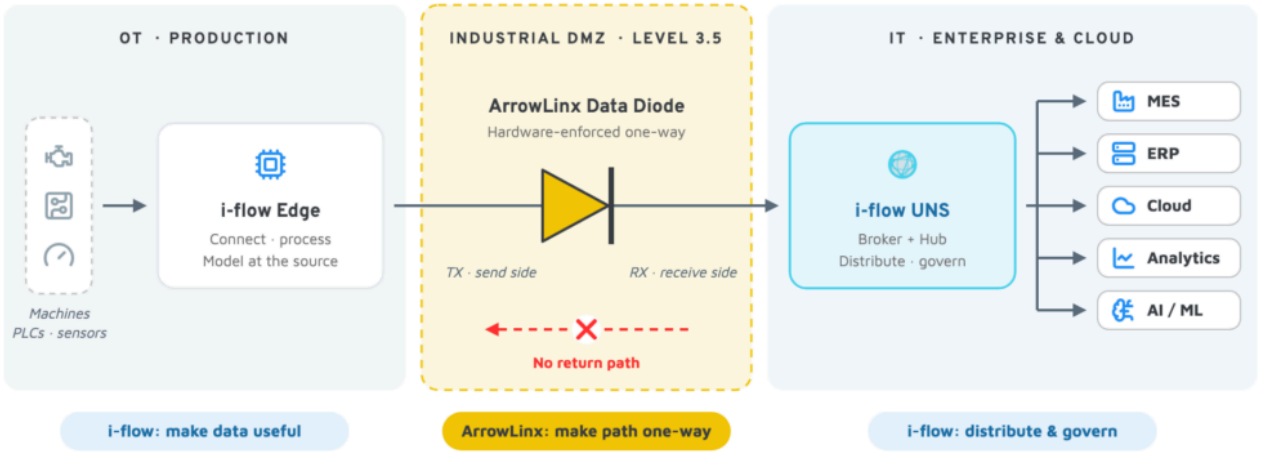
UNS Architecture & Design

Data Diode in the Unified Namespace: Secure OT-to-IT Data Flow

Originally published by i-flow GmbH: i-flow.io/en/ressources/data-diode-in-the-unified-namespace-secure-ot-to-it-data-flow/

Manufacturers can make production data available to every enterprise system without making the production network reachable from them. A data diode in the Unified Namespace (UNS) makes this possible: i-flow connects and models production data on the plant floor, and an ArrowLinx Data Diode from Allied Solutions sits between operational technology (OT) and information technology (IT). Modeled data moves from OT to IT, and no return path exists across the boundary. This article explains the conflict between data availability and OT security and the role of each component. It also shows why the combination of data diode and Unified Namespace separates two jobs that are often mixed together. The second part of this series then covers the detailed reference architecture.



The Challenge: Data Availability vs. OT Security

Every new consumer of production data adds a connection toward the plant, and every connection is potential exposure. A Manufacturing Execution System (MES) needs production context, and Enterprise Resource Planning (ERP) needs operational events. Cloud platforms need telemetry, analytics needs real-time and historical data, and AI needs trustworthy, contextualized information. Meanwhile, production must stay resilient, deterministic and protected from incidents originating outside OT.

A UNS solves the integration problem: data is published once and consumed by many. How this works is explained in our article What is the Unified Namespace (UNS)?. However, the UNS does not by itself solve the network problem. Most industrial protocols and brokers are bidirectional. A UNS connected straight to the plant floor can therefore leave a route back into production.

Firewall or data diode?

Firewalls remain essential to defense in depth. However, a firewall is a policy applied to a path that can carry traffic both ways. Where the business need is strictly one-way, such as exporting telemetry to IT, a hardware-enforced data diode provides a stronger property: the reverse path does not exist.



Design principle: separate data usefulness from network reachability

The combination of a data diode and the Unified Namespace follows a small set of design principles. Together, they define which data leaves production and how the boundary is enforced:

- **Separate data usefulness from network reachability:** Systems can use production data without being able to reach production.
- **Expose only what is needed:** Production systems expose only the data required for business and operational purposes.
- **Model before consumption:** Data is normalized and contextualized before it is consumed across the enterprise.
- **Explicit boundary:** The OT/IT boundary is explicit and independently enforceable.
- **No unnecessary return path:** Where no return traffic is needed, the architecture does not create a return path just because a protocol supports one.
- **Defense in depth:** The architecture complements, never replaces, segmentation, identity management, patching, endpoint protection, monitoring and secure operating procedures.

ArrowLinx Data Diode: Technical Role and Customer Value

The ArrowLinx Data Diode is a hardware-enforced unidirectional gateway for transferring information between isolated or differently trusted networks. The core architectural concept is simple. Data moves from the send side (TX, OT) to the receive side (RX, IT), while the reverse communication path is physically not available across the diode.

What the technology contributes

- **Hardware-enforced direction:** Unidirectional data transfer at a network boundary.
- **Isolation:** The protected OT environment is isolated from return traffic originating in the destination network.
- **Fit for one-way use cases:** The architecture suits scenarios where data export is required but inbound connectivity is not.
- **Separation of functions:** Moving data and enforcing network direction are clearly separate functions.
- **Additional security layer:** The diode complements, not replaces, industrial segmentation, identity, endpoint security, monitoring and operational controls.

Customer value

- **Smaller attack surface:** An otherwise unnecessary return path is removed.
- **Easier to explain:** Plant engineering, IT security and management can follow the security architecture.
- **Digitalization without exposure:** The production network does not need to become broadly reachable from enterprise or cloud environments.
- **Repeatable pattern:** The same reference architecture applies across plants and production lines.

One point matters for the security positioning: a diode does not make an entire plant "unhackable". Rather, the ArrowLinx boundary enforces a one-way communication property that software alone cannot provide on the same network path. Our article on secure UNS architectures describes the further measures that belong in a defense-in-depth concept.

i-flow: Unified Namespace and Industrial Data Infrastructure

i-flow provides a modular industrial data platform built around the UNS concept. Its architecture combines i-flow Edge, i-flow Broker and i-flow Hub. Together, they connect industrial sources, process and contextualize data, publish it through an event-driven data layer and centrally manage distributed deployments.

i-flow Edge

i-flow Edge is the distributed runtime close to the data source. It connects to Programmable Logic Controllers (PLCs), `OPC UA` servers, sensors and other industrial systems. It processes data in real time, harmonizes it and models it into the UNS structure before it leaves the plant floor. Because the same platform runs from edge to enterprise, the data model defined at the source is the one every consumer sees. Existing connectivity such as Kepware or an OPC UA server can still be used as a source for i-flow Edge.

i-flow Broker

i-flow Broker is the event and message layer of the UNS. It follows a publish-once, distribute-everywhere principle. Producers publish contextualized data once, and any number of authorized consumers subscribe without point-to-point integrations. At plant level, it supports `MQTT` and `NATS`. Across plants, `NATS` serves as the global backbone. The differences between the two protocols are covered in our comparison NATS vs. MQTT.

i-flow Hub

i-flow Hub is the central management and governance layer. It provides configuration, monitoring, access control and lifecycle management for distributed i-flow components and the UNS.

Data Diode and Unified Namespace: Clearly Separated Responsibilities

ArrowLinx and i-flow address complementary layers of the industrial data architecture. i-flow answers the question: "How do we acquire, understand, structure and distribute production data?" ArrowLinx answers the security question: "How do we enforce that data can cross this boundary in one direction only?" Each component owns one job. This makes the design easy to explain, validate and replicate.

Layer	Component	Primary responsibility	Value to the customer
Edge connectivity & modeling	i-flow Edge	Connect to PLCs and controllers; process and harmonize data; model it at the source	UNS-ready data from the first hop; low-latency processing close to the machines
Security boundary	ArrowLinx Data Diode	Hardware-enforced one-way transfer from TX (OT) to RX (IT)	Reverse path physically unavailable across the boundary
UNS / messaging	i-flow Broker	Publish once; distribute to many consumers	One governed data source for plant and enterprise applications
Central governance	i-flow Hub	Manage, monitor and govern the distributed data infrastructure	Central control while processing stays close to production
Applications	MES, ERP, cloud, analytics, AI	Subscribe to the information they need	Governed production data with no direct OT reachability

Three stages from machine to enterprise

At a high level, the combination of data diode and Unified Namespace works in three stages:

1. **Connect and model:** i-flow Edge connects to the required sources on the OT side and models the data into the UNS structure.
2. **Cross the boundary one way:** Only the required outbound stream crosses the ArrowLinx Data Diode in the Industrial DMZ (Level 3.5 of the Purdue model).
3. **Distribute and govern:** On the IT side, the i-flow UNS publishes the data for plant applications and, across all plants, for ERP, cloud, analytics and AI. i-flow Hub governs it centrally.

Data is modeled before it crosses. Only harmonized, contextualized information leaves OT. As a result, the boundary carries exactly what the business needs, and the UNS receives data that is ready to use. Because consumers subscribe to the UNS instead of connecting to machines, a new dashboard, model or business system is added on the IT side. Neither the plant floor nor the diode configuration has to change.

The second part of this series shows how this pattern looks in detail: a reference architecture with data diode across all Purdue levels for plants with and without existing Kepware connectivity, including data flow, multi-plant scaling and central governance.

Conclusion: Connect the Factory without Connecting the Threat

Industrial digitalization does not require a trade-off between data availability and OT isolation. A data diode in the Unified Namespace makes production information broadly available while keeping a strong separation between production and the systems that consume its data. i-flow makes industrial data useful. ArrowLinx makes the data path one-way. Together they build a secure bridge from production data to the Unified Namespace, without a network bridge back into production.

Three key takeaways

1. **A UNS solves integration, not network direction:** Publish-once data distribution removes point-to-point interfaces. A hardware-enforced boundary is still needed where the business requirement is strictly one-way.
2. **Clear separation of responsibilities:** i-flow handles connectivity, modeling, publishing and governance. The ArrowLinx Data Diode enforces the direction of data flow between OT and IT.

3. **Model first, then cross:** Because i-flow Edge models data at the source, only the information the business needs crosses the diode, and it arrives in the UNS ready to use.

The combination of data diode and Unified Namespace is not a complete cybersecurity program. It is one clearly defined, enforceable layer within a defense-in-depth architecture. The next article in this series describes the reference architecture in detail.

About Allied Solutions: Allied Solutions (Allied Solutions Pte Ltd) is an industrial automation, digital manufacturing and OT cybersecurity company headquartered in Singapore, with offices in Vietnam, Malaysia and China. At the core of its OT cybersecurity portfolio is the ArrowLinx Data Diode, a hardware-enforced unidirectional gateway certified to Common Criteria EAL4+. With more than 25 years of experience in industrial environments, the company helps manufacturers, utilities and critical infrastructure operators modernise their operations securely, from industrial data connectivity, MES and industrial analytics to the protection of the OT networks that run them.

About i-flow: i-flow is an industrial software company based in southern Germany. The company stands for a new era of self-connecting factories — and the end of manual integration. Its platform connects factories fully automatically, at any scale, worldwide. Over 750 million data operations per day in production-critical environments demonstrate the scalability of the software and the deep trust that customers place in i-flow. This success is based on close collaboration with customers and partners worldwide, including renowned Fortune 500 companies and industry leaders like Bosch.