Integrating Mainframe Data with Big Data Platforms

Unlocking the value of mainframe data and offloading queries to reduce MIPS consumption

Mainframes have been used for over 50 years by large enterprises to manage valuable and sensitive data. From order processing to financial transactions, production and inventory control, payroll and beyond, mainframes continue to support mission-critical applications. The mainframe will remain relevant for the foreseeable future.

But today's enterprises need to integrate mainframe data into modern, data-driven, analytical business processes and the environments that support them. A brute-force approach of directly querying the mainframe can be prove costly because mainframe system billing is based on MIPS (millions of instructions per second.) And many enterprises want to leverage mainframe data continuously for business analysis.

By offloading mainframe data to modern platforms such as Apache Kafka, Apache® Hadoop™ and MPP databases, enterprises can create new analytics possibilities and insights. Integrating with these new environments requires a fresh approach that keeps the data current and available without additional complexity or prohibitive cost.

Attunity Replicate Benefits:

- Accelerate mainframe data analytics
- Address new analytics use cases
- Reduce IT and development resource requirements

Bring Mainframe Data to Life in Hadoop and Kafka

Attunity Replicate software provides an efficient and cost effective way to bring mainframe data into a modern analytics environment using change data capture (CDC) technology for DB2, VSAM, and IMS. Attunity Replicate provides low-latency and low-impact data integration for mainframe databases. With Attunity Replicate, you can extract mainframe data efficiently in real-time and deliver it to a data warehouse, Hadoop data lakes or Apache Kafka.

Hadoop data lakes are the modern enterprise platform for storing large volumes of diverse data, and provide the foundation for modern Big Data analytics. With Attunity Replicate, enterprises can ingest data from many data feeds with no need for custom development. Attunity Replicate can offload large volumes of mainframe data at high speed to Hadoop, with certified integration with all major Hadoop distributions (including Cloudera, Hortonworks and MapR).

This enables enterprises to address compelling new analytics opportunities as they analyze mainframe transactions or other data alongside new structured and unstructured data sets in the data lake.

More and more enterprises also are using Apache Kafka for high-scale, lowlatency ingestion and processing of live data streams. But Kafka's impact on source systems, the complexities associated with custom development, and the need for real-time streaming at scale from many different data sources can be challenging.

With Attunity Replicate, enterprises can enable real-time analytics with live data from mainframe databases by leveraging low-impact CDC technology. They can replicate data to Kafka message brokers, thereby converting mainframe transactions into real-time message streams to one or more Big Data platforms to address a wide range of new analytics use cases.

Whether an enterprise's unique business needs demand direct data ingestion into Hadoop or streaming data ingestion through Kafka, Attunity Replicate delivers several critical benefits.

- **Support for many source systems with low impact:** Attunity Replicate is a single platform that supports many types of legacy sources, including DB2 z/OS, IMS and VSAM. This CDC technology identifies changes by scanning logs, reducing the need to consume costly MF resources.
- **Simplicity (no coding required):** Attunity Replicate enables easy configuration through a wizard-based GUI. Enterprises can quickly set up data feeds without manual coding.
- **High scale for data ingestion/streaming:** Attunity Replicate scales to ingest data from hundreds or thousands of databases, with centralized monitoring and management.
- **Real-time data capture for Kafka:** Attunity Replicate feeds live database changes to Kafka message brokers with low latency, enabling enterprises to broadcast data streams concurrently to multiple Big Data targets.
- **Continuous updates into Big Data environments:** Attunity Replicate ingests incremental datasets continuously with enterprise-class change data capture (CDC) from many transactional sources, enabling major data integration projects.

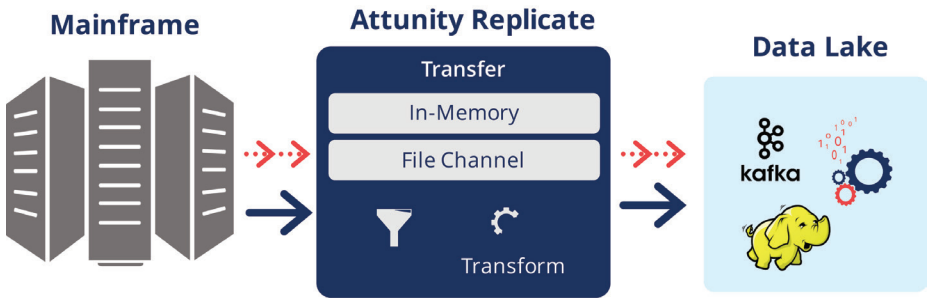

Attunity Replicate Use Cases

- Convert mainframe data into message stream
- Manage thousands of data feeds going into one or more Big Data platforms
- Feed data into Hadoop for analysis with other data types

Americas
866-288-8648
sales@attunity.com

Europe / Middle East / Africa
44 (0) 1932-895024
info-uk@attunity.com

Asia Pacific
(852) 2756-9233
info-hk@attunity.com

attunity.com
@attunity