Nazare

Industrial Big Data Platform

Get Your Industrial Machinery Data with Ease.

Be Original,
Yet Adaptable

Ingkle develops and supplies software for building facility-centric big data infrastructure. **We enable enterprises to instantly utilize real-time collected** machinery **data for their operations, without wasting time on big data infrastructure development.** Through our proprietary big data processing technology, **Ingkle** resolves the limitations of generic databases and cloud solutions regarding performance, security, and scalability that manufacturing industries demand.

With **Ingkle**, data becomes your competitive advantage.

Clients

SAMSUNG LG Electronics HYUNDAI STEEL HYUNDAI Rotem

SěAH Steel Hanwha Systems SUNGWOO HITECH MiCo JB CORP

85%

Top-tier Technical Talent Proportion
(AI & Software)

15+

Patents and Certifications

50+

Client Companies

200%

Revenue Growth Rate
(YoY)

20+

Big Data Platforms Deployed
in Industrial Fields

$2.63M

Total Investment Raised
(Seed Round: $75K, Pre-A Round: $2.48M)

HANYOUNG NUX

RENAULT

rexroth
A Bosch Company

Parker

gm

국방과학연구소
Agency for Defense Development

KITECH
Korea Institute of Industrial Technology

KIMS Korea Institute of
Materials Science

KHNP KOREA HYDRO &
NUCLEAR POWER CO., LTD.

"I'd like to collect data simultaneously at millisecond intervals from dozens of different types of machinery in the factory and analyze and utilize it in real-time with ease. If I could ask questions about the collected machinery data like ChatGPT and get quick answers, it would significantly improve work efficiency."

Manufacturer H's Maintenance Team

"We are already collecting power plant DCS data into a general database, but as the data grows, the performance of collection, storage, and querying drops significantly, causing operational challenges. Is there a solution that offers seamless scalability in storage capacity and performance without compromise?"

Energy Company J's Electrical Engineering Team

"Due to security concerns, we need to build an on-premise infrastructure for machinery data collection. Initially, it should be designed to handle a single machine, but it must allow for flexible server expansion in the future to scale our data infrastructure seamlessly."

Manufacturer M's Production Technology Team

Nazare
solves problems.

Machinery Manufacturer P

"When issues arise with the machinery supplied to our customers, the lack of detailed control data from the time of the incident makes root cause analysis time-consuming. We want to leverage data to enhance our machinery servicing and improve product development."
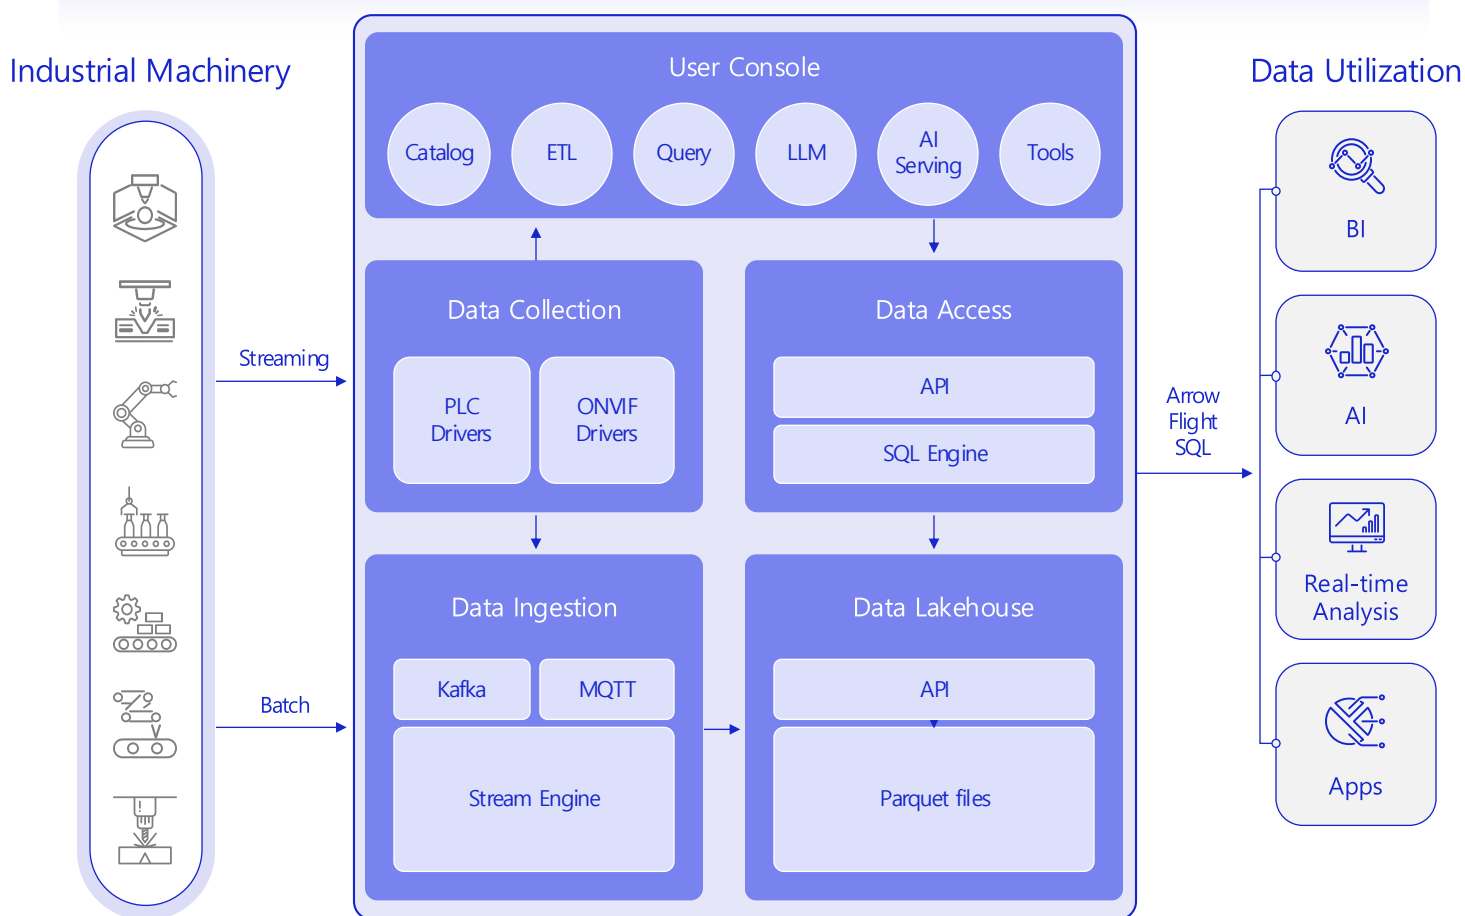
Manufacturing AI Solution Provider W

"We've developed AI models for our clients' machinery, but setting up the environment to collect, store, and run (inference) AI on machinery data in the field is taking too much time and effort."

THE ONLY DATA PLATFORM
FOR INDUSTRIAL MACHINERY

Nazare is an industrial data platform that delivers the performance, flexibility, and scalability needed to precisely collect, rapidly store, transform, and analyze data from industrial machinery. With **Nazare**, customers can seamlessly leverage data to accelerate their core business operations.

Big Data Setup Time	AI Inference Environment Setup	Big Data Setup/Operation Costs
6 MONTHS → 1 WEEK	1 MONTH → 1 DAY	OVER 20X REDUCTION

Nazare Data Platform Overview

Industrial Machinery — Streaming / Batch

User Console: Catalog, ETL, Query, LLM, AI Serving, Tools

Data Collection: PLC Drivers, ONVIF Drivers

Data Ingestion: Kafka, MQTT, Stream Engine

Data Access: API, SQL Engine

Data Lakehouse: API, Parquet files

Data Utilization (via Arrow Flight SQL): BI, AI, Real-time Analysis, Apps

Collect

Nazare collects data at millisecond intervals from multiple industrial machineries. It directly communicates with PLC and DCS controllers through a Fieldbus protocol configuration tool. Additionally, it supports integration with existing DAQ and SCADA machinery via MQTT or Kafka.

Manage

Nazare securely stores over 10 million data points per second in a unified big data repository. It offers horizontal scalability beyond petabyte volumes while ensuring consistent SQL query performance capable of responding to over 100 million data points per second. Without duplicating original data, it enables diverse workloads for BI and AI using the latest data. Intuitive ETL and aggregation can be configured through SQL and user-defined functions.

Access

Nazare provides access to real-time and collected data through a management console and built-in open-source tools for monitoring and analysis. External access to big data is enabled via APIs available in multiple languages, including Python, Java, R, Go, and Rust. For optimal performance, it supports FlightSQL based on Apache Arrow.

1. Collect

Effortlessly collect data from any machinery, at any scale.

Supports a wide range of Fieldbus protocols

Simultaneously collects millisecond-level data from multiple industrial machineries. Through the protocol configuration in the management console, it directly connects to controllers like PLCs and DCS to retrieve user-specified data. It supports open protocols such as OPC-UA and Modbus, as well as proprietary protocols from leading controller manufacturers like Fanuc, Mitsubishi, Siemens, Omron, AB, and LS.
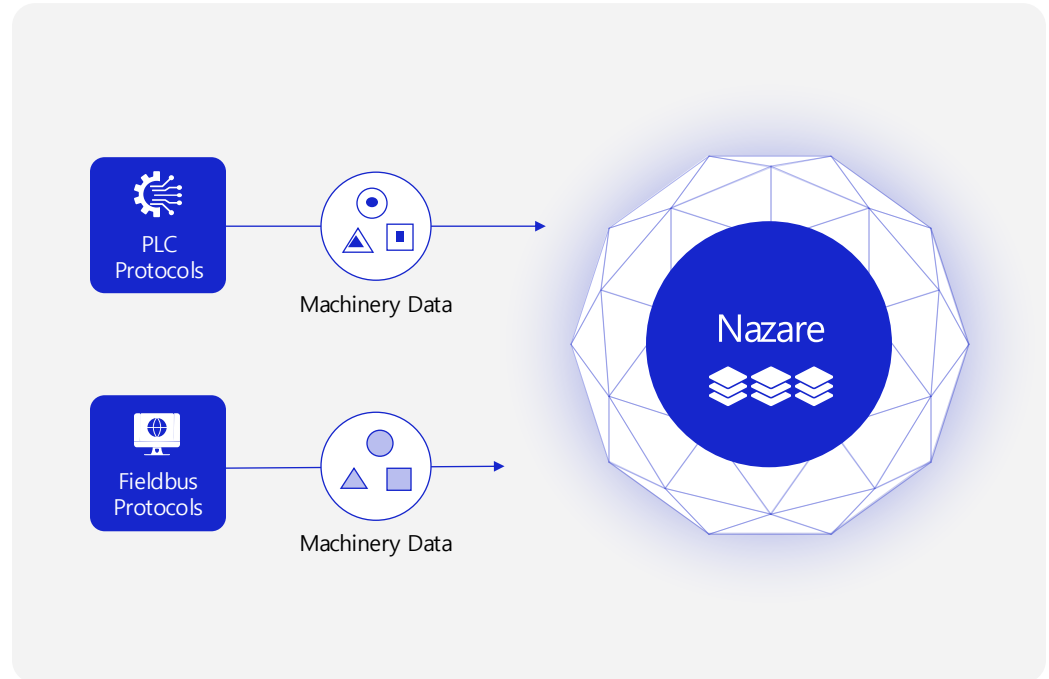

Flexible Data Collection Methods Virtual Edge and External Edge

Nazare supports both "Virtual Edge", where edge software runs inside the server to collect data, and "External Edge", where edge software is installed on external collection machineries that communicate with the **Nazare** server. Users can choose the edge type that best suits the characteristics of their industrial site, enabling seamless connection with machinery and efficient data collection.
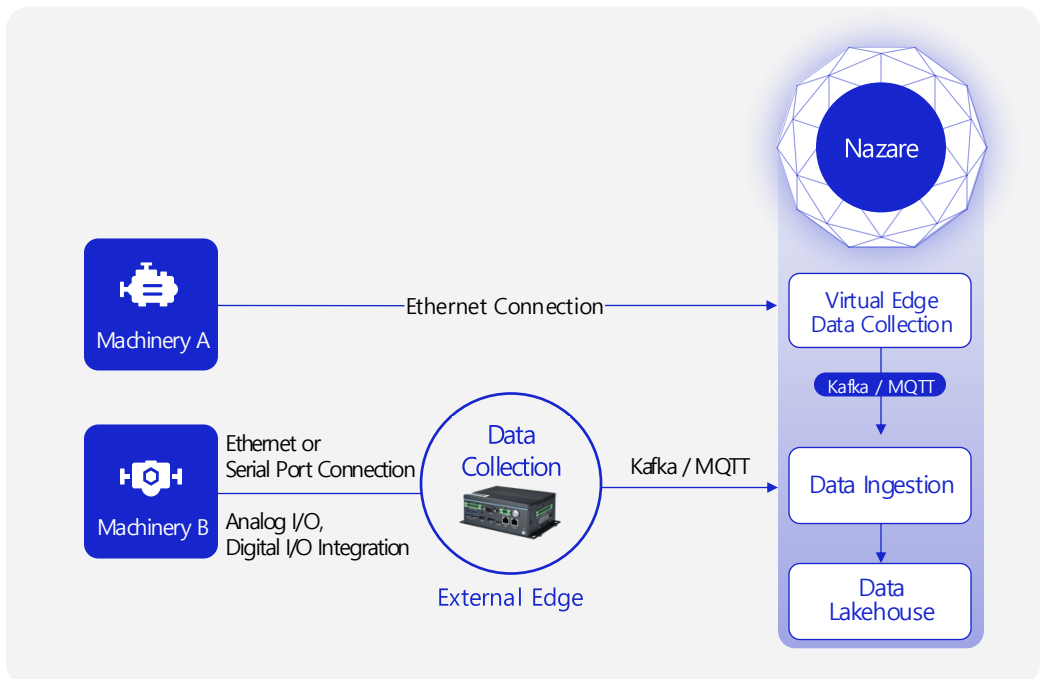
Seamless Integration with Third-Party OT Machineries and Software

Nazare integrates with third-party OT data acquisition machineries (DAQ, edge gateways) and OT software (SCADA, HMI, Historian) through MQTT and Kafka protocols to collect machinery data. It allows you to leverage your existing data collection infrastructure while effortlessly setting up a foundational environment for big data analytics and AI applications.
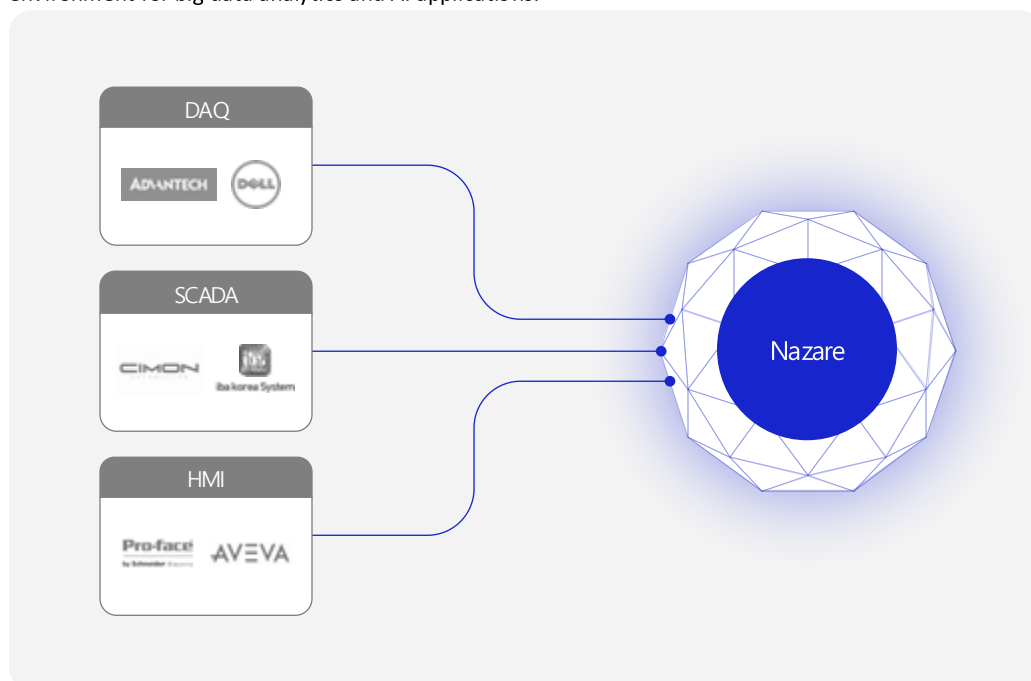
Collection of All Data Types

Nazare efficiently collects and stores not only time-series data generated from machinery controllers and sensors but also unstructured and semi-structured data such as images, videos, and text. This enables the consolidation of all industrial data into a single unified repository, significantly expanding the scope of data utilization and analysis.
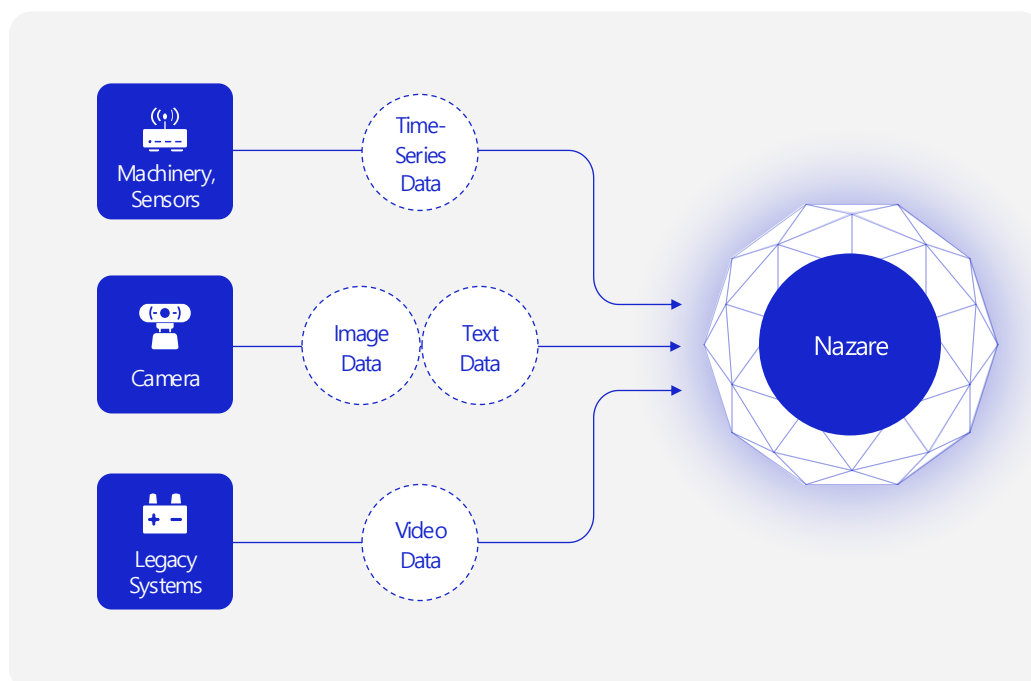
2. Manage

Easily handle large-scale data with fast processing and quick response times.

| Data Management Powered by a Lakehouse Architecture

Nazare manages data using a Lakehouse architecture, combining the scalability of a data lake with the ACID transactions and data management capabilities of a data warehouse. This enables seamless execution of Business Intelligence (BI) and Machine Learning (ML) workloads on all stored data without the need for duplication.
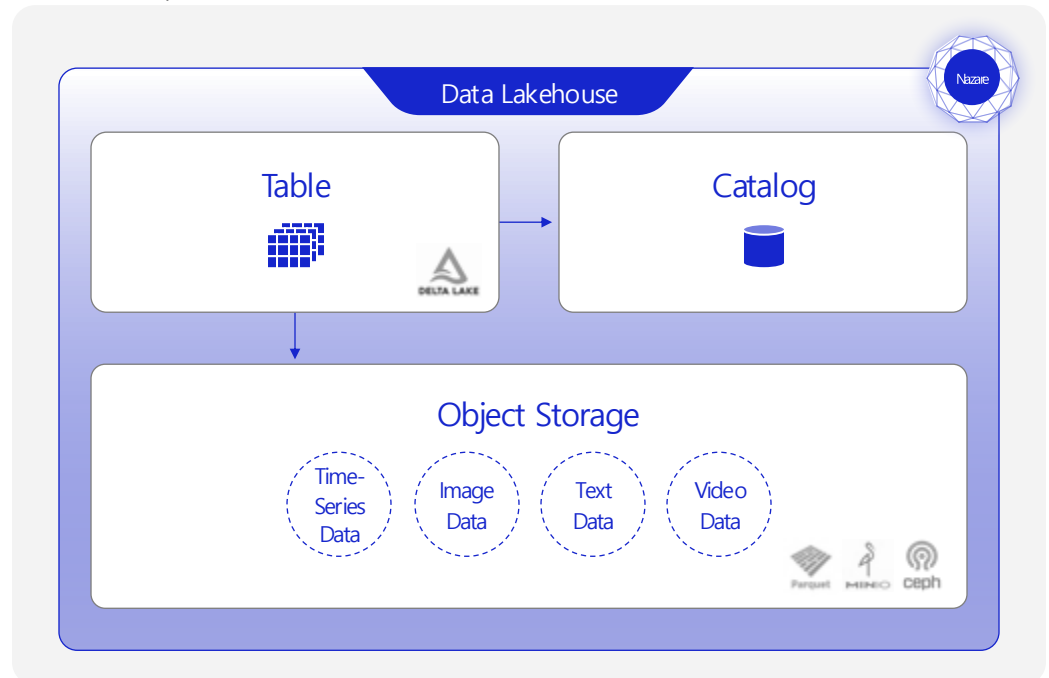
## | Real-Time Stream Processing for Large-Scale Data Collection and Transformation

Nazare can collect, store, and transform over 10 million data points per second from industrial machinery, powered by Ingkle's proprietary high-speed big data processing engine built on Rust. This enables users to analyze data in real-time, gain immediate insights, and make more accurate predictions and faster decisions.
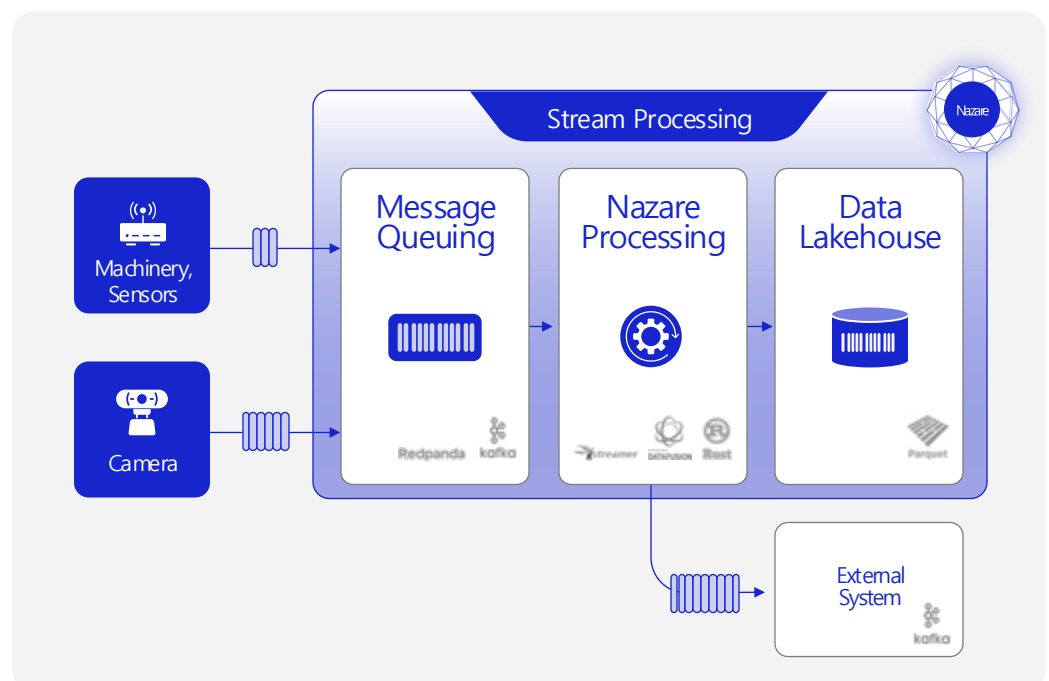
Ultra-Fast SQL Query Performance Powered by Apache Arrow

Nazare leverages the in-memory columnar format of Apache Arrow and the FlightSQL protocol to deliver exceptionally fast SQL query performance on large datasets. Its query speed is over 20 times faster on average than traditional relational databases (RDBs), processing over 100 million data points per second, making it ideal for real-time, interactive data analysis.
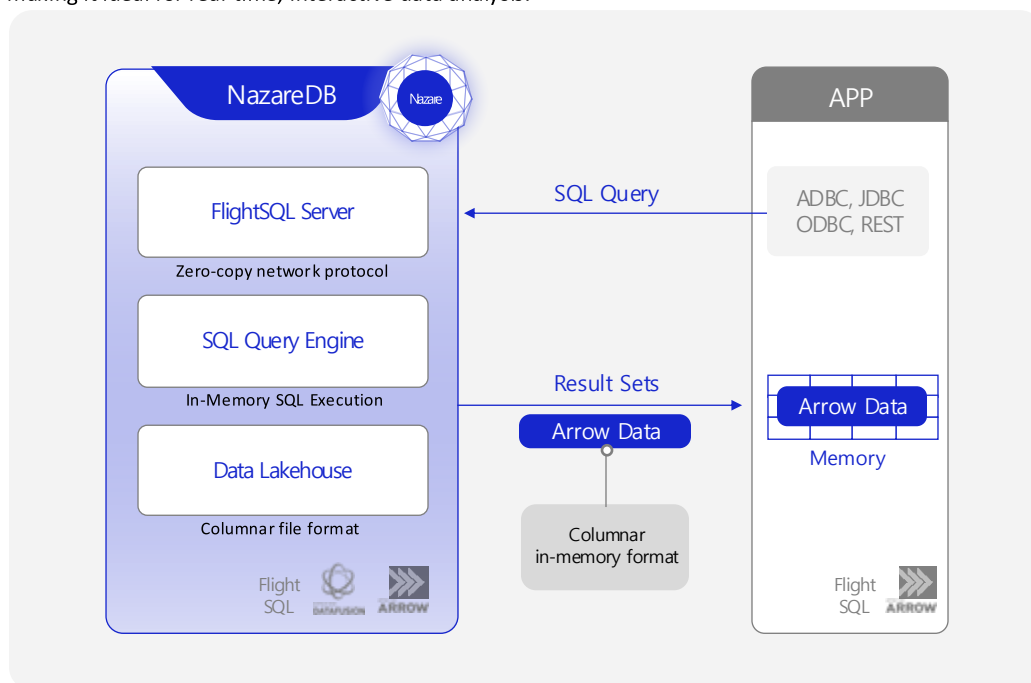
Data Engineering with SQL and UDFs (User-Defined Functions)

Nazare enables intuitive and declarative data engineering using SQL and Python-based User-Defined Functions (UDFs). This allows users to perform a wide range of tasks, such as condition checks and data transformations, on both real-time and large-scale collected datasets. Additionally, data aggregation and exporting to external systems can be executed seamlessly, simplifying complex workflows.
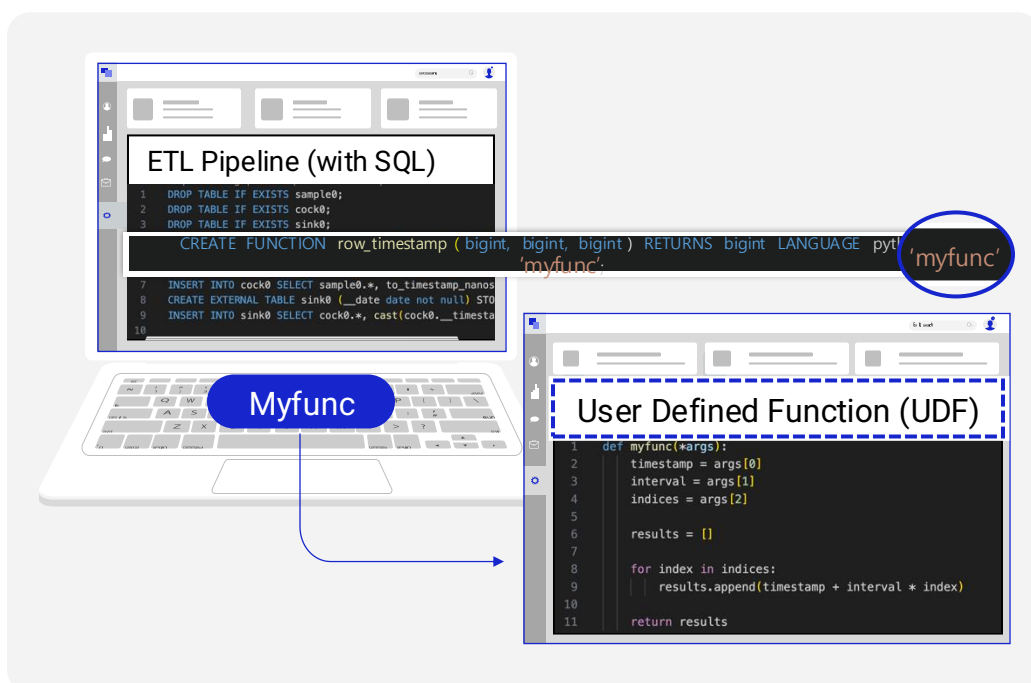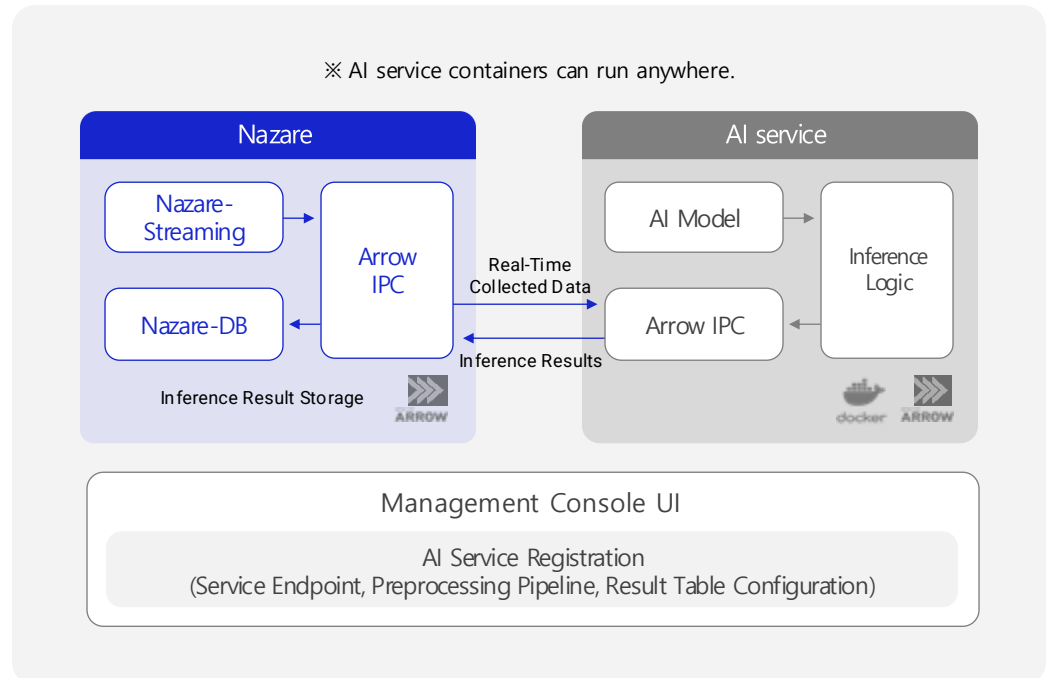
3. Access

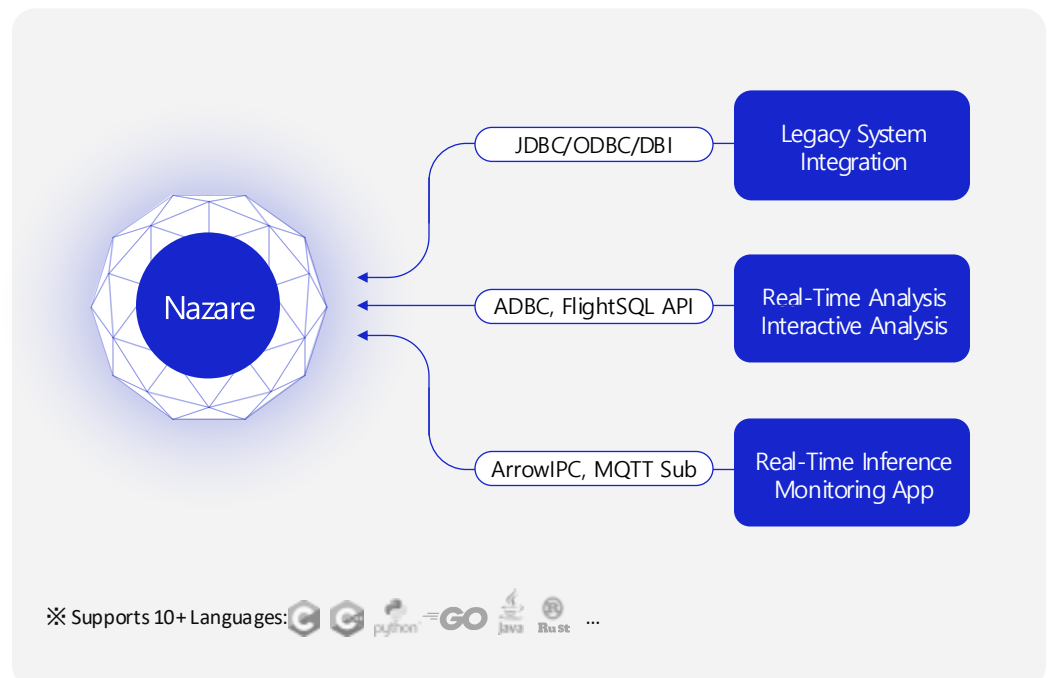## Leverage big data in real-time.

| Real-Time AI Model Inference

Nazare provides a seamless way to perform AI model inference using real-time collected data. By deploying AI models and their execution functions as Docker containers, users can obtain rapid inference results on real-time stream data collected by **Nazare**. The pipeline configuration in **Nazare** allows users to specify the container deployment location and the storage destination for inference results.

※ AI service containers can run anywhere.

Nazare		AI service

Nazare-Streaming

Arrow IPC — Real-Time Collected Data

Nazare-DB

Inference Results

Inference Result Storage

AI Model → Inference Logic

Arrow IPC

Management Console UI

AI Service Registration
(Service Endpoint, Preprocessing Pipeline, Result Table Configuration)

| Provides APIs for Real-Time and Historical Data Access

With **Nazare** APIs, users can develop various applications based on machinery data. Real-time data can be accessed in second intervals via the MQTT protocol, while historical data is available through Arrow FlightSQL, PostgreSQL protocol libraries, or REST APIs. Arrow FlightSQL, in particular, utilizes a binary format optimized for streaming, delivering the fastest response times.

Nazare

JDBC/ODBC/DBI → Legacy System Integration

ADBC, FlightSQL API → Real-Time Analysis Interactive Analysis

ArrowIPC, MQTT Sub → Real-Time Inference Monitoring App

※ Supports 10+ Languages: ...

Real-Time Data Analysis and Visualization Environment

Nazare provides a range of open-source tools for free, enabling users to perform data analysis and visualization on their own. With JupyterLab, users can run Python-based real-time interactive analysis, while the built-in open-source tools facilitate effective Business Intelligence (BI) tasks and dashboard visualizations.
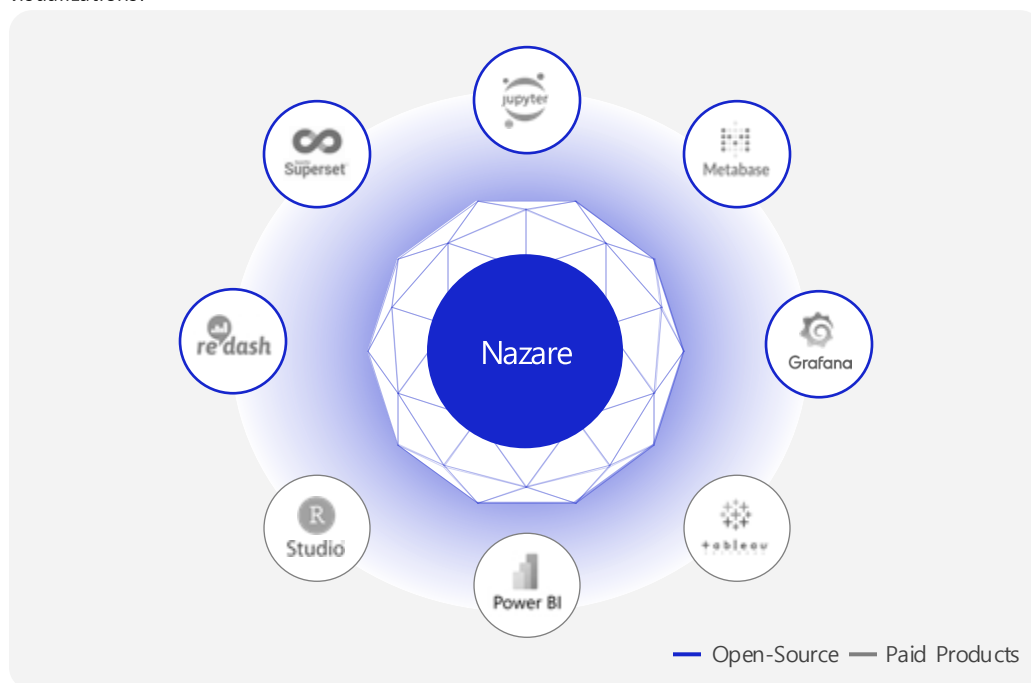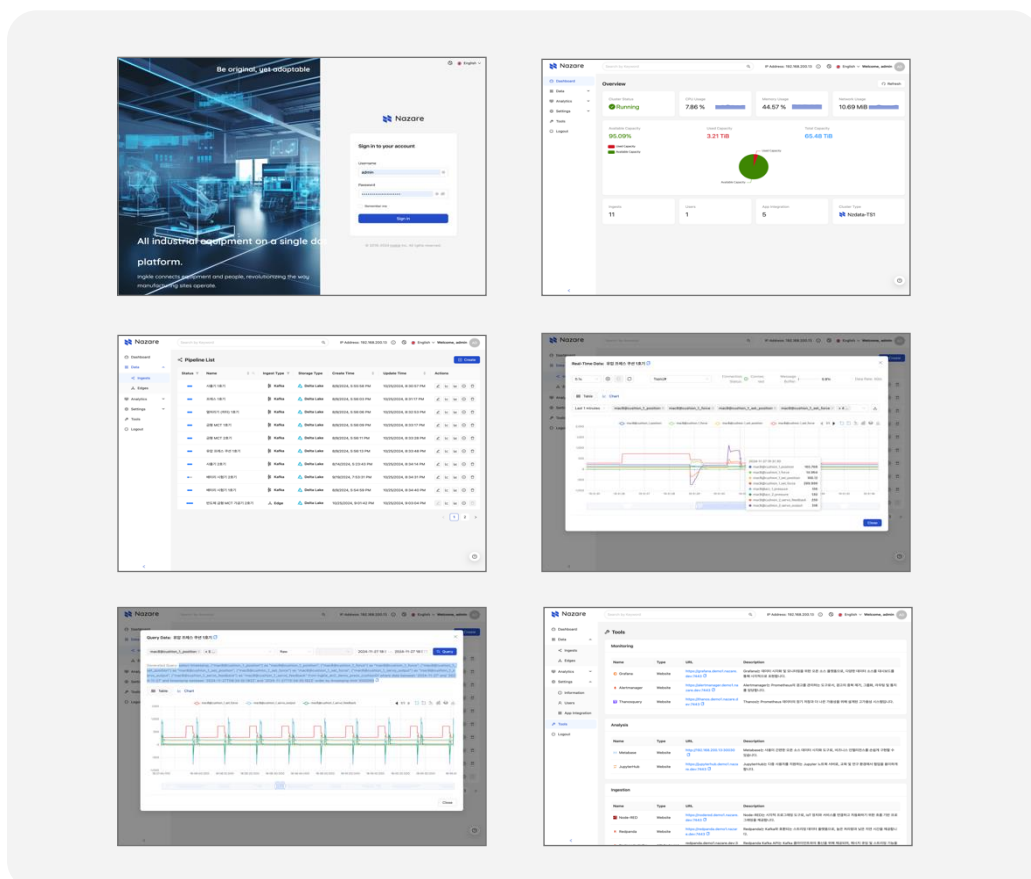

User-Friendly Management Console UI

Reasons to Choose **Nazare**
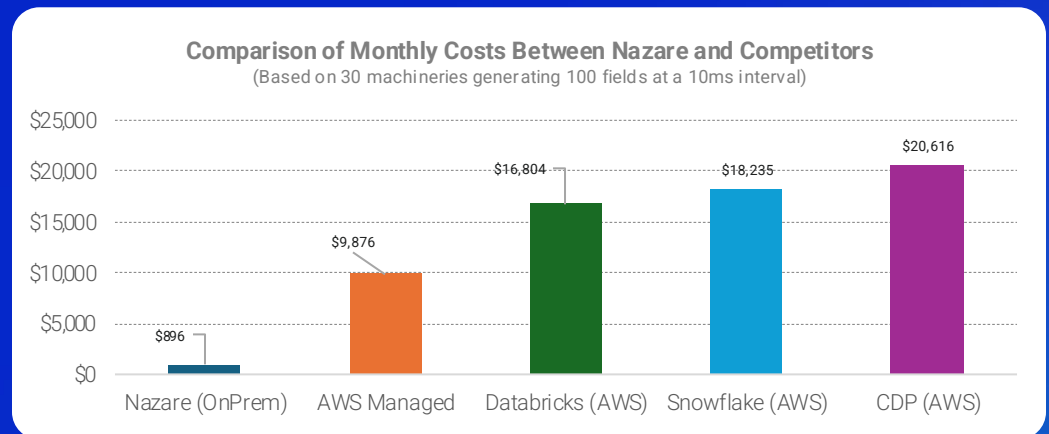
Nazare enhances your business competitiveness to a new level with integrated data management, cost savings, and powerful performance.

1
Unified Data Management

When companies introduce IT solutions from multiple vendors to improve factory production efficiency, redundancy and inefficiency often arise in data collection, storage, and management. This becomes a major cause of wasted time and costs. By implementing **Nazare**, companies can manage all machinery data from the manufacturing site in a single repository, eliminating data redundancy and inefficiencies, while enabling them to proactively asset and utilize their data.
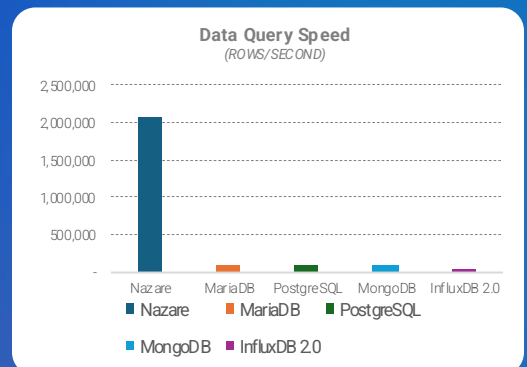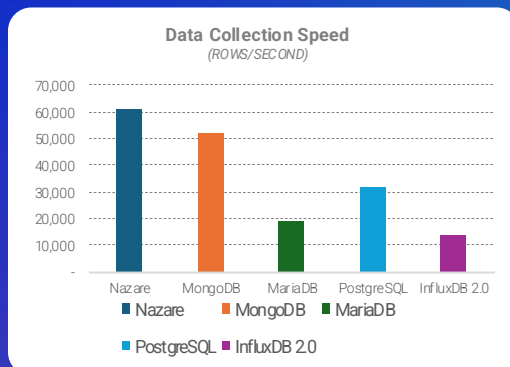
2
Cost-Efficient Operations

While public cloud is typically considered for handling big data, manufacturing companies often prefer on-premises solutions due to concerns about intellectual property leakage. However, traditional RDBs have storage limitations at the terabyte level, and their performance deteriorates as data accumulates. Hadoop, on the other hand, requires at least six months to set up and costs millions, with poor performance in real-time stream processing. In contrast, **Nazare** can be implemented within 7 days, costs more than 20 times less than cloud solutions, and offers a new alternative that solves the issues of both on-premises and cloud solutions.

Comparison of Monthly Costs Between Nazare and Competitors
(Based on 30 machineries generating 100 fields at a 10ms interval)

	Nazare (OnPrem)	AWS Managed	Databricks (AWS)	Snowflake (AWS)	CDP (AWS)
Cost	$896	$9,876	$16,804	$18,235	$20,616

3
Outstanding Data Processing Performance

Nazare guarantees consistently high performance for data collection (ingest) and querying, regardless of the data volume. This enables businesses to develop or integrate real-time applications and make faster data-driven decisions. Additionally, it allows for developing long-term strategies for real-time data utilization.

Data Collection Speed
(ROWS/SECOND)

	Nazare	MongoDB	MariaDB	PostgreSQL	InfluxDB 2.0

Legend: Nazare, MongoDB, MariaDB, PostgreSQL, InfluxDB 2.0

Data Query Speed
(ROWS/SECOND)

	Nazare	MariaDB	PostgreSQL	MongoDB	InfluxDB 2.0

Legend: Nazare, MariaDB, PostgreSQL, MongoDB, InfluxDB 2.0

※ TEST CRITERIA: 63 COLUMNS PER ROW

4
Maximized Memory Efficiency

Nazare demonstrates exceptional memory efficiency, with only 0.12 GiB of RSS memory usage for 100 fields collected at a 10ms interval throughout a 1-hour query. This is significantly lower than other database solutions such as MongoDB (3.15 GiB), MariaDB (1.22 GiB), and PostgreSQL (1.61 GiB). InfluxDB2, on the other hand, consumes a staggering 35.49 GiB. **Nazare**'s efficient memory usage allows businesses to handle large datasets without unnecessary resource consumption, ensuring optimal performance and cost savings.

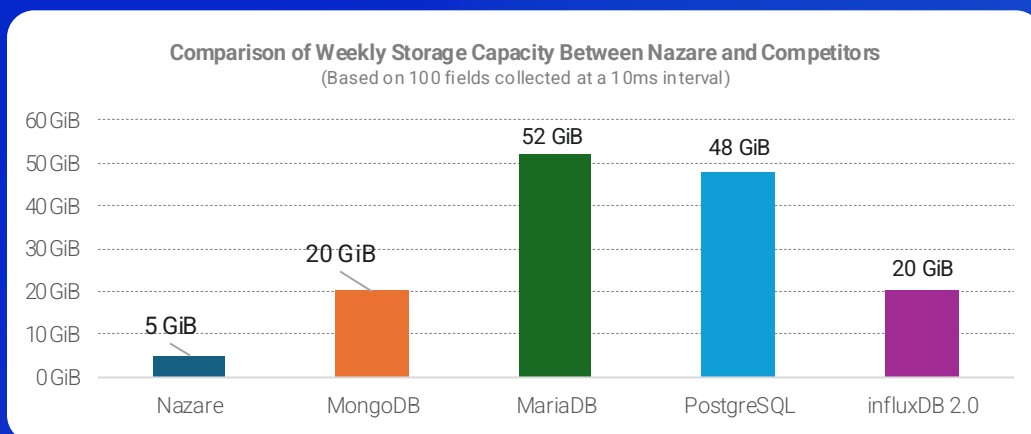**Comparison of Client Max RSS Memory Usage Between Nazare and Competitors**
(Based on 100 fields collected at a 10ms interval over 1 hour)

NazareDB	MongoDB	MariaDB	PostgreSQL	InfluxDB 2.0
0.12 GiB	3.15 GiB	1.22 GiB	1.61 GiB	35.49 GiB
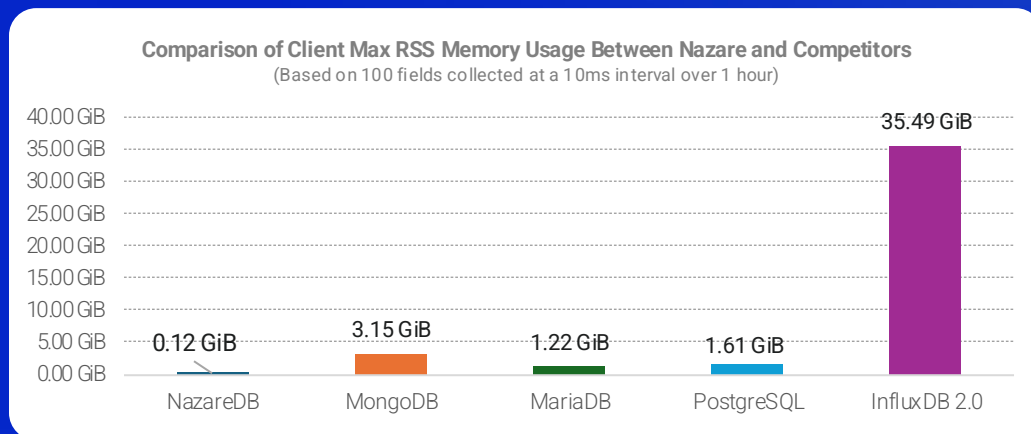
5
Efficient Data Storage

Nazare leverages a data lake architecture and stores data in a columnar Parquet format, which provides significant compression advantages over traditional RDS. This high compression ratio, combined with seamless scalability and fast data retrieval, makes Nazare an ideal choice for handling large-scale data efficiently. **Nazare** offers a cost-effective, high-performance system that simplifies real-time data management compared to traditional storage solutions.

Comparison of Weekly Storage Capacity Between Nazare and Competitors
(Based on 100 fields collected at a 10ms interval)

Nazare	MongoDB	MariaDB	PostgreSQL	influxDB 2.0
5 GiB	20 GiB	52 GiB	48 GiB	20 GiB

Expected Benefits of Implementation

Through the **Nazare** platform, companies can effectively address the diverse challenges they face and strengthen their competitiveness.

Manufacturer

- Real-Time Data Analysis, Process Optimization
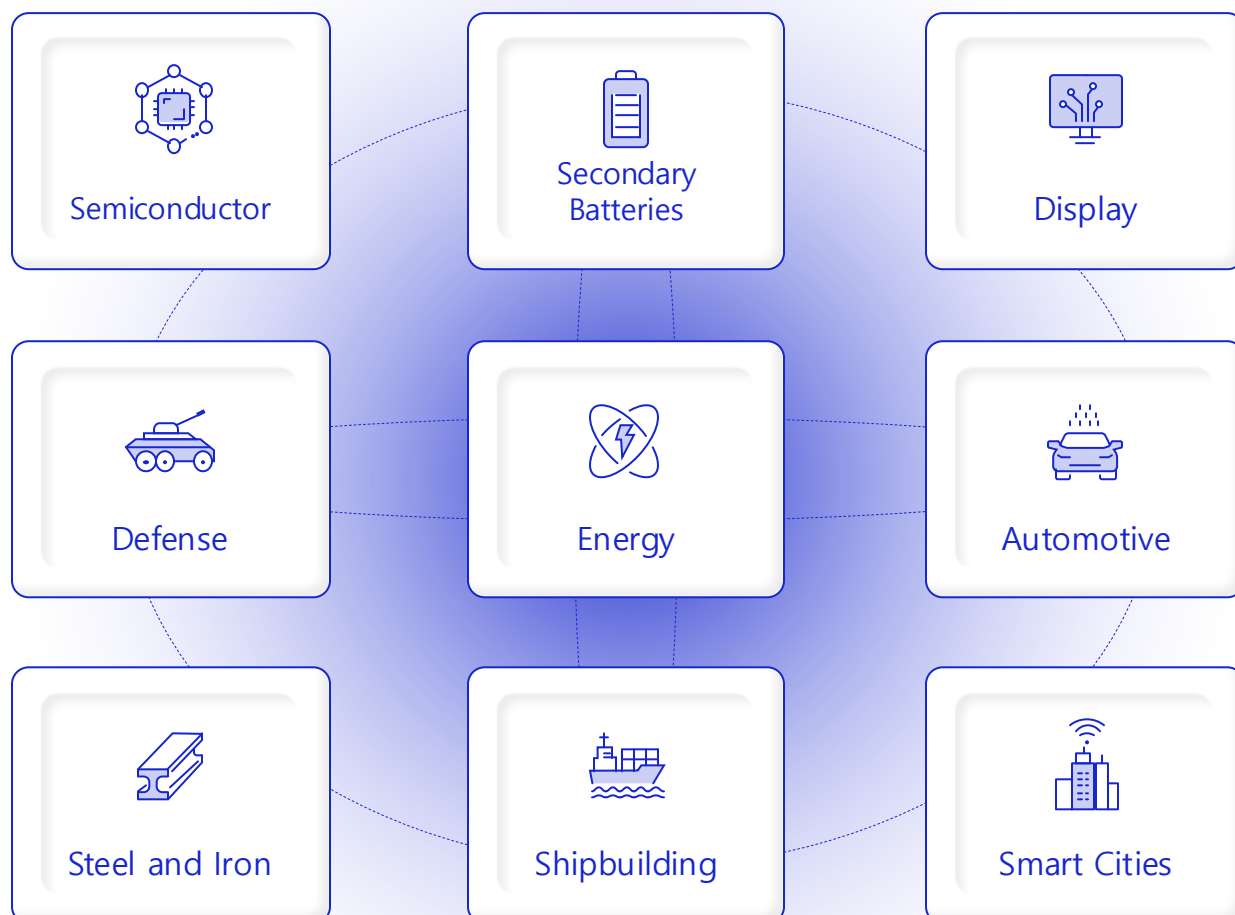- Strengthening Data-Driven Decision Making
- Improving Machinery Utilization and Maintenance Efficiency
- Internalizing a Cost-Effective Manufacturing Data Infrastructure

Machinery Manufacturer

- Improving Machinery Lifespan and Supporting R&D
- Enhancing Real-Time After-Sales Service
- Responding to Customer Data Requests

Nazare

AI Solution Provider

- Automating On-Site Data Collection and Preprocessing
- Reducing Time to Build Real-Time AI Inference Environment

Smart Factory Solution Provider

- Accelerating Solution Implementation Speed
- Standardizing Integration with Systems such as MES/ERP

Research Institute

- Integrated Management of R&D Research Data
- Promoting Data Sharing and Collaboration

Applicable Industries

Nazare is applied across various industries.

Semiconductor	Secondary Batteries	Display
Defense	Energy	Automotive
Steel and Iron	Shipbuilding	Smart Cities

Key Clients

- **Semiconductor:** MiCo, 21st Century
- **Secondary Batteries:** EcoPro BM, EcoPro HN, EcoPro EM
- **Defense:** Hanwha Systems
- **Automotive:** Renault (Headquarters), Seongwoo Hitec, JD Tech, SL Corporation - GM Changwon, Daeil Industrial
- **Steel and Iron:** Hyundai Steel, SeAH Steel, Hyundai RB, Husteel
- **Energy/Environment:** JB, KSG Technology Environment, Donghwa Enterprise, Donghwa Entech
- **Electronics:** Samsung Electronics, LG Electronics
- **Machinery Manufacturing:** Hyundai Rotem, Bosch Rexroth Korea, Parker Korea, Motrol, Hanyoung Nox
- **Research Institutes:** Defense Science Research Institute, Korea Institute of Industrial Technology, Korea Institute of Materials Science

Technical Specifications and Installation Requirements

Nazare data platform consists of an open-source framework, Ingkle's proprietary data processing engine, and a management console. Since **Nazare** stores data in the open big data format, Parquet, it allows for easy sharing and utilization of data across different systems.

S/W Technical Specifications

Data Ingestion	Redpanda, EMQX, NazareStreaming
Data Query	DataFusion, Arrow, FlightSQL, NazareDB
Data Storage	Deltalake, Ceph, Parquet, PostgreSQL
Data Tools	JupyterHub, Metabase, Grafana, Redash, Superset
Infrastructure	Cilium, MetalLB, Istio, Kubernetes, Linux
Operation	Loki, Grafana, Prometheus, Nzboard
Access Managing	Keycloak, Vault

※ ▬▬▬ Ingkle's proprietary core technology.

※ **Nazare** automatically deploys and configures all internal services using Ansible and Helm.

Minimum Recommended H/W Specifications

|	
CPU	12 Core
RAM	128 GB
Disk	1TB x 2 Disks (1 for OS, 1 for services like Kafka, MQTT, etc.) 1TB x 3 Disks (for data storage and disk redundancy)

Purchasing **Nazare**

Nazare can be accessed directly through Ingkle's appliance products, or it can be installed on servers from a manufacturer of your choice or a specific server product specified by your company. The **Nazare** platform can start with a single server.

※ Can be installed in on-premises, cloud, or hybrid environments.

Nazare

Supply Procedure

Interview → Server Selection → Delivery → On-Site Installation
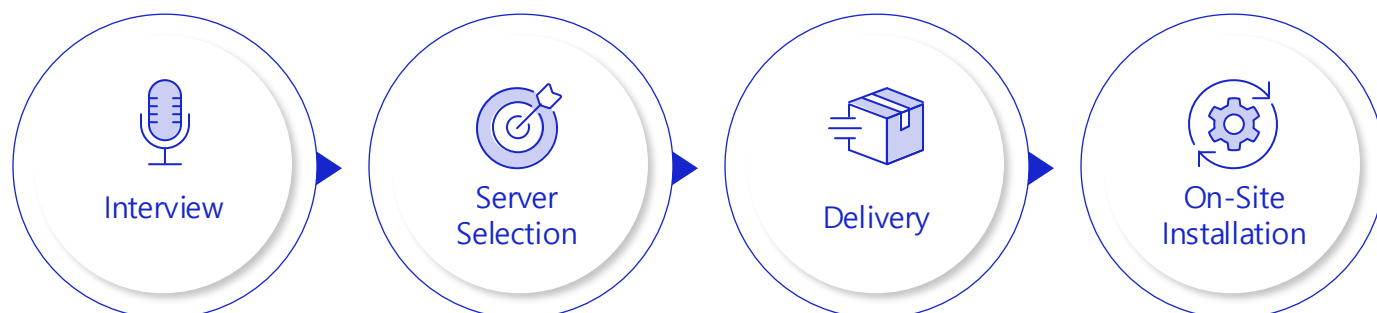
Server

Ingkle Appliance Server

AMD 16 Core, 4.5GHz
128 GB RAM
72 TB Disk
(Additional options available)

Nzdata TS-01

※ Nzdata is an Ingkle tower appliance server with **Nazare** pre-installed and optimized. Once you purchase Nzdata, it is ready for immediate use like an electronic product, allowing you to transform your manufacturing site into a big data environment within a day of receipt.

Customer-Specified Server

HP, Dell, Lenovo, and all other manufacturers

※ Supplied after **Nazare** installation and optimization are completed.

Configuration

Single Node Cluster

※ **Nazare** can be used with a single server. This setup is suitable for small-scale environments, where data collection, storage, and analysis are all performed on one server, making installation and management simple. It is ideal when you want to minimize initial setup costs while quickly processing data.

Multi-Node Cluster

※ **Nazare** can be used with multiple servers. This setup is suitable for manufacturing environments where large volumes of data are generated simultaneously from multiple machineries. The system automatically distributes processing and storage across several servers. Servers can be elastically added as needed to scale the infrastructure.

17

Technical Support

Ingkle is committed to providing the best technical support to help customers achieve rapid results with **Nazare**. Customers can receive fast and professional assistance for any technical issues that may arise before or after using the product.

DATA

When implementing Nazare, Ingkle actively supports various tasks related to the machinery. These include:

- Safely migrating machinery data from the existing database to Nazare.
- Directly setting up machinery connections through the Nazare management console.
- Developing collection driver software and integrating it with Nazare for collecting data from special machinery or systems.
- Configuring the PLC I/O address for communication between the machinery and Nazare.
- Installing gateway collection machineries to enable the connection between the machinery and Nazare.

INFRA

We solve various complex challenges related to big data infrastructure, operations, and integration with other systems.

- Building big data infrastructure using Nazare.
- Operating and managing big data infrastructure just like an internal IT team.
- Safely integrating existing data pipeline systems or preferred open-source technologies with Nazare.

SOLUTION

We provide solution software that performs specific functions required by customers using the Nazare data platform.

- Real-time machinery anomaly detection based on data trends (patterns) using HMI and SCADA.
- Machinery monitoring based on Large Language Model (LLM).
- Next-generation historian system that performs trend analysis and statistical aggregation at high speeds, leveraging Nazare's query response performance.

※ The cost of technical support may vary depending on the scope and content of the services and can be negotiated.

Our Company

At **Ingkle**, we provide software solutions that enable companies to build big data infrastructure centered around machinery quickly and easily. Our platform solves the limitations of traditional databases and cloud solutions, addressing the performance, security, and scalability needs of the manufacturing industry through big data processing technologies we have developed independently. With our cutting-edge technology, we aim to become a global leader in the industrial data platform market, driving customer growth and innovation.

Be Original,
Yet Adaptable

Our Mission

Our mission is to transform manufacturing sites by seamlessly connecting machinery and people.

'Be Original, Yet Adaptable" is **Ingkle**'s slogan. We aim to bring creativity to life by harmoniously combining it with practical realities, opening up new possibilities. We provide flexible, customer-tailored solutions, and together with our clients, we work to create a better future through continuous innovation and adaptation.

Our Values

Customer-Centric Decision Making

We make all decisions by prioritizing customer needs and addressing their challenges first.

Sustainable Growth

We are committed to long-term growth and development, focusing on the future, not just short-term business results.

Technology Innovation

We create new possibilities and open new markets through the development of innovative technologies.

Constant Challenge

We learn from our failures, pushing ourselves to always achieve our goals by continuously challenging the status quo.

Transparent Communication

We believe in sharing information openly to find and implement the best solutions collaboratively.

Trust and Cooperation

We build strong partnerships based on trust with our customers and partners to create successful outcomes.

ingkle

Headquarters:
Green Startup Town 3F, 32 Cheonancheon 4-gil, Seobuk-gu, Cheonan-si, Chungcheongnam-do, South Korea

Research Center:
#504, Building D, Digital Empire Building, 16, Deokyeong-daero 1556beon-gil, Yeongtong-gu, Suwon-si, Gyeonggi-do, South Korea