

IBM InfoSphere replication server, version 9.5: replicating enterprise data

Highlights

- **Support business continuity and workload distribution across IBM® DB2® instances**
- **Synchronize mainframe data with data in distributed DB2, Oracle, IBM Informix®, Sybase, Microsoft® SQL Server™ and Teradata databases**
- **Manage data distribution and consolidation between headquarters and retail or branch sites**
- **Deliver more current data to data warehouses for tactical decision making**
- **Support synchronization of multiple copies of VSAM data**

Distribute, consolidate and synchronize information

The complexity of today's computing environment is staggering. Whether measured by the number of databases, the pace of new application deployment, the drive toward real-time business intelligence or demanding service-level requirements, IT organizations have their hands full.

Each new application brings not only its own data, but some that might overlap with data in other application domains and, as a result, must be synchronized. Business intelligence must now support tactical decisions based on real-time data. And businesses are driven to increase redundancy not only for availability, but also for regional performance and low-cost capacity.

IBM data replication offerings are designed to address these requirements:

- *IBM InfoSphere® Replication Server, Version 9.5*
- *IBM InfoSphere Replication Server for z/OS®, Version 9.1*
- *IBM InfoSphere Classic Replication Server for z/OS, Version 9.5*

InfoSphere Replication Server: Two replication approaches deliver flexibility and power

InfoSphere Replication Server for z/OS and InfoSphere Replication Server for Linux®, UNIX® and Microsoft Windows® provide two approaches for asynchronous, log-based replication: queue-based replication architecture and an SQL replication architecture. These approaches can be used independently or together for maximum flexibility and function. Both architectures support data sharing configurations for DB2 for z/OS and partitioned database environments for DB2 on Linux, UNIX and Windows—automatically merging the logs.

Regardless of the replication approach, administrators can use a wizard-driven GUI, command-line processor, script-driven processes or all three options to configure the replication environment. Integrated monitoring and statistics make it easier to react to problems and maintain system health.

SQL replication architecture

InfoSphere Replication Server for z/OS and InfoSphere Replication Server for Linux, UNIX and Microsoft Windows support an SQL replication architecture. This architecture is designed to maximize flexibility in managing scheduling, transformation and distribution topologies for populating warehouses or datamarts, maintaining data consistency between applications, or efficiently managing distribution and consolidation scenarios among headquarters and branch or retail sites. They also support replication among mixed relational databases.

InfoSphere Replication Server products include these SQL replication capabilities:

- *Data can be distributed from one database to many and consolidated from many databases to one.*
- *Data can be filtered either horizontally or vertically.*
- *Transformation of data can be performed in-line using standard SQL or stored procedures.*
- *Data movement can be automated on a specific schedule, at designated intervals, continuously or in an event-driven manner.*
- *Data movement can be managed table-at-a-time for scenarios such as warehouse loading during batch windows, or it can be managed with transaction consistency for data that is never offline.*

Queue-based replication architecture

InfoSphere Replication Server products also provide a queue-based architecture designed to achieve high-volume, low-latency data replication with managed conflict detection and resolution. This approach enables backup systems to be used productively, and it enables geographically dispersed systems to share application workloads.

The architecture supports DB2 instances on Linux, UNIX, Windows and IBM z/OS platforms as sources and targets. It also supports Oracle, Sybase, Microsoft SQL Server and Informix instances as targets. It supports VSAM files, IBM IMS® databases and CA-IDMS and Software AG Adabas databases as sources when used in conjunction with InfoSphere Classic Replication Server.

With InfoSphere Replication Server for z/OS and InfoSphere Replication Server for Linux, UNIX and Windows, special focus is placed on the new Data studio web console's health monitoring of the queue based replication environment. InfoSphere Replication Server products include these queue-based replication capabilities:

- *Committed changes are published to IBM InfoSphere MQ message queues, which provide a high-performance, robust and reliable data delivery mechanism.*
- *A sophisticated apply engine determines transaction dependencies and replays transactions on target systems to maximize parallelism and minimize latency.*
- *Conflict detection and resolution features allow backup systems to be used for productive work so that application workload can be distributed across multiple servers.*
- *Data can be filtered so that only the data of interest is replicated.*
- *Stored procedures can be invoked by the apply process so that data can be transformed as it is replicated.*




Figure 1. Queue-based replication architecture features low latency, high throughput and managed conflict detection and resolution

- *With Version 9, new target types enable changed data histories, which are essential for data auditing and analysis purposes.*

InfoSphere Classic Replication Server: Expanding replication to nonrelational mainframe data

InfoSphere Classic Replication Server for z/OS provides asynchronous, log-based capture for mainframe data designed to feed the queue-based replication architecture of IBM System z™, Linux, UNIX and Windows versions of InfoSphere Replication Server. It also provides a stand-alone native VSAM to VSAM replication capability.

First, InfoSphere Classic Replication Server supports both data access and changed-data capture of VSAM,

IMS, CA-IDMS and Software AG Adabas. It accesses mainframe data to support the InfoSphere Replication Server's queue replication and initial load of a target database. It uses various data-specific logging and journaling mechanisms in conjunction with WebSphere MQ for z/OS to feed ongoing changed-data to queue the replication parallel apply mechanism for continuous updating of target data stores.

To support this unidirectional movement of mainframe data to InfoSphere replication server, InfoSphere Classic Replication Server shares mainframe data virtualization technology with IBM InfoSphere Classic Federation Server. This virtualization technology helps ensure



that all replicated data is reformatted into a consistent relational format, regardless of the source data format. The data reformatting is driven by metadata definitions that map the nonrelational data sources to logical relational table constructs. Based on this metadata, the nonrelational source data is made to look like DB2 data before it is delivered to InfoSphere MQ queues and to the InfoSphere Replication Server apply processing mechanism.

A second role for InfoSphere classic replication server is the native, unidirectional replication of VSAM data. This provides a low latency mechanism for synchronizing the content of VSAM files across systems or on the same system. This is especially useful to ensure high availability of VSAM data.

IBM Information Server delivers information you can trust

InfoSphere Replication Server and InfoSphere Classic Replication Server offerings are companion products to IBM Information Server, an innovative new software platform that helps you

derive more value from the complex, heterogeneous information spread across your systems. It enables your organization to integrate disparate data and deliver trusted information wherever and whenever needed, in line and in context, to specific people, applications and processes.

IBM Information Server helps business and IT personnel to collaborate to understand the meaning, structure and content of any type of information across any sources. It also provides breakthrough productivity and performance for cleansing, transforming and moving this information consistently and securely throughout the enterprise, so it can be accessed and used in new ways to drive innovation, help increase operational efficiency and lower risk.

For more information

To learn more about replication and other information integration solutions from IBM, contact your IBM marketing representative or IBM Business Partner, or visit ibm.com/software/data/integration

© Copyright IBM Corporation 2006

IBM Software Group
Route 100
Somers, NY 10589

Printed in the United States of America
December 2006
All Rights Reserved

IBM, the IBM logo, DB2, IMS, Informix, System z, InfoSphere and z/OS are trademarks of International Business Machines Corporation in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, SQL Server and Windows are trademarks of Microsoft Corporation in the United States, other countries or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product or service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. Offerings are subject to change, extension or withdrawal without notice.

All statements regarding IBM future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.

The information contained in this document is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this document, it is provided "as is" without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this document or any other documents. Nothing contained in this document is intended to, nor shall have the effect of, creating any warranties or representations from IBM Software.

TAKE BACK CONTROL WITH **Information Management**