

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 an off-the-shelf Industrial DataOps solution to accelerate and scale the use of your 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.

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 pipelines 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 pipeline and the Model Validation stage to assess incoming data payloads against a model definition. Iterate through data sets or paginate through APIs until specified conditions are met with the For Each and While stages. Reuse and share data processing routines using the Subpipeline stage.

Critical Features for Industrial DataOps

FLEXIBLE 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. The Intelligence Hub can also be deployed in on-premises data centers or in the cloud. Deploy the software on bare metal, a virtual machine, or as a container.

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

Use DataOps for AI and AI for DataOps in HighByte Intelligence Hub. Build pipelines that can be exposed as custom MCP "tools" for agentic AI workflows. With the Industrial MCP Server, AI agents can access all connected industrial systems in the Intelligence Hub and make real-time or historical data requests on them. AI can also be used to scale Industrial DataOps. Use an LLM of your choice, such as Amazon Bedrock, Azure OpenAI, or Google Gemini, to map disparate data sources into data models and accelerate data contextualization. Conversationally build and debug configuration with the embedded Modeling and Pipeline AI Agent or through external agents via MCP configuration tools.


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.



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 hierarchal namespace. Catalog data sources and modeled datasets 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 to drive the topic namespace in an MQTT broker, asset hierarchy in a historian, tag structure in a SCADA or IIoT platform, and more. Namespaces within a remote hub can be federated into a central hub to build data pipelines that draw from distributed sources and serve enterprise applications and services.

Core Components for Enterprise Administration

SECURITY

The Intelligence Hub applies security at three points: the connections that carry data, the users who access it, and the credentials that authorize it. 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 2.0 (SAML 2.0), Windows Active Directory (LDAP), and OpenID Connect (OIDC). Securely move project configurations across environments with different connection credentials using external secrets support and the dedicated configuration construct for secure secret referencing.

CENTRAL HUB

A central hub offers both configuration management and data exchange with remote hubs, simplifying distributed architectures. Administrators can connect multiple hubs to a single host that acts as the central hub. Once connected, they 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 see the namespace of a remote hub and map its nodes into a central namespace. A central hub can browse that namespace, run Smart Queries, and build pipelines that subscribe to and draw from remote data sources as if they were local. This approach simplifies distributed deployments and makes modeled factory data discoverable and actionable on demand.

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 multiple Intelligence Hub runtimes in High Availability (HA) mode to ensure reliable and resilient industrial data operations. All hubs stay synchronized with the latest project configuration and state. If the primary hub goes offline, a secondary hub takes over seamlessly. Active Pipelines resume in their current state, minimizing the risk of data loss.

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. Configuration is also highly portable, offering the ability to import and export JSON-based configuration files within the browser-based configuration UI.

In the Intelligence Hub, you can 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 integration, analytics, and AI use cases with a digital infrastructure solution built for scale.

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

Technical Specifications

CONNECTIVITY

Connector	Inbound	Outbound
AI		
Amazon Bedrock	✓	✓
Anthropic	✓	✓
Azure OpenAI	✓	✓
Google Gemini	✓	✓
OpenAI	✓	✓
AWS		
Amazon Data Firehose		✓
Amazon Kinesis Data Streams		✓
Amazon Redshift	✓	✓
Amazon S3	✓	✓
Amazon S3 Tables		✓
AWS IoT SiteWise		✓
Databricks		
Databricks SQL	✓	
Databricks Storage		✓
Databricks Zerobus		✓
Files		
Apache Parquet	✓	✓
CSV	✓	✓
File	✓	✓
Google		
Google Cloud PubSub		✓
Google BigQuery	✓	✓
HTTP		
Industrial Information eXchange (i3X)	✓	
Model Context Protocol (MCP)	✓	✓
REST Client	✓	✓
Webhook	✓	
Inductive Automation		
Ignition Module	✓	✓

Technical Specifications

CONNECTIVITY

Connector	Inbound	Outbound
Microsoft Azure		
Azure Blob Storage	✓	✓
Azure Event Hubs	✓	✓
Azure IoT Edge	✓	✓
Azure IoT Hub	✓	✓
Microsoft OneLake		✓
Modbus		
Modbus TCP	✓	✓
OPC		
OPC UA TCP	✓	✓
Snowflake		
Snowflake Streaming		✓
Snowflake SQL	✓	✓
SQL		
JDBC Driver	✓	✓
Microsoft SQL Server	✓	✓
MySQL	✓	✓
Oracle Database	✓	✓
PostgreSQL	✓	✓
SQLite	✓	✓
Streaming & Messaging		
Apache Kafka	✓	✓
MQTT	✓	✓
NATS	✓	✓
Sparkplug	✓	✓
Time Series		
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 25 (LTS) or OpenJDK 25 (LTS)
- ✓ 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 licensed as an annual subscription and deployed on bare metal, a virtual machine, or as a container. Please visit highbyte.com/pricing to view licensing options.



Next Steps

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

About HighByte

HighByte is an industrial software company addressing the data architecture and integration challenges faced by global manufacturers as they digitally transform. HighByte Intelligence Hub, the company's proven Industrial DataOps software, provides modeled, ready-to-use data to the Cloud using a codeless interface to speed integration time and accelerate AI use cases. Learn more at highbyte.com.





© 2026 HighByte, Inc. All rights reserved.

HighByte is a registered trademark of HighByte, Inc.