



Highlights

- Provides a comprehensive solution for dynamic, near-real-time, incremental delivery of transactional relational data to a broad spectrum of database and big data targets
 - Integrates real-time inventory and sales information with business applications
 - Enables integration of product and customer data across multiple branches, outlets and offices
 - Supports 24x7 worldwide operations with reliable, exceptional performance
-

IBM data replication

To compete in today's fast-paced business climate, enterprises need accurate and frequent sales and customer reports to make real-time operational decisions about pricing, merchandising and inventory management. They also require greater agility to respond to business events as they happen, and more visibility into business activities so information and systems are optimized for peak efficiency and performance. By making use of data capture and business intelligence to integrate and apply data across the enterprise, organizations can capitalize on emerging opportunities and build a competitive advantage.

The IBM® data replication portfolio is designed to address these issues through a highly flexible one-stop shop for high-volume, robust, secure information replication across heterogeneous data stores.

The portfolio leverages real-time data replication to support high availability, database migration, application consolidation, dynamic warehousing, master data management (MDM), service-oriented architecture (SOA), business analytics, and extract-transform-load (ETL) or data quality processes. It also delivers outstanding capabilities for loading real-time information into a data warehouse or operational data store, which can help organizations enhance business agility and visibility into key processes. In addition to direct delivery of changing data to Hadoop, IBM data replication can deliver data to Apache Kafka, providing an ideal big data landing zone and point of enterprise integration for changing transactional source data. Using a Kafka landing zone simplifies the delivery of operational data to multiple consumers such as Hadoop/HDFS, cloud data stores and event-driven consumers.



Leverage a comprehensive solution for data replication

The data replication portfolio includes multiple technologies to support the broadest set of source/target pairings as well as a comprehensive set of data replication topologies. These include:

- Change data capture (CDC)
- Q Replication
- SQL Replication

CDC (Figure 1) addresses a broad spectrum of data synchronization challenges by delivering changed data from a relational database source directly to:

- A target relational database management system (RDBMS) in support of a data warehouse, operational data store or cross-application integration
- A target Hadoop or Kafka cluster in support of enterprise information hubs and other big data initiatives
- A cloud-based version of one of the above
- An IBM PureData® appliance for high-performance analytics
- IBM Information Server DataStage® and QualityStage® for transformation, rich data integration and data quality purposes
- IBM Master Data Management Server for enterprise data integration initiatives

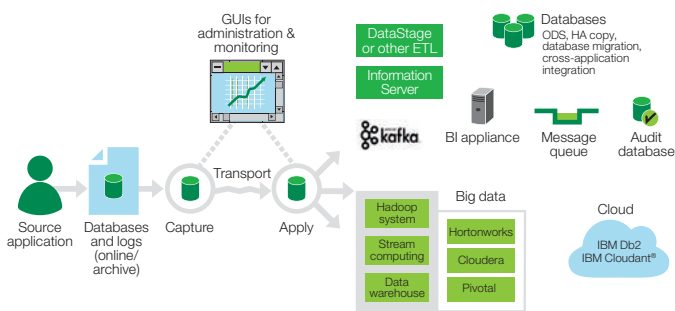


Figure 1. CDC: Log-based change data capture.

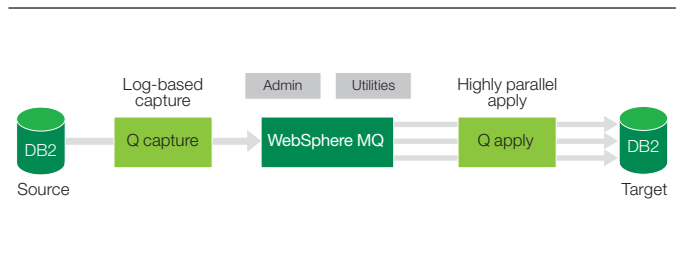


Figure 2. Q Replication architecture.

Q Replication (Figure 2) offers queue-based replication by leveraging IBM WebSphere® MQ for data delivery and fast recovery for high availability after a system outage. Q Replication is ideally suited to the synchronization of IBM Db2® data across the data center or around the globe.

Key specific offerings in the IBM data replication portfolio include:

- IBM InfoSphere® Data Replication (and all of its permutations: for Db2 for z/OS®, IBM IMS™ for z/OS, VSAM for z/OS and non-RDBMS targets)
- IBM Data Replication (and its permutation targets for z/OS sources)
- IBM InfoSphere Classic Change Data Capture for z/OS (IMS, VSAM)

Customers have the flexibility to acquire and deploy components to fit their needs.

Make decisions based on the latest, most accurate information

The IBM data replication portfolio provides real-time feeds of changed data for data warehouse or MDM projects, enabling operational and tactical business decision-making based on the latest information. It dynamically routes data for consumption by one or more applications to help ensure accurate, reliable data is used across the enterprise.

Reduce costs by facilitating application consolidation

As businesses cope with economic uncertainty by cutting costs, IBM data replication helps accelerate application consolidation projects by eliminating multiple application instances on different systems, databases and operating systems. This approach can help reduce maintenance costs dramatically.

Help reduce risk through exceptional performance

Consolidating and moving the data contained within applications can be a sensitive task. Consolidation can be challenging because most system environments comprise a variety of operating systems and databases. Also, if consolidation or migration projects require significant downtime, it could defeat the potential cost savings generated by the initiative.

IBM data replication addresses this issue by allowing data to be migrated from one database or platform to another while end users are still using applications—without downtime or productivity loss for the business. It also supports a broad range of operating systems and databases to help maximize the value of investments for use in future projects.

Instead of using triggers or performing queries against the database, IBM data replication software reads the native database log to capture changes. For businesses faced with shrinking batch windows or overutilized applications, this log-based CDC approach helps ensure the performance of even the most demanding mission-critical applications running on the source system is not adversely affected. In addition, data replication supports near-zero latency for pervasive integration projects.

Data replication in the real world

Data replication is a critical capability for a variety of industries. For example:

- **Financial services companies** must leverage data from legacy financial systems using open platforms that support efficient internet-based communications, statement processing and other applications with minimal impact on key systems.
- **Retailers** must have the ability to sell in stores, over the web and by phone—and they need consistent, up-to-date information that is accessible across all three channels. Data replication can help ensure, for example, that wedding registry purchases are not duplicated and that customers can check product availability and order status over the phone or the web.
- **Shipping companies** must allow customers to access real-time shipment information online, but they must also significantly reduce the impact on their own mission-critical logistics systems to support efficient, continuous operations.
- **Government agencies** need to provide public access to up-to-date information on criminal proceedings on the internet, but they cannot risk unauthorized access to sensitive information on offenders.
- **Health insurers** need to offload queries, reports and backups to avoid negatively affecting performance and user response times on essential claims processing systems.

Enable consistent, trusted data delivery

The downtime costs for the inaccessibility of mission-critical data and applications range from contractual fines and lost productivity to a loss of credibility—all of which may result in lost customers. With IBM data replication software, data can be synchronized between two systems to provide continuous availability. If the primary system is impacted by a planned or unplanned outage, a secondary system is available to serve customers and keep the business up and running.

IBM data replication enables continuous delivery of data to support critical business operations. Transaction consistency is maintained throughout the process to preserve units of work and referential integrity. The software supports full transaction granularity with before-and-after images of all transactional changes. The platform is data event-aware, so it can be used to trigger specific business processes. In addition, fault tolerance capabilities allow organizations to recover to the last committed transaction.

Why IBM?

- The IBM data replication portfolio provides better performance and scalability
- IBM offers broad functionality, support for a wide range of replication endpoints and support for complex implementations for large enterprise customers
- IBM has extensive mainframe knowledge and is uniquely positioned to provide data replication solutions from the mainframe environment
- IBM offers deep industry expertise across most major industries in the worldwide economy

For more information

For more information about IBM data replication, contact your IBM representative, or visit:

ibm.com/analytics/us/en/technology/data-replication

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition.

For more information, visit: ibm.com/financing



© Copyright IBM Corporation 2017

IBM Analytics
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
July 2017

IBM, the IBM logo, ibm.com, DataStage, Db2, IMS, InfoSphere, PureData, QualityStage, WebSphere, and z/OS are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

This document is current as of the initial date of publication and may be changed by IBM at any time.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.



Please Recycle
