

A scalable architecture for OT data, context, and intelligence

What this paper covers

01	Executive Summary	03
	The edge as a deliberate architectural layer, not a gateway	
02	The Industrial Data Problem	12
	Why connectivity alone does not create usable information	
03	The Hidden Economics of Sensor-to-Cloud	19
	How industrial data volumes multiply into infrastructure cost	
04	Reference Architecture	25
	Five logical layers, one source, many purpose-driven data products	

Executive Summary

The purpose of the edge layer is not simply to move data. It is to transform operational signals into governed, contextualized, consumable industrial information before it reaches the cloud.

Industrial organizations are generating operational data at an unprecedented scale. Modern manufacturing facilities contain thousands—and increasingly tens of thousands—of sensors, instruments, controllers, drives, machines, analyzers, and software applications continuously producing data about the state and performance of operations. At the same time, cloud computing, enterprise analytics, digital transformation programs, and artificial intelligence are creating an expanding capability to leverage this operational technology (OT) data to drive business improvements.

A common architectural response has been to connect industrial data sources directly, or through relatively simple gateways, to cloud platforms where data is ingested, stored, contextualized, transformed, and ultimately consumed. Cloud data architectures such as lakehouses and medallion models provide powerful capabilities for enterprise data management and analytics. However, applying these patterns to industrial data without considering the unique characteristics of OT can result in unnecessarily large volumes of raw data being transported, stored, and repeatedly processed in the cloud leading to significant costs.

The key consideration emerging is what data should be sent to the cloud, in what form, at what frequency, and after what processing and contextualization have occurred?

This paper proposes an Industrial Edge Data Plane as a deliberate architectural layer between industrial automation systems and enterprise or cloud platforms. Rather than treating the edge as a collection of simple protocol gateways—or deploying an individual edge device for every sensor or machine—the Industrial Edge Data Plane provides a small number of strategically located, managed computing nodes within a manufacturing facility.

Depending on the scale, network architecture, criticality, and segmentation of a site, one or several edge nodes can service hundreds or thousands of underlying industrial data sources.




Figure 1 — The edge data plane sits between plant systems and enterprise consumers.

The Industrial Edge Data Plane performs a progression of functions:




Figure 2 — The core data progression, underpinned by continuous governance.

Raw signals can be collected from PLCs, DCSs, SCADA systems, historians, smart instruments, drives, analyzers, MES/MOM applications, quality systems, maintenance platforms, and other industrial sources using established OT protocols and interfaces. At the edge, these signals can be normalized into common structures, associated with equipment and process context, filtered according to their intended use, enriched with production information, analyzed locally where appropriate, and routed to one or multiple downstream consumers.

This represents an important shift from a sensor-to-cloud architecture toward a source-to-context-to-consumer architecture.

Context before consumption

Industrial data derives much of its value from context. A raw value such as 43.7 associated with a controller tag such as FIC101.PV may be meaningful to the controls engineer who designed the system, but it has limited value to an enterprise data scientist, maintenance application, AI model, or cross-site analytics platform. The same observation becomes considerably more useful when it is associated with its industrial meaning:

Measurement → **Instrument** → **Equipment** → **Process Cell** → **Production Phase** →
Batch → **Product**

Creating this context at or near the source provides an opportunity to establish consistent industrial information models before data is distributed throughout the enterprise.

Standards and semantic models such as ISA-95, ISA-88, OPC UA information models, and other industry-specific models can provide the foundation for this contextualization. The result is that downstream systems receive industrial data products rather than collections of unexplained tags.

Moving the cloud boundary

The Industrial Edge Data Plane also changes the economics of industrial data. Consider a facility containing thousands of continuously changing process variables sampled every second or faster. Sending every observation to a cloud platform can generate billions of records annually.

Across multiple facilities, the volume can become enormous. The organization may then incur costs for network transmission, ingestion, storage, transformation, compute, replication, retention, and governance—often before determining whether the majority of those observations have meaningful enterprise value.

Cloud medallion architectures commonly organize data through Bronze, Silver, and Gold layers. In a conventional implementation, raw industrial data may first enter the cloud Bronze layer before being cleaned, contextualized, aggregated, and transformed.

The Industrial Edge Data Plane challenges the assumption that the cloud must always be the first location where those activities occur. For many industrial use cases, deadband filtering, change-of-state processing, aggregation, contextualization, event detection, compression, and other transformations can occur closer to the source.

High-frequency raw data can remain locally available when operational or forensic requirements justify it, while contextualized events, summaries, exceptions, production information, and selected time-series data are transmitted to enterprise platforms.

The objective is not simply to minimize data volume. It is to move the right data, with the right context, to the right consumer, at the right time.

This approach can significantly reduce unnecessary cloud ingestion and processing while simultaneously improving the quality and usability of the data that does reach enterprise platforms.

The edge as an OT computing platform

The economic case for industrial edge infrastructure becomes even stronger when the edge is viewed as more than a data gateway. The same managed computing infrastructure used to contextualize and route industrial data can support OT-specific applications requiring proximity to machines and processes. These may include condition-based maintenance, anomaly detection, predictive analytics, machine health monitoring, energy optimization, OEE calculations, quality monitoring, protocol conversion, local visualization, and AI/ML inference.

Some of these applications have operational requirements that make exclusive dependence on cloud infrastructure undesirable. Network latency, intermittent connectivity, cybersecurity boundaries, high-frequency data requirements, or the need for continued operation during loss of enterprise connectivity may all favor local execution. The edge therefore becomes a managed OT application and data platform rather than simply another gateway.

This creates a second source of economic value. Edge infrastructure can both reduce the cost of moving and processing unnecessary data in the cloud and enable operational applications that create measurable value within the plant.

A governed boundary between OT and IT

The Industrial Edge Data Plane also establishes an important organizational boundary. OT systems contain specialized information about equipment, processes, production states, control strategies, and operating conditions. The meaning of this data is frequently best understood by the engineers, operators, maintenance personnel, and manufacturing specialists responsible for those systems.

Enterprise IT organizations, conversely, provide the platforms and disciplines necessary to manage information at enterprise scale. Successful industrial data architectures require both.

The edge provides a logical point at which OT can maintain responsibility for source connectivity, industrial semantics, operational context, data quality, and appropriate collection policies while IT provides enterprise identity, cybersecurity governance, cloud infrastructure, analytics platforms, and broader data governance.

OT RETAINS	IT PROVIDES
Source connectivity	Enterprise identity
Industrial semantics	Cybersecurity governance
Operational context	Cloud infrastructure
Data quality	Analytics platforms
Collection policy	Data governance

OT should participate in governing the meaning and quality of operational data before that data becomes an enterprise asset.

A reference technology stack

The architecture described in this paper is technology-independent, but it can be implemented today using commercially available industrial platforms.

A representative implementation uses HighByte Intelligence Hub as the DataOps and contextualization layer for acquiring, modeling, transforming, and routing industrial data. ZEDEDATA provides centralized management, security, orchestration, and lifecycle management for the underlying distributed edge computing infrastructure. Sorba.ai provides edge-capable analytics and machine learning that allow selected industrial intelligence workloads to execute close to the processes they monitor.




Figure 3 — The same three responsibilities deployed across multiple sites, each delivering curated data to shared cloud platforms.

Keeping these capabilities logically separated is important. It allows organizations to evolve individual components without tying the industrial data architecture to a single cloud provider, analytics platform, application vendor, or proprietary industrial ecosystem.

This paper does not advocate moving enterprise analytics away from the cloud or attempting to recreate hyperscale computing infrastructure inside manufacturing plants. Cloud platforms remain exceptionally well suited to enterprise analytics, cross-site comparison, long-term data management, large-scale AI model training, business intelligence, collaboration, and integration of operational data with financial, supply-chain, commercial, and other enterprise information.

The Industrial Edge Data Plane complements these capabilities. It establishes an architectural decision point between the physical operation and the enterprise data ecosystem—determining what should remain local, what should be processed locally, what should be transmitted, and what context should accompany it. The resulting architecture allows organizations to use cloud computing more deliberately rather than simply using more of it.

Ultimately, the future of industrial data architecture should not be defined by how much data an organization can move to the cloud. It should be defined by how effectively that organization can transform operational data into useful information and deliver that information wherever it creates the greatest value.

The Industrial Data Problem

Industrial operations have never lacked data. What has changed is the scale at which organizations now want to consume it.

PLCs, DCSs, SCADA systems, historians, drives, smart instruments, analyzers, machine controllers, MES/MOM systems, quality applications, and maintenance systems continuously generate information about the state and performance of manufacturing processes. What has changed is the scale at which organizations now want to consume this information and the number of enterprise applications seeking access to it.

Cloud analytics, enterprise data platforms, artificial intelligence, predictive maintenance, digital twins, and cross-site performance initiatives are creating demand for OT data well beyond the control systems and historians for which it was originally generated. This creates a fundamental architectural challenge: how does an organization transform large volumes of highly specialized operational data into trusted, contextualized information that can be efficiently consumed across the enterprise?

OT data is growing faster than its usefulness

A modern manufacturing facility may contain tens of thousands of continuously changing process variables. A single production asset can generate hundreds or thousands of measurements related to process conditions, machine states, energy consumption, alarms, quality parameters, and equipment health. When these signals are sampled at one-second, sub-second, or event-driven frequencies, even a single facility can generate billions of observations annually.

Connecting each of these sources is relatively straightforward until you consider how you will maintain all these connections in the aggregate. Additionally, creating useful information from them is considerably more difficult.

This distinction becomes increasingly important as industrial organizations pursue enterprise-scale data initiatives. Simply making more OT data available does not necessarily create more business value. Without deliberate decisions about what data is required, how frequently it should be collected, how long it should be retained, and which consumers require it, organizations risk creating enormous collections of operational data with limited practical utility.

The objective should not be to collect all available data, but to make the appropriate data available for its intended purpose.

Connectivity does not create context

Industrial data is inherently contextual. Consider the value of 43.7. By itself, the value communicates almost nothing. Associating it with the control-system tag FIC101.PV provides additional information, but that identifier may only be meaningful to engineers familiar with the source automation system.

THE SAME VALUE, IN CONTEXT	
Value	43.7
Measurement	Flow measurement
Instrument	FIC-101
Equipment	Transfer Pump System
Process cell	Process Cell A
Phase	Production Phase, part of a Batch

This context is essential for enterprise analytics, AI, condition-based maintenance, performance management, and cross-site comparison. Historically, much of this knowledge has been embedded implicitly within control-system tag structures, historian configurations, SCADA applications, engineering documentation, or the experience of plant personnel.

When raw OT data is exported to enterprise platforms, that context must often be reconstructed downstream. The result can be duplicated engineering effort, inconsistent interpretations of the same equipment or process, and data pipelines that are tightly coupled to individual source systems.

A scalable industrial architecture must therefore address not only how data is connected, but also where its meaning is established and governed.

The sensor-to-cloud assumption

The rapid adoption of cloud platforms has introduced a seemingly straightforward architectural model.




Figure 4 — The conventional sensor-to-cloud model.

This model has been championed by corporate IT groups who typically control the cloud platforms and have access to vendor analytical tools offered by these hyper-scaler vendors. This architecture can also encourage an assumption that raw industrial data should first be transported to the cloud and decisions about its usefulness can be made later.

Industrial data presents a different challenge than traditional enterprise transactional data. It is frequently high-volume, continuously generated, time-series oriented, and highly dependent upon multiple dimensions of process context. A large percentage of that data may have immediate operational value but limited long-term enterprise value.

For example, high-frequency vibration, motor current, temperature, pressure, or machine-state information may be essential for local condition monitoring or anomaly detection. The enterprise, however, may only require equipment health indicators, detected anomalies, production events, aggregated performance measures, or selected historical observations.

Sending every underlying observation to an enterprise platform simply because it is technically possible can therefore create significant infrastructure without creating proportional business value.

What information does each consumer truly require, and where should raw operational data be transformed into that information?

Industrial data has multiple consumers

Another limitation of a simple sensor-to-cloud model is the assumption that the cloud is the primary destination for industrial data. In practice, the same operational information may serve many consumers:

Operators and production supervisors	Enterprise historians and data platforms
Maintenance and reliability applications	Cloud analytics and business intelligence
Crewing and labor/skills systems	AI and machine-learning applications
Material usage and quality systems	Energy and sustainability applications

Some of these consumers are local to the manufacturing operation. Others operate at the site, business-unit, or enterprise level. The architecture must therefore support one source serving multiple consumers without requiring every consumer to independently rediscover, connect to, and interpret that source.

This creates the need for an intermediary data layer capable of acquiring industrial information once, establishing its context, and delivering fit-for-purpose representations to multiple destinations.

The architectural problem

The industrial data challenge is ultimately not a connectivity problem. Modern technologies provide numerous ways to connect an industrial asset to a cloud platform. The more challenging questions require governance:

What should be collected?

What does it mean, and who governs that meaning?

Where should it be processed?

How frequently should it move, and how long retained?

Which consumers need raw data, and which need context?

Without an architectural layer that addresses these questions, organizations risk building increasingly complex collections of point-to-point integrations and cloud ingestion pipelines. InflexionPoint proposes and promotes an alternative, establishing an Industrial Edge Data Plane between operational systems and enterprise consumers—a layer specifically responsible for transforming raw industrial signals into governed, contextualized, and purpose-driven information before determining where that information should go.

The Hidden Economics of Sensor-to-Cloud Architectures

The organization is not simply paying to store data. It is creating a chain of infrastructure and processing costs that multiply as data moves through the architecture.

Cloud platforms have transformed the economics of enterprise computing by allowing organizations to consume storage, processing, analytics, and artificial intelligence as scalable services. For many business applications, this model provides tremendous flexibility and value.

Industrial data, however, presents a different economic profile. OT systems can produce very large quantities of continuously generated, high-frequency data. When an architecture assumes that this data should first be transported to the cloud and evaluated later, the organization is not simply paying to store data.

The industrial data multiplication effect

Consider a hypothetical manufacturing facility with 10,000 operational data points collected once per second. That represents 864 million observations per day, or approximately 315 billion observations per year. Across ten comparable manufacturing sites, the architecture could be handling more than 3 trillion observations annually.




Figure 5 — Increases compound rather than add, growing downstream cost exponentially.

A condition-monitoring application sampling selected measurements at 100 milliseconds produces ten times the observations of the same signals sampled once per second. At enterprise scale, seemingly small architectural decisions about collection frequency and data granularity can have substantial economic consequences.

The cost is not simply storage


The cost of a sensor-to-cloud architecture should not be evaluated solely by the price of cloud object storage. Before industrial data becomes useful information, it may pass through multiple services:

Network transport	Cloud ingestion	Raw storage	Processing
Transformation	Contextualization	Analytical storage	Consumption

Each stage potentially introduces infrastructure, licensing, compute, engineering, and operational costs. A raw OT observation may therefore be handled multiple times. It may first be transmitted across the plant and enterprise network, ingested through a cloud service, written to raw storage, read by a transformation process, written to another data structure, joined with contextual information, aggregated into another representation, and finally queried by an analytics or AI application.

The medallion architecture and industrial data

The Bronze-Silver-Gold, or “medallion,” architecture has become a common pattern for modern enterprise data platforms. Conceptually, the pattern is straightforward:



If every high-frequency OT observation must first enter a cloud Bronze layer before filtering, normalization, aggregation, or contextualization occurs, the organization may incur cloud infrastructure costs for large quantities of data that have little enduring enterprise value.

Consider a motor monitored for condition-based maintenance. The local analytical application may require high-frequency motor current, vibration, temperature, and speed data to detect abnormal behavior; however, an enterprise reliability platform may not require every underlying observation. These are not the same requirement and should not be treated equally.

Moving the Bronze boundary toward the edge

One of the tenets of InflexionPoint’s Industrial Edge Data Plane creates an opportunity to move selected data-management functions closer to their operational sources.

TRADITIONAL — SENSOR TO CLOUD			
OT source	Raw cloud ingestion	Transform and add context in the cloud	Consume

INDUSTRIAL EDGE DATA PLANE			
OT source	Edge acquisition, filtering, context	Cloud Bronze, then Silver and Gold	Consume

Figure 6 — The edge does not eliminate the cloud Bronze layer. It changes what qualifies as appropriate Bronze data.

The edge can perform functions such as deadband filtering, change-of-state detection, aggregation, resampling, normalization, contextualization, compression, event detection, and store-and-forward before data is transmitted upstream.

For some use cases, raw high-resolution data should absolutely be transmitted and retained. Regulatory requirements, process investigations, advanced analytics, scientific applications, AI model development, or forensic analysis may justify retaining significant volumes of source data.

For other use cases, the enterprise may require only contextualized events, production summaries, exceptions, statistical features, or lower-frequency time-series data. The architecture therefore allows different data products from the same source to have different destinations, frequencies, resolutions, and retention policies.

The objective of this edge architecture is to reduce data volume and, with correct governance and discipline, promote purpose-driven data selection.

One-second values	Required by an optimization application
One-minute average	Sufficient for enterprise performance reporting
Change-of-state event	Sufficient for calculating equipment utilization
Spectral features	Transmitted upstream while raw vibration stays local

The Industrial Edge Data Plane and Reference Architecture

The Industrial Edge Data Plane is an OT data and application platform, not another gateway.

It is a managed computing and data-services layer positioned close to industrial operations. Its purpose is to acquire data from OT systems, establish its operational context, apply appropriate processing and governance, support selected local applications, and deliver fit-for-purpose information to enterprise and cloud consumers.

This represents an important distinction from both traditional industrial gateways and cloud-centric data acquisition architectures. A basic gateway primarily provides connectivity between systems or protocols. The Industrial Edge Data Plane provides a broader set of persistent data and computing services that become part of the plant's operational architecture.

Conceptually, its responsibilities follow a progression:

OT source Connect to equipment and systems	Edge acquisition Collect and manage at the edge	Filter and context Process, enrich, govern	Purpose pipeline Route to the right consumers	Cloud and enterprise Deliver fit-for-purpose data
--	---	--	---	---

Figure 7 — From source to context to consumer.

An Industrial Edge Data Plane can typically be implemented using a relatively small number of appropriately sized computing nodes distributed according to the operational and cybersecurity architecture of the facility. A smaller manufacturing site may require only one or two nodes. A larger facility may deploy several nodes based on production areas, network segmentation, availability requirements, data volumes, or IEC 62443 zones and conduits.

Thousands of tags, hundreds of instruments, and dozens of production assets may be serviced by a small number of managed edge nodes. Edge infrastructure does not necessarily scale linearly with the number of connected assets.

Reference architecture

The Industrial Edge Data Plane can be represented as five logical layers. Each layer has a distinct responsibility.




Figure 8 — Five logical layers, bottom to top.

Layer 1 — Industrial (OT) sources

The architecture begins with the systems that create and manage operational information: PLCs and PACs, distributed control systems, SCADA systems, process and enterprise historians, smart instruments and analyzers, drives and motor-control systems, machine and robot controllers, vision systems, MES/MOM applications, quality and laboratory systems, maintenance and asset-management systems, and utility and environmental systems.

Connectivity may use OPC UA, MQTT, EtherNet/IP, Modbus TCP, REST APIs, SQL, vendor-specific interfaces, and other industrial protocols. Importantly, implementation does not require replacement of these systems. It provides a common data-services layer above existing automation investments.

Layer 2 — Managed edge infrastructure

This layer consists of industrialized computing resources deployed within the plant and managed as a coordinated fleet rather than as individually administered PCs or unmanaged gateways. It should provide centralized device and application management, secure workload deployment, application lifecycle management, device identity and authentication, workload isolation, health and availability monitoring, remote updates, local storage, cybersecurity policy enforcement, and resilience during loss of enterprise or cloud connectivity.

Without centralized lifecycle management, edge computing can simply create another population of difficult-to-maintain industrial computers. A representative technology for this layer is ZEDEDA.

Layer 3 — Industrial DataOps and contextualization

Above the infrastructure layer sits the core data-services capability. This layer connects to industrial sources and transforms their proprietary or source-specific representations into governed industrial information. Its responsibilities include acquire, model, normalize, contextualize, transform, and route. This is where the architecture begins to separate the meaning of industrial information from the systems that originally produced it.

A PLC tag such as Line3.PLC2.FIC101.PV may be transformed into an information model identifying:

Measurement	Flow rate
Instrument	FIC-101
Equipment	Transfer Skid 2
Area	Formulation
Process	Material transfer
Batch	B20260810-017
Product	Product A

Representative technologies for this layer include HighByte Intelligence Hub, AVEVA CONNECT Edge, and Litmus. Standards-based models further strengthen it: ISA-95 and ISA-88 hierarchies, OPC UA information models, CESMII Smart Manufacturing Profiles, and industry-specific semantic models. The objective is not simply to rename tags. It is to establish persistent meaning that can travel with the data.

Layer 4 — Edge intelligence and OT applications

Once industrial data has been acquired and contextualized, selected analytical workloads can execute directly within the edge environment: condition-based maintenance, anomaly detection, equipment-health calculations, predictive analytics, energy optimization, quality monitoring, process-performance calculations, OEE and machine-state analysis, AI/ML inference, and event detection and notification.

Representative technologies include Sorba.AI, Velotic CSense, and AVEVA CONNECT Advanced Analytics. High-frequency source data required for an analytical model does not necessarily have to leave the facility. A local model may continuously analyze thousands of observations and produce a relatively small number of meaningful results.

PROCESSED LOCALLY	SENT UPSTREAM
High-frequency vibration	Asset health score
Motor current	Detected anomaly
Speed	Failure probability
Temperature	Recommended action

Move the insight when moving all of the underlying data provides little additional value.

Layer 5 — Enterprise and cloud consumers

The Industrial Edge Data Plane does not replace enterprise data platforms or cloud computing. It provides them with better-defined industrial information. Potential consumers include enterprise data lakes and lakehouses, Microsoft Azure and Fabric, AWS, Google Cloud, Databricks, Snowflake, enterprise historians, business-intelligence platforms, AI and machine-learning environments, digital twins, enterprise knowledge graphs, MES/MOM and manufacturing applications, and maintenance and reliability platforms.


These systems receive data products appropriate to their requirements rather than identical copies of every source signal.

Cloud analytics platform	One-minute contextualized time-series data
Maintenance system	Equipment-state changes and detected anomalies
Performance platform	Production events and calculated KPIs
Data-science environment	High-resolution datasets, temporarily, for model development

The edge therefore becomes a distribution and policy boundary, not a destination.

One source, multiple data products

One of the most important characteristics of the Industrial Edge Data Plane is its ability to create multiple representations of the same source information. Consider a motor producing current, speed, vibration, and temperature measurements. The same source data might support four simultaneous pipelines:



The source is acquired once, its meaning is established once, and governed representations are distributed according to consumer requirements. This reduces point-to-point integration while creating an opportunity for consistent semantics and data quality across applications.

Separating the architecture from the technology

Although this paper uses ZEDED, HighByte Intelligence Hub, and Sorba.ai as representative technologies, the Industrial Edge Data Plane should not be defined by any individual product. The architectural capabilities are more important than the implementation technologies. The three principal capabilities can be summarized as:

Edge infrastructure management

Securely deploy, manage, monitor, and maintain distributed computing resources

Industrial DataOps

Acquire, model, contextualize, transform, govern, and distribute industrial information

Edge intelligence

Analyze operational information locally and transform high-volume observations into actionable insights

Separating these concerns creates flexibility. Infrastructure management can evolve independently from contextualization technology. Analytical applications can be added without redesigning connectivity. Cloud destinations can change without requiring fundamental changes to the underlying OT systems.

A shared OT computing and data platform

The broader significance of the Industrial Edge Data Plane is that it establishes a common platform upon which multiple OT and enterprise initiatives can be built. Rather than creating separate infrastructure for each predictive-maintenance, energy-management, data-lake, AI, or digital-transformation project, the organization establishes reusable capabilities for connectivity, context, compute, governance, and distribution.




Figure 9 — A new application consumes an existing contextualized data product rather than another direct connection.

This changes the architecture from a collection of projects into a reusable industrial capability. Over time, the edge becomes the location where the organization establishes a consistent understanding of its operational information and determines how that information should be used and distributed.

OT systems generate data. The Industrial Edge Data Plane establishes meaning and determines purpose. Enterprise and cloud platforms consume the resulting information.

The future of industrial data architecture should not be defined by how much data an organization can move to the cloud. It should be defined by how effectively that organization can transform operational data into useful information and deliver it wherever it creates the greatest value.