

## Solution Brief


- Critical features
- Core components
- Measurable benefits
- Technical specifications

# Make Your Industrial Data Useful

HighByte Intelligence Hub is an Industrial DataOps software solution designed specifically for industrial data orchestration, observability, quality, and governance.

DataOps is the orchestration of people, processes, and technology to securely deliver trusted, ready-to-use data to all who require it. HighByte Intelligence Hub provides industrial companies with an off-the-shelf DataOps solution to accelerate and scale the use of operational data throughout the enterprise by contextualizing, standardizing, and securing this valuable information.


Run the software at the Edge to merge and model telemetry, transactional, time series, and file data into a single payload. Observe and govern the delivery of contextualized, correlated information to all the applications that require it using a low-code application designed for scale. Together, we can make your industrial data useful for whatever digital transformation project comes your way, including generative and agentic AI use cases.



# Critical Features for Industrial DataOps

## CODELESS INTEGRATION

Collect and publish data over open standards and native connections—eliminating the need for custom-coded integrations. Easily configure and manage multiple connections and their respective inputs and outputs within the script-free interface. Collect data from SQL and REST source systems using dynamic requests leveraging inputs from other systems. Quickly integrate data from specialty systems and devices. Merge data from multiple systems into a complex modeled payload.

A screenshot of the HighByte Intelligence Hub interface. The left sidebar shows navigation options: Configure (Connections, Conditions, Models, Instances, Namespaces, Pipelines, Functions), Manage (Project, Users, Roles, Tags, Settings), and Tools (Log, UNS Client). The main area is titled 'Models' and shows a 'DiscoveryModel' card. The card has a 'Discovery' section with a 'Discovery' button. Below it are 'Overview' and 'Usage' tabs, with 'Overview' selected. The 'Details' section shows an 'Attributes' table with three rows: 'param', 'vibration', and 'temperature'. Each row has a delete button. Below the table are sections for 'Attribute Type' (Simple) and 'Simple Type' (Any). A 'Description' field is at the bottom. Top right of the card has 'Create Instance', 'Edit', 'Copy', and 'Delete' buttons.

## DATA CONDITIONING

Collect and condition raw input data and then pass conditioned data to model instances or pipelines. Filter data through a deadband condition to reduce the jitter in a source sensor or measurement. Filter the data through an aggregate to buffer higher resolution data and provide statistical calculations using average, min, max, count, and delta at a slower rate to characterize the specified time period.

Manipulate and transform raw input data into a usable format. Alarm on bad quality or stale data. Use the built-in transformation engine to standardize and normalize data for comparison and application mismatches. The transformation engine enables you to perform calculations, execute logic to define new “virtual property” values, and decompose complex strings at the Edge to improve data usability and reduce transmission volume.

Define global JavaScript functions or load third-party JavaScript or Node packages, then use them in any expression within the Intelligence Hub.


## DATA MODELING

Represent machines, products, processes, and systems with intelligent data models suited to your needs. Contextualize thousands of industrial data points by merging them with information from other systems, adding meta data, standardizing data attribute names and lists, and normalizing units of measure. Model hundreds of common assets in minutes with templated inputs and instances and manage models through an intuitive attribute tree that enables model nesting.

# Critical Features for Industrial DataOps

## DATA ORCHESTRATION

Create data flows for raw data, modeled information, or files between connections on an interval or event basis. Enable store and forward to buffer data to disk if target connection is lost. Manage data pipelines within HighByte Intelligence Hub, monitoring state and processing metrics at each step. Use Pipeline Debug to compare payloads, statistics, and state side-by-side. See and be alerted to connection failures and easily monitor the Intelligence Hub at scale using third-party observability platforms.



## DATA PROCESSING

Use the graphical Pipelines builder to curate complex data payloads for everything from MQTT brokers to historians to data warehouses and track the transformation of data through the pipeline. Use stages to read, filter, buffer, transform, format, and compress payloads. Read within and across Namespace hierarchies using the Smart Query stage. Use the On Change stage to enable event-based delivery and report-by-exception of any data source. Use the Switch stage to introduce conditional logic to your flow and the Model Validation stage to assess incoming data payloads against a model definition. Share data processing on-demand by calling between pipelines or securely exposing pipelines to the REST Data Server.

# Critical Features for Industrial DataOps

## EDGE DEPLOYMENT

Run HighByte Intelligence Hub on your choice of light weight hardware platforms including single board computers, industrial switches, IoT gateways, and industrial data servers at the Edge. Deploy as an individual software installation or Docker image to rapidly deploy and upgrade system software components.

## REST DATA SERVER

The REST Data Server acts as an API gateway for industrial data residing in OT systems, so any application or service with an HTTP client can securely request OT data in raw or modeled form directly from the Intelligence Hub—without requiring domain knowledge of the underlying systems. This API securely exposes the Intelligence Hub's connections, models, instances, and pipelines as well as the underlying values. Use the REST Data Server with Callable Pipelines to gain an “Industrial Data API Builder.” Connect to the REST Data Server to access the Intelligence Hub as a transactional, request-and-response interface and programmatically browse the full Industrial DataOps infrastructure.

## INDUSTRIAL AI

Contextualized industrial data from Intelligence Hub is accessible through the Industrial MCP Server, which exposes pipelines as “tools” to be used in agentic AI workflows. With the industrial MCP Server, AI use cases can access all connected industrial systems including devices, SCADA, MES, Historians, ERP, and LIMS, and make real-time or historical data requests on them. AI can also be used to scale DataOps efficiency through connections with Amazon Bedrock, Azure OpenAI, Google Gemini, OpenAI, and local LLMs.


## UNS CLIENT

Use the UNS client to visually discover and interrogate contents residing in any MQTT broker, negating the need for external testing clients. Simply select a connection and instantly visualize the namespace. The UNS Client can automatically detect and visualize message payloads including JSON, Sparkplug, text, and raw binary, and decode Protobuf to make payloads human readable for Sparkplug users. In addition to topic and message inspection, the UNS Client can also publish messages to topics.

# Critical Features for Industrial DataOps

## MQTT BROKER

Use the embedded MQTT broker that's tightly integrated with the Intelligence Hub's core data integration and contextualization capabilities to design a namespace that provides a real-time view of the state of business. The MQTT v3.1.1 and v5 compliant broker was built from the ground up by HighByte. It supports both JSON and Sparkplug payloads and can be quickly enabled in the UI. The broker is a critical component to rapidly building a local unified namespace (UNS) inside the factory.



The screenshot shows the HighByte Intelligence Hub interface. On the left, the navigation menu includes 'Configure' (Connections, Conditions, Models, Instances, Namespaces, Pipelines, Functions), 'Manage' (Project, Users, Roles, Tags, Settings), and 'Tools' (Log, UNS Client). The 'Namespaces' section is expanded, showing a hierarchy: Site1, Site2, Site3, Area1, and a selected namespace 'Injection'. The 'Details' panel for 'Injection' shows the name 'Injection', description 'Injection process - Material "a"', and a reference to 'Injection' instance. The 'Smart Query Test' panel on the right shows a query: 'Site3/Area1/Line1/Injection/\*\*'. The 'Where' section is empty. The 'Select' section shows a path 'Injection'. The 'As' section is set to 'Hierarchy'. The results pane displays a JSON structure representing the namespace hierarchy: 'Site3' with 'Area1' and 'Line1' under 'Injection', each with 'Pressure', 'Temperature', and 'Cycles' fields.

## NAMESPACES

Namespaces provides a dedicated space to visually organize datasets and their relationships. Use Namespaces to design and govern a Unified Namespace (UNS) or any destination system with a hierarchical namespace. Catalog data sources and modeled data sets into a logical hierarchy for seamless integration with many data consuming entities. Use Namespaces to take a modeled, contextualized representation of your operations and use this as a basis to drive the topic namespace in an MQTT broker, asset hierarchy in a historian, tag structure in a SCADA or IIoT platform, and more.

# Core Components for Enterprise Administration

## SECURITY

Exchange data using the built-in security of connection protocols. Authenticate users and their roles through third-party identity providers with Security Assertion Markup Language (SAML) and Windows Active Directory (LDAP). Securely move project configurations across environments with different connection credentials using external secrets support and the dedicated configuration construct for secure secret referencing.

## CENTRAL MANAGEMENT

Connect multiple hubs to a single host that acts as the central hub. Once connected, administrators can log in to the central hub and easily switch between hubs to configure and monitor individual hub activity and compare configurations for differences. Administrators can also synchronize models, connections, or complete projects between hubs. The configuration is truly portable. Administrators can import and export JSON-based configuration files directly within the browser-based configuration UI.


## PERMISSIONS

Create unique user names and passwords for each user. Assign a user to a role with a pre-defined set of permission claims or assign a user their own unique permission claims. Use Active Directory to manage authentication, authorization of users, and application settings. Create and maintain certificates in the hub configuration to authenticate and secure data transfer with other systems.

## AUDITING

Access the change management, traceability, and accountability required in heavily regulated industries. Audit logging allows all configuration creations, modifications, or deletions to be logged to the event log, including any cascading changes on related configuration. View all log events in full detail through the configuration and filter by type, source, or message text. Automatically back up the runtime's configuration file to a backup directory at a specified frequency and maintain a specified maximum number of these encrypted files.

# Core Components for Enterprise Administration



## HIGH AVAILABILITY

Deploy redundant Intelligence Hubs for critical data and pair the hubs as a primary and secondary. A secondary hub will be ready in a warm state ready to enable flows when it detects the primary is no longer responding.

## PROJECT MANAGEMENT

Automate the backup, version control, and deployment of an Intelligence Hub instance configuration from one or more Git repositories to streamline DevOps workflows. Using Git with the Intelligence Hub provides distributed version control of configurations, ranging from single sites to the enterprise.

Within the Intelligence Hub, use tags to define ad hoc collections of items in the configuration. These tags can easily be used to identify configuration related to a use case or a section of the factory or to limit permissions to specific people. The tags enable easy grouping to filter the UI, synchronize across hubs, and grant permissions within a hub.

## Measurable Benefits for Operational Technology (OT), IT, and Line of Business

Accelerate analytics and other Industry 4.0 use cases with a digital infrastructure solution built for scale.

- ✓ Reduce system integration time from months to hours
- ✓ Improve data curation and preparation for AI and ML applications
- ✓ Scale operations metrics and analytics across the enterprise
- ✓ Reduce information wait time for business functions
- ✓ Eliminate time spent troubleshooting broken integrations
- ✓ Empower operators with insights from the Cloud
- ✓ Improve system-wide security and data governance
- ✓ Meet system integrity and regulatory traceability requirements
- ✓ Reduce Cloud ingest, processing, and storage costs and complexity
- ✓ Optimize data payloads for specific target applications and use cases

# Technical Specifications

## CONNECTIVITY

Connector	Inbound	Outbound
<b>AI</b>		
Amazon Bedrock	✓	✓
Azure AI	✓	✓
Google Gemini	✓	✓
OpenAI	✓	✓
<b>AWS</b>		
Amazon Data Firehose		✓
Amazon Kinesis Data Streams		✓
Amazon Redshift	✓	✓
Amazon S3	✓	✓
Amazon S3 Tables		✓
AWS IoT SiteWise		✓
<b>Azure</b>		
Azure Blob Storage		✓
Azure Event Hubs	✓	✓
Azure IoT Edge	✓	✓
Azure IoT Hub	✓	✓
<b>Databricks</b>		
Databricks Client		✓
<b>Files</b>		
Apache Parquet	✓	✓
CSV	✓	✓
File	✓	✓
<b>Google Cloud Platform</b>		
Google Cloud Pub/Sub		✓
Google BigQuery	✓	✓
<b>HTTP</b>		
REST Client	✓	✓
Webhook	✓	
<b>Inductive Automation</b>		
Ignition Module	✓	✓

# Technical Specifications

## CONNECTIVITY

Connector		Inbound	Outbound
<b>Modbus</b>			
Modbus TCP		✓	✓
<b>MQTT</b>			
MQTT		✓	✓
Sparkplug		✓	✓
<b>OPC</b>			
OPC UA TCP		✓	✓
<b>Snowflake Data Cloud</b>			
Snowflake Streaming			✓
Snowflake SQL		✓	✓
<b>SQL</b>			
JDBC Driver		✓	✓
Microsoft SQL Server		✓	✓
MySQL		✓	✓
Oracle Database		✓	✓
PostgreSQL		✓	✓
SQLite		✓	✓
<b>Streaming</b>			
Apache Kafka		✓	✓
<b>Time Series</b>			
Aspen InfoPlus.21		✓	
InfluxDB		✓	✓
PI System		✓	✓
TimescaleDB		✓	✓



HighByte Intelligence Hub can also connect bi-directionally to AWS IoT Core and AWS IoT Greengrass through the MQTT connector.

# Technical Specifications

## SUPPORTING OPERATING SYSTEMS

- Windows Server 2016/2019/2022/2025
- Windows 10/11
- Linux [any Linux distribution capable of running a JVM; tested with Ubuntu]
- macOS

## SYSTEM REQUIREMENTS

- Java SE 21 or OpenJDK 21 [or newer]
- HTTP server [for hosting frontend]
- 1.4 GHz processor
- 1 GB RAM
- 1 GB available disk space
- Network capable [TCP]

These are minimal system requirements. Actual requirements will vary based on product configuration.

## SOFTWARE DELIVERY

HighByte Intelligence Hub is an on-premises application configured remotely through a web browser or a REST-based API. The software is available as an annual subscription. Please visit [highbyte.com/pricing](https://highbyte.com/pricing) to view licensing options.

## Next Steps

Interested in learning more? Please contact [sales@highbyte.com](mailto:sales@highbyte.com) to request additional information, schedule a demo, or join our free trial program.

## About HighByte

HighByte is an industrial software company in Portland, Maine USA building solutions that address the data architecture and integration challenges created by Industry 4.0. HighByte Intelligence Hub, the company's award-winning Industrial DataOps software, provides modeled, ready-to-use data to the Cloud using a codeless interface to speed integration time and accelerate analytics. Learn more at [highbyte.com](http://highbyte.com).



© 2025 HighByte, Inc. All rights reserved.

HighByte is a registered trademark of HighByte, Inc.