



Accelerate Your Data Pipeline




The fabric of business decisions is changing. Historical transaction reporting and batch analyses of data at rest are giving way to predictive data science, streaming analytics, machine learning and AI. Enterprises must embrace these new paradigms to make more informed, real time decisions to tune operations, win customers and capitalize on new opportunities. Your peers and competitors are making this shift.

Data lake, streaming and cloud platforms have emerged as the modern data architectures of choice. But they require a scale, speed and variety of data that break traditional integration tools. These tools are slow and require manually coded commands that tie up your best programmers.

There is a new way. Attunity Replicate is the foundation for modern, efficient data pipelines, providing automated, real time and universal data integration across all major data lakes, streaming systems, databases, data warehouses and mainframe systems, on premises and in the cloud. With Attunity Replicate, you can deliver more data to the business, ready for agile analytics, while reducing effort and cost.

Attunity Replicate moves data at high speed from source to target, all managed through a simple graphical interface that completely automates end-to-end replication. With streamlined and agentless configuration, administrators and data architects can easily set up, control, and monitor bulk loads and real-time updates with enterprise-class change data capture (CDC). Using Attunity Replicate you gain global visibility into and centralized control of data replication across distributed, hybrid environments.



BENEFITS

- Faster migrations and analytics projects
- Reduced skill and training requirements
- Increased agility to meet business needs

ENTERPRISE USE CASES

- Data Lake ingestion
- Data streaming with Kafka
- Mainframe data modernization with Big Data Platforms
- Real-time Data Warehousing
- Query offload and live reporting
- Enabling cloud analytics

KEY CAPABILITIES

- Advanced change data technology (CDC) addressing transactional, streaming and batch architectures
- Centralized and simple "Click to Load" interface
- Automated end-to-end data replication
- Agentless connections to end points throughout your environment
- Industry's broadest platform support, covering all major RDBMS and legacy systems, data warehouse, Hadoop and NoSQL platforms on premises and in the cloud



Ease of use and automation

An intuitive web-based console makes it simple to configure, control and monitor replication tasks across all sources and targets without needing a deep understanding of the environment or coding requirements. Execute initial data loads, then transition automatically to continuous CDC for instant updates. Attunity Replicate automatically generates target databases based on metadata definitions in the source schema. Any DDL changes made to that schema, such as table/column additions or changes to data types, can be replicated dynamically to the target.

Intelligent management and control

Attunity Enterprise Manager (AEM) allows you to design, execute and monitor thousands of Attunity Replicate tasks across distributed data center and cloud environments, enabling you to scale easily and monitor data flow in real-time through KPIs and alerts. Customize your views by grouping, searching, filtering and drilling down on key tasks to control complex replication activities across the enterprise. You also can leverage comprehensive analytics for capacity planning and performance monitoring, and integrate with microservices architectures through REST and .NET APIS.

Industry's broadest platform coverage

Attunity Replicate supports all major source and target systems for data replication, including relational databases, data warehouse systems, Hadoop, streaming systems, cloud targets and mainframe systems. Attunity integrates data across these platforms seamlessly thanks to years of innovation, optimization and strategic partnerships.

Zero-footprint architecture

Attunity Replicate provides agentless data replication for mainstream database systems. Using platform native utilities, remote change log scanners and secured client connections eliminates the need for intrusive agents, triggers or timestamps on sources and targets, dramatically reducing your administrative burden and improving source production performance. For example, using IBM's Data Server Driver to connect to DB2 z/OS results in significant production performance benefits. Lab tests showed reductions of 85% in source MSU (million service units), 75% in replication latency and 95% in loading time, when compared with agent-based technologies.

Instant updates with flexible, enterprise-class change data capture technology

Our flexible CDC has several options. Attunity Replicate can copy and deliver transactions sequentially in the same order in which they were applied to the source, or create optimized batches of changes to further optimize throughput and latency. Users also can integrate Attunity Replicate CDC with native data warehouse loaders, or feed CDC changes to message streams.

Secure and high-speed cloud data transfer

Attunity Replicate compresses, encrypts and transfers data in parallel streams into, across and out of cloud architectures, including Amazon Web Services, Microsoft Azure and Google cloud. Users also can now feed data from any source to cloud-optimized Snowflake SQL data warehouses.

Easy ingestion of structured data into Hadoop

Attunity Replicate delivers high-performance data loading and publication to Data Lake ecosystems through native APIs, and is certified with all the leading Hadoop distributions including Cloudera, Hortonworks and MapR. You also can ensure transactional consistency in Hadoop targets such as Hive data stores, partitioning updates by time to ensure each transaction is processed holistically and thereby eliminate the risk of partial or competing entries.

Universal stream generation

Attunity Replicate enables databases to easily publish events to all major streaming services, including Kafka, Confluent, Amazon Kinesis, Azure Event Hub and MapR-ES. Users can support multi-topic and multi-partition streaming, as well as flexible JSON and AVRO file formats. They also can separate data and metadata by topic to integrate metadata more easily with various schema registries.

Enabling SAP analytics

Attunity Replicate for SAP easily and securely transfers SAP data, documents, and transactions into Hadoop and other external platforms, and automatically incorporates SAP text and descriptions into target table and column names. Attunity Replicate supports data collection from all core SAP business applications such as ERP, HR, CRM, industry-specific applications, and loading data into SAP HANA.

verizon

"INGESTING DATA FROM SAP AND OTHER SOURCES INTO OUR DATA LAKE IN UNDER AN HOUR IS INVALUABLE. OUR FINANCIAL ANALYTICS AND MODELING TEAMS NOW HAVE INSTANT VISIBILITY FOR REALTIME REPORTING ACROSS THE BUSINESS."

ARVIND RAJAGOPALAN, DIRECTOR,
GLOBAL TECHNOLOGY SERVICES, VERIZON

