

IBM WebSphere Information Integrator Version 8.2 — Replication and Event Publishing

Highlights

- ***Provide a high-availability option for IBM DB2 Universal Database (UDB) customers who need high data throughput and low latency***
- ***Initiate business processes based on DB2 UDB changed-data events***
- ***Manage data distribution and consolidation between headquarters and retail or branch sites***
- ***Support business continuity and workload distribution across DB2 UDB instances and locations***
- ***Deliver more current data to data warehouses for tactical decision making***

IBM WebSphere Information Integration

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 list goes on, and the challenges grow unceasingly if businesses cannot ensure that they have access to authoritative, consistent, timely and complete information.

The IBM WebSphere® Information Integration platform integrates and transforms any data and content to deliver information you can trust for your critical business initiatives. It provides breakthrough productivity, flexibility and performance, so you and your customers and partners have the right information for running and growing your businesses. It helps you understand, cleanse and enhance information, while governing its quality to ultimately provide authoritative information. Integrated across the extended enterprise and delivered when you need it, this consistent, timely and complete information can enrich business processes, enable key contextual insights and inspire confident business decision-making.

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 the demanding service-level requirements, IT organizations have their hands full. Each new application brings its own data, including some that may overlap with data in other application domains and must therefore be synchronized. Business intelligence applications 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 WebSphere Information Integrator Replication Edition Version 8.2 and IBM WebSphere Information Integrator Event Publisher Version 8.2 address these requirements.

Two Replication Architectures Meet Your Needs

IBM WebSphere Information Integration software supports two styles of asynchronous log-based replication: a queue-based replication architecture which complements a structured query language (SQL) replication architecture. Featuring high volume, low latency, and managed conflict detection and resolution, the queue-based replication architecture helps enable high availability scenarios that allow backup systems to be productive and geographically dispersed systems to share application workloads. Replication among IBM DB2® Universal Database instances on Linux®, UNIX®, Microsoft® Windows® and IBM z/OS® operating systems provides business continuity for applications on these systems. On the distributed platform, replication support for high availability massively parallel processing (MPP) configurations using the DB2 Database Partitioning Feature is supported; replication automatically merges the logs from partitioned databases to help ensure data consistency. In addition, on distributed platforms, the queue-based replication architecture supports Oracle, Sybase, Microsoft SQL Server 2000, and Informix Dynamic Server as targets for data apply.

WebSphere Information Integrator Replication Edition offers extensive queue-based replication capabilities:

- *Changes to the database are captured from the DB2 UDB log, minimizing impact to the production systems*
- *Committed changes are published to IBM WebSphere MQ, the industry-leading message queuing system*
- *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 do productive work so that application workload can be distributed across multiple servers*
- *Data can be filtered so that only data of interest is replicated*

Alternatively, the SQL replication architecture maximizes 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 configurations. The SQL replication architecture of WebSphere Information Integrator Replication Edition supports replication among mixed relational databases.

Key features include:

- *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 via standard SQL or stored procedures*
- *Data movement can be automated on a specific schedule, at designated intervals, continuously or event-driven*
- *Data movement can be managed table-at-a-time, such as for warehouse loading during batch windows or with transaction consistency for data that is never offline*

Regardless of the replication approach, your administrators can use a wizard-driven graphical user interface (GUI), command-line processor and/or script-driven processes to configure the replication environment. Integrated monitoring and statistics make it easier to react to problems and maintain system health.




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

Data Event Publishing Facilitates Business Integration

WebSphere Information Integration Event Publisher makes it easy to link data events with business processes. Changed-data events can be captured in distributed DB2 UDB systems as well as z/OS data sources including DB2 UDB, IBM IMS™ (Information Management System), VSAM (Virtual Storage Access Method), Computer Associates Advantage CA-IDMS/DB, and Software AG Adabas. These changes can then be published to WebSphere MQ, giving you additional business integration options for your enterprise.

Changed-data Event Publishers Accelerate Data Warehousing, Enterprise Integration

Changed-data event publishing provides efficient environments for the following scenarios:

- *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 an extract, transform and load (ETL) product that then populates the data store*

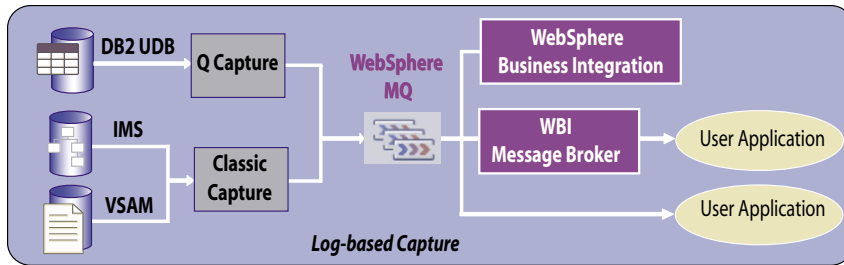

Robust, Real Time Changed-data Capture

WebSphere Information Integration event publishing captures data changes and then publishes them to WebSphere MQ as self-describing extensible markup language (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 Information Integrator Event Publisher complements and extends 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*
- *Removing data event capture overhead from the transaction path*
- *Providing a single integration point—the source data—for data events that may 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*



© Copyright IBM Corporation 2005

IBM Software Group
Route 100
Somers, NY 10589
U.S.A.

Printed in the United States of America
October 2005
All Rights Reserved

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

System Requirements:

IBM WebSphere Information Integrator Version 8.2—Replication and Event Publishing

WebSphere Information Integrator Replication Edition and WebSphere Information Integrator Event Publisher Edition support the following operating systems: IBM AIX®, HP-UX, Linux®, Solaris and Microsoft® Windows®.

Products for the z/OS platform include IBM WebSphere Information Integrator Replication for z/OS, IBM WebSphere Information Integrator Event Publisher for DB2 Universal Database for z/OS, IBM WebSphere Information Integrator Classic Event Publisher for IMS, IBM WebSphere Information Integrator Classic Event Publisher for CA-IDMS and IBM WebSphere Information Integrator Classic Event Publisher for VSAM.

Products for the z/OS platform include IBM WebSphere Information Integrator Replication for z/OS, IBM WebSphere Information Integrator Event Publisher for DB2 Universal Database for z/OS, IBM WebSphere Information Integrator Classic Event Publisher for IMS, IBM WebSphere Information Integrator Classic Event Publisher for CA-IDMS and IBM WebSphere Information Integrator Classic Event Publisher for VSAM.

For current, detailed hardware and software system requirements for these and other WebSphere Information Integrator products, visit ibm.com/software/data/integration

For More Information

To learn more about the technologies and products behind IBM information integration solutions, contact your IBM marketing representative or IBM Business Partner, or visit ibm.com/software/data/integration.

IBM, the IBM logo, the On Demand Business logo, AIX, DB2, DB2 Universal Database, IMS, 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 registered trademarks of Microsoft Corporation in the United States, other countries or both.

Java and all Java-based trademarks are trademarks 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.

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.

♻️ Printed in the United States on recycled paper containing 10% recovered post-consumer fiber.

The IBM home page on the Internet can be found at ibm.com