SOLUTION BRIEF

INDUSTRIAL DATA FABRIC
MOBILIZE DATA FROM INDUSTRIAL ASSETS AND SYSTEMS AT SCALE

Industrial Data Fabric (IDF) is a reference architecture developed by AWS with a supporting ecosystem of partners to enable the convergence of IT/OT data in modern manufacturing environments. The architecture spans from edge to cloud to democratize operations data across the enterprise.

IDF removes the heavy lifting associated with infrastructure management and maintenance so customers can focus on deploying use cases and driving business value. IDF provides the foundation for industrial companies to streamline manufacturing and supply chain operations and scale Industry 4.0 use cases like predictive asset maintenance and enterprise OEE.

Snowflake and HighByte have invested in technical integrations under the IDF framework alongside Snowflake's solution accelerators, including OEE Insights Framework, Anomaly Detection, and Energy Optimization to accelerate customer adoption of Smart Manufacturing initiatives.

As a component of the IDF framework, HighByte Intelligence Hub offers a low-code / no-code approach to curating, modeling, and contextualizing machine, transactional, and time series (historical) data from various edge data sources into a single payload. The Intelligence Hub provides OT data within Snowflake's highly efficient table format.

IDF reduces the total cost of ownership for managing industrial data at scale while accelerating innovation for manufacturing and supply chain operations including analytics and AI.

Predictive Maintenance

Production Metrics

Quality

Traceability

Sustainability

Supply Chain

SNOWFLAKE FOR MANUFACTURING

Expand supply chain visibility for manufacturers, power Industry 4.0 smart manufacturing initiatives, and unlock insights from IoT, OT, or time series data.

POWER INDUSTRY 4.0 AND SMART MANUFACTURING ANALYTICS

Accelerate Industry 4.0 initiatives with Snowflakes elastic scale and native support for semi-structured, structured, and unstructured high-volume manufacturing data.

IMPROVE SUPPLY CHAIN RESILIENCY

Enhance supply chain visibility through data collaboration with key partners, improving product quality and on-time delivery with seamless and secure cross-cloud data sharing.

ENHANCE CUSTOMER EXPERIENCE AND CREATE NEW REVENUE STREAMS

Gain access to high-volume, connected product (IoT) data and service history for optimal product performance. Create the next generation of products and services with new digital feedback loops and insights mined from IoT and consumer-insight data.

MEETING THE CUSTOMER WHERE THEY ARE IN THEIR TRANSFORMATION JOURNEY

In cases of customers using multiple hyperscalers, Snowflake will remain the single source of truth for IT / OT data in manufacturing and other domains like energy to drive insights.

SNOWFLAKE-FIRST MODEL

All data is ingested into Snowflake directly, blended into IT/OT contextualized datasets in Snowflake or S3 or any AWS datastore for further analytics. Snowflake is the Single Source of Truth for Industrial Data. Takes full advantage of the connectors in Snowflake. Snowflake's Snowpipe streaming is designed for ingesting time-series data with near real-time latencies. Working in conjunction with Amazon MSK and Kinesis, Snowpipe streaming can help data workers achieve an attractive cost/latency profile in a variety of use cases such as IoT, CDC for OLTP, and SIEM.

S3 DATA-LAKE FIRST MODEL

All raw data is loaded into Amazon S3 first and then loaded into Snowflake. S3 is the Single Source of Truth. Snowflake is used as an industrial analytics application and can use open-source iceberg table format to access the data in S3 buckets.

HYBRID MODEL

Snowflake ingests large datasets (e.g. SaaS data) and other high-volume data sources (e.g., IoT data) that are loaded into Amazon S3 first before pre-processed data is loaded into Snowflake. You can create a unified Industrial Data Catalog to act as a Single Source of truth. Use Snowpark to write code in your preferred language and run that code directly on Snowflake.

IDF Solution using HighByte Intelligence Hub, the Snowflake Data Cloud, and AWS services

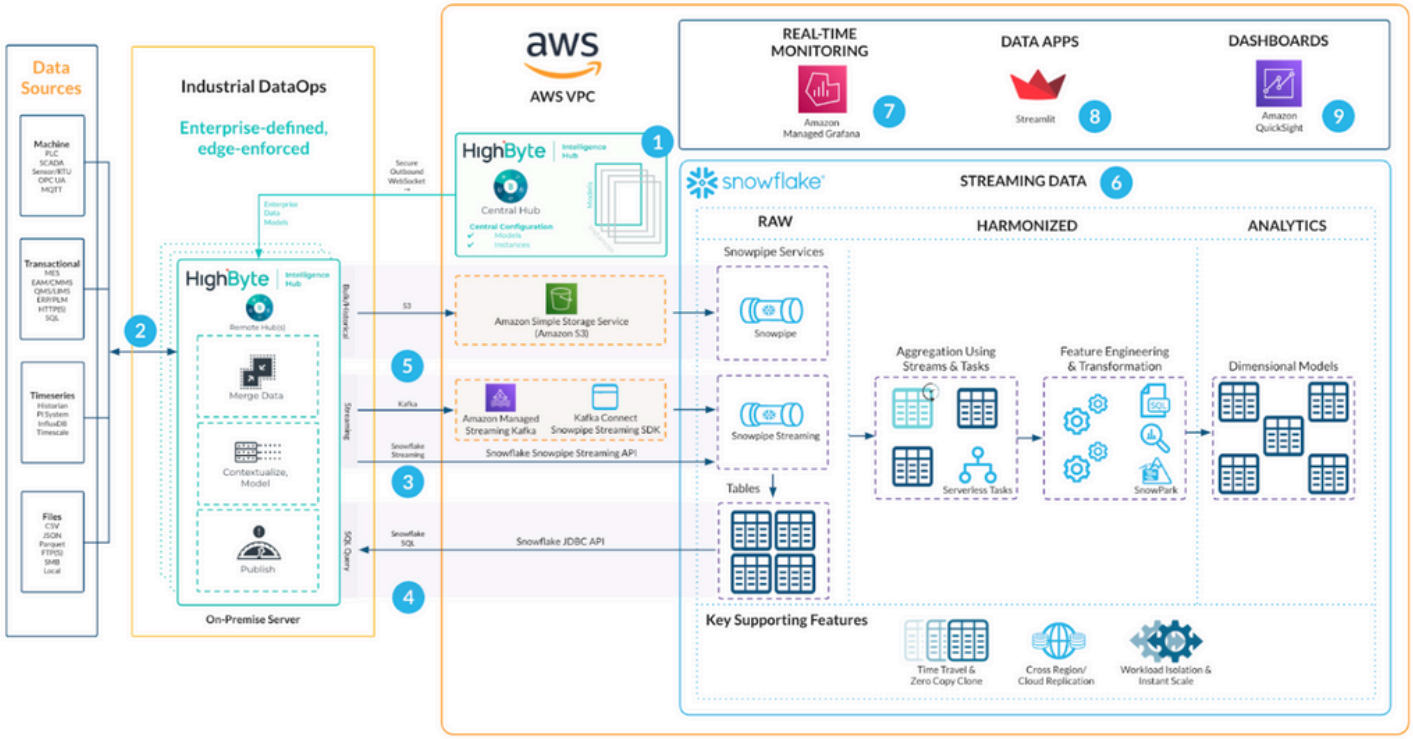

DATA-PROCESS FLOW

1. Enterprise standardized and contextualized data models are defined in central Intelligence Hub and pushed to remote on-premises Hubs at each site.
2. HighByte ingests OT data from industrial data sources including PI, Kepware, Ignition, MES, SCADA, bespoke databases etc.
3. HighByte includes a low-latency streaming connection direct to Snowpipe. Connection creates tables and streams data.
4. HighByte's JDBC-based connection to Snowflake provides direct table access from Snowflake to OT systems.
5. The Intelligence Hub also has out of the box connectivity to Amazon S3 for large historical datasets and Kafka for streaming datasets where the customer already has Kafka infrastructure in place.
6. Data gets transformed using Snowpark to leverage the power of Java or Python. Data can then be aggregated and prepared using SQL for analytics.
7. Real time dashboards on streaming data with Amazon Managed Grafana.
8. Streamlit Apps can be used to display real time dashboards from the raw data zone, as well as historical analysis from the Analytics zone.
9. Data gets analyzed and visualized using separate compute and BI tools like Amazon QuickSight.

Smart Manufacturing Strategy for Global Leader in Food & Agriculture

- Leading agriculture company implemented the AWS and Snowflake solution as part of their Scaling Smart Manufacturing initiative
- Company had more than 800 datasets across 30 plants, 798 asset twins built, and more than 510,000 tags consumed
- Solution adoption resulted in Operational gains at each plant that were realized in 2-4 weeks vs. 6 months.

ABOUT SNOWFLAKE

Snowflake enables every organization to mobilize their data with Snowflake's Data Cloud. Customers use the Data Cloud to unite siloed data, discover and securely share data, and execute diverse analytic workloads. Wherever data or users live, Snowflake delivers a single data experience that spans multiple clouds and geographies. Thousands of customers across many industries, including 543 of the 2022 Forbes Global 2000 (G2K) as of October 31, 2022, use Snowflake Data Cloud to power their businesses. Learn more at snowflake.com.

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.