

IBM WebSphere Replication Server, Version 9.1 and IBM WebSphere Data Event Publisher, Version 9.1

Highlights

- **Support business continuity and workload distribution across IBM DB2® instances.**
- **Synchronize mainframe DB2 data with data in distributed DB2, Oracle, IBM Informix®, Sybase, Microsoft® SQL Server and Teradata databases.**
- **Manage data distribution and consolidation between head-quarters and retail or branch sites.**
- **Deliver more current data to data warehouses for tactical decision making.**
- **Initiate business processes based on changed data events.**

IBM Information Integration Solutions

Organizations face an information challenge. Where is it? How do I get it when I need it in the form I need? What does it mean? What insight can I gain from it? Can I trust it? How do I control it? The challenges continue to grow if businesses cannot ensure that they have access to authoritative, consistent, timely and complete information.

IBM Information Integration Solutions help you derive more value from the complex, heterogeneous information spread across your systems. They enable 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. They help business and IT personnel to collaborate to understand the meaning, structure and content of any type of information across any sources. They provide breakthrough productivity and performance for

cleansing, transforming and moving this information consistently and securely throughout the enterprise, so you can access and use information in new ways to drive innovation, increase operational efficiency and lower risk.

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 Information Integration Solutions include four data replication and data event publishing offerings designed to address these requirements:

- *IBM WebSphere Replication Server, Version 9.1*
- *IBM WebSphere Replication Server for z/OS®, Version 9.1*
- *IBM WebSphere Data Event Publisher, Version 9.1*
- *IBM WebSphere Data Event Publisher for z/OS, Version 9.1*

IBM WebSphere Replication Server: Two replication approaches deliver flexibility and power

IBM WebSphere Replication Server for z/OS and *IBM WebSphere Replication Server for Linux®, UNIX®* and *Microsoft Windows®* provide two approaches for asynchronous, log-based replication: a 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, your 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.

Queue-based replication architecture

WebSphere Replication Server products provide a queue-based architecture that is 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 IBM DB2 instances on Linux, UNIX, Microsoft Windows and IBM z/OS platforms as sources and targets. It also supports Oracle, Sybase, Microsoft SQL Server and Informix instances as targets.

With *WebSphere Replication Server for z/OS* and *WebSphere Replication Server for Linux, UNIX* and *Windows*, special focus is placed on new and enhanced tools for configuring and monitoring the queue-based replication environment.

WebSphere Replication Server products include these queue-based replication capabilities:

- *Committed changes are published to IBM WebSphere 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.*
- *With Version 9, new target types enable changed data histories, which are essential for data auditing and analysis purposes.*




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

SQL replication architecture

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

WebSphere 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 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.*

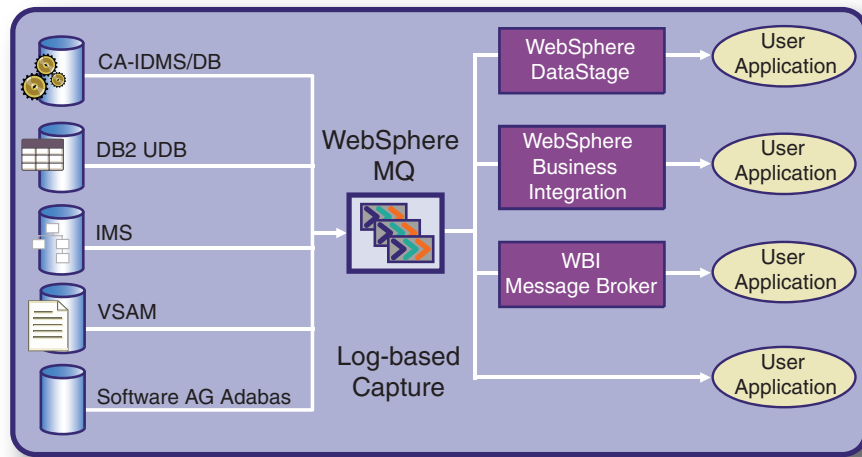


Figure 2: Event publishing captures changed-data events and publishes them as WebSphere MQ messages that can be used by other applications to drive subsequent processing.

IBM WebSphere Data Event Publisher: Data event publishing facilitates business integration

WebSphere information integration event publishing makes it easy to link data events with business processes. Event publishing is available through IBM WebSphere Data Event Publisher for z/OS and IBM WebSphere Data Event Publisher for Linux, UNIX and Microsoft Windows. Data event publishing capabilities are also provided by the IBM WebSphere Classic Data Event Publisher for z/OS family of products, which supports IBM IMS™ (Information Management System), VSAM (Virtual Storage Access Method), Computer Associates Advantage CA-IDMS/DB, and Software AG Adabas.

Changed-data events accelerate data warehousing, enterprise integration

Event publishing provides efficient changed-data publishing environments for scenarios like the following:

- *Application-to-application integration – Changed-data event*

publishing makes it possible to push operational customer data to a packaged customer relationship management (CRM) application.

- *Business process initiation*
 - *Changed-data event publishing allows a customer record to initiate a welcome e-mail, credit verification and an update to the CRM system.*
- *Critical data event monitoring*
 - *Changed-data event monitoring allows events, such as low inventory levels, to trigger a process, such as product restocking workflow.*
- *Data population – Changed-data event publishing can feed a data warehouse, data mart or operational data store by pushing changed data to IBM WebSphere DataStage® or another data integration product that then populates the data store.*

Robust, real-time changed-data capture

WebSphere Data Event Publisher products capture data changes and then publish them to WebSphere MQ as self-describing XML messages. These changed data events can then be used by WebSphere application



and process integration middleware or by a Java™ Message Service-aware application, tool or message broker to drive subsequent processing. This loosely coupled integration helps ensure that each application can be changed independently of every other application.

WebSphere Data Event Publisher products complement and extend your investments in service oriented architecture, enterprise application integration and extract/transform/load (ETL) infrastructure by:

- *Eliminating the hand coding typically required to detect data changes made by operational applications*
- *Removing data event capture overhead from the transaction path by removing it from the operational application*
- *Providing a single integration point –the source data–for data events that might be initiated by multiple applications, making the data event capture independent of the applications and their evolution*
- *Making the data integration independent of the structure or processing flow of the applications involved*
- *Reducing latency for legacy data delivery through ETL tools*

IBM Information Server

IBM WebSphere Replication Server and IBM WebSphere Data Event Publisher 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, increase operational efficiency and lower risk.

For more information

To learn more about replication, event publishing, and other information integration offerings 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
U.S.A.

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

IBM, the IBM logo, the On Demand Business logo, DataStage, DB2, IMS, Informix, WebSphere 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 and Windows are trademarks of Microsoft Corporation in the United States, other countries or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

TAKE BACK CONTROL WITH **Information Management**