




SYNADIA

How Intelecy Built a Super Fast Bi-Directional Data Pipeline with Synadia NGS and JetStream



Intelecy

Case Study



Industry
Industrial



Location
Norway



Employees
15

Overview

Intelecy is No-Code AI platform that helps industrial companies optimize factory and plant processes with real-time machine learning insights. These insights allow process engineers to manually or automatically improve outcomes, quality and sustainability without requiring significant custom code or hardware. Designed to help factories move towards greater automation and true “lights out” operational capabilities, Intelecy securely collects data from tens of thousands of sensors in a factory and pushes that data into the cloud using NATS and the NGS global data fabric. The Intelecy platform then runs that data through machine learning models customized for each industrial process in a plant and returns analysis and insights back to the sensors and controlling systems in near real time, with less than one-second round-trip latencies. Intelecy can ingest and process data in all widely used industrial formats including SCADA, HMI, DCS, and MES. Intelecy customers come from a variety of industrial sectors including food and beverage, materials, power generation, and mining. Intelecy runs on a cloud native architecture with loosely coupled services utilizing containers orchestrated by Kubernetes/Nomad, the Envoy proxy, gRPC for messaging, ClickHouse for time series data storage, and NATS for data streaming and intelligent persistence.

About Intelecy

Customers: TINE, Heidelberg Cement, IFF, AM Nutrition, Eidsiva Bioenergi, Glencore Nikkelverk

Tech Stack: Kubernetes, Nomad, ClickHouse, gRPC, NATS, Kafka, Microsoft Azure, Microsoft Windows, Linux

Application Language: Go

Background

Intelecy is operating in a sector that changes slowly due to the risks of production interruption; shutting down a factory or a production line, even for a few minutes, is costly. The industrial sector also operates using a wide array of systems and sensors, requiring flexible integrations and ability to interact with and control several types of different operational technology (OT) assets and protocols. For the most part, as well, factories are running with older technology paradigms and designs; many use historians, a data logger that functions as a crude time series database.

At its core, Intelecy seeks to provide a bi-directional data pipeline that allows industrial companies to leverage cutting-edge machine learning for on-the-fly insights to optimize operations. “To build vision, we knew we needed two things,” explains Jonathann Camp, Intelecy’s Chief Security Officer.

“*We needed something that could handle streaming of data of various message sizes and also gracefully handle intermittent connectivity.*”

Jonathann Camp, Intelecy’s Chief Security Officer

“First, we needed a simple extensible method for commanding something that sits behind a firewall inside a factory and has no externally initiated connections and cannot be controlled from the outside in any meaningful way. Second, we needed something that could handle streaming of data of various message sizes and also gracefully handle intermittent connectivity.”

The Intelecy Application Architecture

Intelecy's application layer is written in Go and composed of multiple services running on containers orchestrated by Kubernetes / Nomad. Intelecy puts a gateway — a Windows virtual machine — on the plants' networks. The gateway connects securely to the gateway controller running in an Azure Cloud instance. Once data from a plant enters the Intelecy platform, it passes through a data pipeline to normalize the data structure and prepare it for machine learning models. Intelecy also pushes the data in its raw form into a cloud-based object store. From the machine learning models and the data pipeline, plant data goes to a ClickHouse time series database. This database replaces traditional factory historians and can serve as the historical record for plant process and sensor data. From ClickHouse, Intelecy offers API access.

Users, including process engineers and continuous improvement teams, can either view insights and data directly on Intelecy's front-end or integrate the data and machine learning insights into other systems of record or presentation layers and analytics engines. A crucial difference from competitors is Intelecy's bidirectional capability; unlike other industrial process analytics systems, Intelecy can write data back to the control units in the factories. This enables automation and closed-loop optimization based on rules created by engineering teams or machine learning models. "For example, a dairy can use a machine learning forecast model to set temperature control systems using the real-time predicted state of the cheese production process," explains Camp. "The process engineering team can use Intelecy insights to identify trends and then automate temperature adjustments on a vat of yogurt to ensure quality and output are not compromised.

Intelecy Needed:

- ✓ **Low Latency:** Less than 2 seconds round-trip time for processing streamed data and returning ML-driven insights to customers
- ✓ **Segmentation and Security:** Secure and simple customer segmentation of data streaming in a multi-tenant environment
- ✓ **Portability:** A cloud-native and edge-friendly solution that was portable and cloud agnostic

Challenges: Latency, Segmentation, Portability

As a design convention, Intelecy wanted to leverage managed services running in the cloud as much as possible for non-core functionality. The engineering team also realized quickly that they would need to make modifications on the original architecture to reduce latency of data insight delivery, improve segmentation to handle multi-tenancy. In addition, they wanted to update the design of their application and services stack to be as cloud agnostic and container-centric as possible, avoiding inclusion of cloud-specific services that could create a lock-in or a bottleneck.

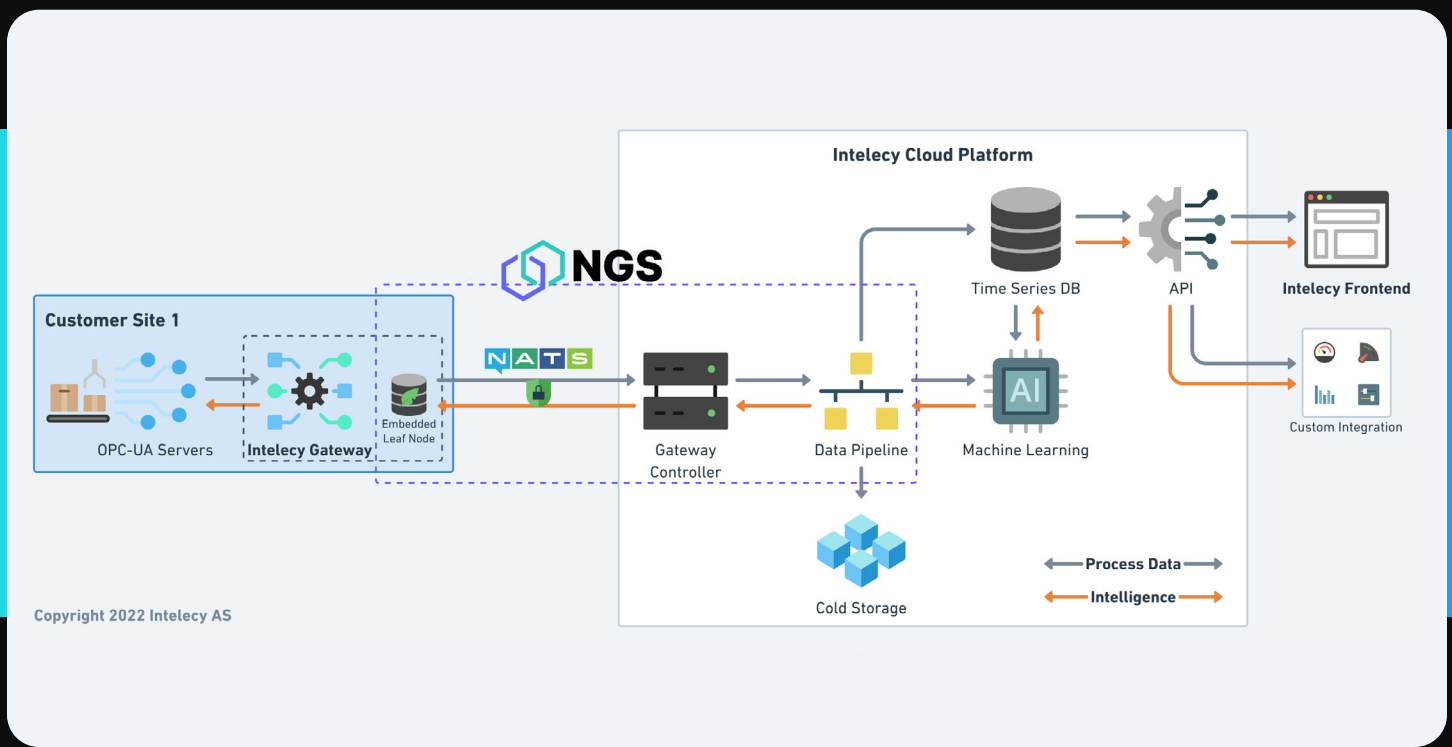
Intelecy initially designed its platform using poll-based systems that were tightly tied to Microsoft Azure services. Those services included SQL data stores but not time series, which was a requirement for many industrial customers. In addition, the poll-based architecture could return processed data and insights back to customers with a delay of 15 to 30 minutes. When customers began to ask for faster insights, Intelecy began to look at alternatives to traditional pub-sub poll-based systems and considered how to architect a more edge-centric application to take advantage of cloud native principles and software.

Camp wanted flexibility, both in terms of restructuring the client topology (or adding and deleting clients) without restarts, and in terms of



message sizes. Equally important, Intelecy wanted to be able to scale their data streaming footprint up and down quickly, without having to restart the system or make big configuration changes when new clients came online or clients were destroyed. Lastly, Intelecy wanted to ensure a high level of security for any data streaming service due to accommodate the requirements of customers. Camp disliked the security conventions of most data streaming services, which implemented security as a secondary element on top of the data streaming layer. This would mean additional overhead and management

Why Intelec Deployed NATS, NGS and JetStream



The Intelec team had worked with NATS before and immediately considered it as the data streaming solution for their application requirements. Camp liked the simplicity of NATS and was interested in the NATS utility SaaS offering, NGS, from Synadia. NGS would eliminate the need for Intelec to maintain a data streaming infrastructure and only consume as much capacity and bandwidth as it needed. Camp also liked the manner in which Synadia’s team (the core maintainers of NATS) worked on the software in public and interacted with issues, PRs and community members. “The way someone develops software tells you a lot about its quality. And the fact that the creators of NATS are so heavily involved in the code on a day-to-day basis gave me confidence,” says Camp.

Intelec did look at a handful of other streaming and messaging applications, including some using MQTT, but found them to be less flexible and requiring more resources for stand up and maintenance. Intelec briefly looked at Apache Kafka but decided against using it extensively for data streaming. “Kafka has requirements that it puts on the client,” says Camp. “The client has to be aware of a lot of things. I never liked that design paradigm. I wanted the client to be really simple.” The worker queues model, as well, exposed complexity and made changing the topology of the application challenging.

After a quick decision process, Intelecy opted to replace its polling-based pub-sub data fabric with NGS. “The hosted service closed the deal,” says Camp. “We are a small startup. We don’t have time to manage everything.” For its application, Intelecy uses a data plane to handle data streaming back and forth to plants and the various sensors, as well as a control plane to manage these interactions. Intelecy constructed a global control plane using RPC and protobufs that could issue commands to all Intelecy gateways simultaneously. The control plane is primarily designed to deliver intelligence back to the plant floor running Intelecy agents to move or receive data. “I can send a request to all the gateways around the world but RPC can also go the other way,” explains Camp. “When a gateway comes online, it uses RPC to register with the gateway controller.”

BENEFITS:

Intelecy Technology Benefits and Results


- ✓ 15-30x reduction in roundtrip latency
- ✓ Cloud and infrastructure agnostic - no lock-in
- ✓ Requires zero maintenance
- ✓ Improved resiliency and higher availability
- ✓ Improved security for multi-tenancy
- ✓ Eliminated need for complex proxy or service mesh
- ✓ Enables combination of data streaming and batch
- ✓ Handles large message sizes

Intelecy Business Benefits

- ✓ Onboard and manage customers easily
- ✓ Improved security
- ✓ Less technology complexity
- ✓ Reduced operational overhead
- ✓ Easy scalability

The data plane is primarily a fabric of NATS services running on NGS. For Intelec, the gateway controller connects to the gateway, which in turn is linked to the factory or plant sensor web and controllers. These endpoints often number in the tens of thousands. The Intelec gate fetches data in real-time from those sensors in one of the data formats common to industrial processes and writes that process data back to NGS to a specific topic. “On the data pipeline, we have multiple types of workers listening to various topics,” says Camp. “So, for example, the log ingestion service listens for the log topic and when the logs come back from the gateway, it grabs them and sends them to DataDog.” For machine learning insights in real-time, Intelec has live models that are running with associated NATS clients in NGS. The machine learning containers are triggered by a specific subject linked to a customer’s process data.

Although originally written with a BoltDB caching layer, Intelec added persistence when Synadia added the JetStream persistence and durability layer to NGS. JetStream made NGS far more useful to Intelec. “For process data that comes back, granular control over durability was our missing piece,” says Camp. “We have some buffering capability built into the gateway but we’ve never liked that because we don’t like being responsible for synchronizing stores. JetStream completely solves that.”



The “least-privilege by default” security model for NATS appealed to Camp, in part because it simplified multi-tenant security. Intelecy could easily segregate customers by limiting their NATS clients to subjects associated with that specific organization. NATS (and, by extension, NGS) incorporates all permissions directly into JSON Web Tokens. This eliminates the need for additional permissioning infrastructure and makes it simple to use NGS for Zero Trust architectures where permission is continuously requested and granted.

Because NATS is so performant, Intelecy was able to reduce round trip times from 15 to 30 minutes down to 1 second, meeting the most rigorous customer demands. NGS is also cloud agnostic, allowing Intelecy to move data streaming from one cloud environment to another in a matter of minutes. Scaling NGS up and down, or adding new clients and new nodes to a cluster did not require a restart or any changes to how the service was configured. “I can scale up and down just by adding or removing workers listening for a subject on messages,” says Camp. “I don’t have to configure anything. They all just join the queue and distribute the work.”

In the future, Intelecy may make use of the key value store capabilities in JetStream to store tables and structured data that changes rarely but is frequently accessed. NGS LeafNodes, another innovative feature of NATS, will allow Intelecy to set up high availability configurations at factories with two gateways running. Says Camp, “We tested NATS Leaf Nodes. They were easy to set up and just worked. That’s why we like NGS and NATS.”

“ We tested NATS Leaf Nodes. They were easy to set up and just worked. That’s why we like NGS and NATS.”

Jonathann Camp, Intelecy’s Chief Security Officer