

IBM DB2 Information Integrator, Version 8.1

Highlights

- **Accesses distributed structured and unstructured data for greater return on information assets**
- **Supports both replication and federated data access to match business needs**
- **Minimizes integration risk and cost by leveraging existing data structures, platforms, skills and tools**
- **Integrates with wider variety of data sources than predecessor products such as IBM DB2 Relational Connect, IBM DB2 Life Sciences Data Connect and IBM DB2 DataJoiner®**

Access real-time information across the enterprise and beyond

As businesses become increasingly interconnected, competitive success demands the ability to integrate and correlate information in real time—not only within the enterprise, but also across the value chain. However, the information that decision makers require often resides in diverse source systems (such as databases, spreadsheets and flat files) and may be distributed across several operating environments. Developers have long struggled to integrate information through a variety of piecemeal—and often costly—approaches.

IBM DB2® Information Integrator, V8.1 is an important step in increasing the efficiency and reducing the cost of information integration. This new middleware technology allows analytical tools, portal implementations, customer data integration solutions and other applications to access disparate, distributed data as if it resided in a single database.

A fundamental underpinning of e-business on demand, DB2 Information Integrator helps businesses increase efficiencies, improve responsiveness to customers and suppliers, make better, more timely decisions and act more promptly on new opportunities.

Combine structured and unstructured data for competitive edge

DB2 Information Integrator provides out-of-the-box access to a broad range of structured and unstructured data sources. These include relational databases, flat files, XML documents, spreadsheets, content repositories, Web sites, and Web services and message queues, with more to come.



With DB2 Information Integrator, Version 8.1, you are finally able to bring all of your information assets together, for real-time decision making.

Make the most of existing data, platforms and skills

The unique ability of this offering to integrate both structured and unstructured data opens the door to new types of applications and allows greater returns on existing information assets. For instance, when relational data, such as client account information, is combined with unstructured data, such as claims images, analyst reports or Web research, decision makers can acquire a more complete picture of their business options. And companies can provide more valuable service to customers and partners through consolidated content and information delivery.

Match access strategy to business value

Unlike other information integration offerings, DB2 Information Integrator gives developers a choice of how to access far-flung data. When it makes sense to keep the data physically close to the application (for better data access performance and availability), DB2 Information Integrator performs hub-based

replication and caching. This centralized approach gives predictable query performance and repeatable access to a stable view of data with controlled refresh intervals. Utilizing a hub-based approach to data consolidation means having the flexibility to support configurations with multiple data sources going to one or many data copies and vice versa.

Alternatively, when data diversity, volume or ownership make replication too expensive or impractical, DB2 Information Integrator lets developers leave distributed data where it is, yet access it transparently as though it were a single source. This federated access-in-place approach is appropriate for the predominantly read-only access scenarios common to enterprise-wide reporting, business intelligence, portal infrastructures and customer data integration.

With either replication or the access-in-place approach, companies can integrate information through the use of existing data structures, data management systems and servers. This significantly lowers the risk and cost of information integration projects.

Flexibility of multiple integration approaches

DB2 Information Integrator may be implemented in any of the following ways:

- *As a federated data server, enabling administrators to configure relational and non-relational data source access and define integrated views across diverse and distributed data. Applications can query the integrated views as if they were a single database. A developer toolkit is provided to add access to other sources. Federated data server capabilities include:*
 - *Graphical tools to configure access to source data*
 - *Integrated views across multiple sources*
 - *Standard structured query language (SQL) application programming interface (API) plus rich text semantics for query*
 - *Database or Web service clients*
 - *Output as standard SQL answer sets or XML documents*
 - *Cost-based distributed query optimization*
 - *Data caching of relational data via Materialized Query Tables (MQTs)*

- *Structured/semi-structured* — IBM DB2 Universal Database™, IBM Informix® Dynamic Server, IBM Informix Extended Parallel Server, Microsoft SQL Server, Oracle, Sybase SQL Server, Sybase Adaptive Server Enterprises, Teradata, Open Database Connectivity (ODBC), OLE DB, IBM WebSphere® MQ message queues, Web services, Microsoft Access, Microsoft Excel spreadsheets, flat files, XML documents
- *Rich content* — IBM Lotus® Notes®, IBM Lotus Document Manager (Domino.Doc®), IBM Lotus Discovery Server™, IBM Lotus Team Workplace (QuickPlace®) 3.0, IBM Lotus Instant Messaging (Sametime®), IBM Websphere Portal Search Engine, Microsoft Index Server, Microsoft Site Server, Microsoft Exchange, IBM SecureWay LDAP Directory Server directories, file systems, IBM DB2 Information Integrator for Content sources (IBM DB2 Content Manager, IBM DB2 Content Manager OnDemand, IBM ImagePlus®, and more) and over 18 Web search sites including Yahoo!, Lycos, Excite, HotBot, Google, Alta Vista, ABCNews.com, AOL.com Search, Business Wire, CNN, Canada.com, EuroSeek, EAST Search, GoTo.com, NBCi, PR Newswire WebCrawler, Yahoo News Headline Search.
- *Life sciences sources* — data sources accessible by Entrez, Blast, HMMer, BioRS, Documentum

- As a Replication Server for mixed relational databases
 - Data can be replicated between: DB2, Informix, Microsoft, Oracle, Sybase
 - Teradata and Informix Extended Parallel Server supported as a replication target
 - Distribution of data from one database to many and consolidation of data from many databases to one
 - Transformation performed in-line via standard SQL or stored procedures
 - Automated on a specific schedule, designated intervals, or continuously
 - Table-at-a-time or a set of tables
- As a local database server, for customers who choose a centralized access strategy for data availability or performance. Having a full-function, local database store available simplifies the creation of new integration-oriented applications, providing local storage for many purposes.

Leverage industry standards and existing skills and tools

Based on industry standards such as SQL, XML, Java™ and other Web services protocols, DB2 Information Integrator enables broad system and application interoperability. It also helps minimize training requirements for information integration projects and lowers client-specific coding and maintenance costs. At IBM, for

example, in-house studies for joining disparate data sources using DB2 Information Integrator have revealed code savings of 40 percent to 65 percent.

With DB2 Information Integrator, developers can leverage existing SQL skills to implement integration projects. The power of SQL can be applied to relational and non-relational data, to join and merge data, compute statistical functions, aggregate data, use online analytical processing features, and compose or transform XML documents. DB2 Information Integrator also fits transparently behind widely used integrated development environments, portal infrastructures and analytical tools.

Rely on proven IBM technology

DB2 Information Integrator, V8.1 is based on proven IBM technologies, such as SQL query optimization, that are the result of 25 years of information management research and development. The value customers reap from these technologies is growing every year, as IBM continues to invest in them.

As part of a scalable, cross-platform business integration infrastructure, DB2 Information Integrator works with the IBM WebSphere business integration portfolio, including IBM WebSphere Business Integration, IBM WebSphere Portal, IBM WebSphere MQ and IBM WebSphere Studio.

IBM DB2 Information Integrator products

DB2 Information Integrator, V8.1 is offered in five editions:

DB2 Information Integrator Replication

Edition is designed for the data replication needs of medium to large businesses. This edition provides an ideal integration solution for situations that require replication of relational data for query performance, accessibility and locality. The integrated administration and management facilities in DB2 Information Integrator Replication Edition provide end-to-end copy management between popular relational data sources such as DB2 Universal Database, Informix Dynamic Server, Microsoft SQL Server, Oracle and others. DB2 Information Integrator Replication Edition can be deployed on UNIX®, Windows® and Intel® Linux environments. This edition is licensed by the number of processors and data source connections.

DB2 Information Integrator Standard Edition offers both data replication functions and federated data access for a rich set of information integration functions in a single package. Federated data access gives applications integrated access to diverse and distributed data sources including relational databases, XML documents, flat files, text documents, content sources, the Web and



DB2 Information Integrator provides a comprehensive solution that brings together information from multiple sources, as if it existed in a single database.

a variety of other sources. Using DB2 Information Integrator Standard Edition, businesses can extend reporting and analysis applications with real-time data, semi-structured data, non-relational data, or rich content, yielding more current and comprehensive results. Businesses can also reduce costs and deploy integration projects more quickly using their standard development environments. DB2 Information Integrator Standard Edition can be deployed on UNIX, Windows and Intel Linux environments. This edition is licensed by the number of processors and data source connections.

DB2 Information Integrator Advanced Edition

provides all of the federated data access and data replication functions of DB2 Information Integrator Standard Edition and includes a full use license of DB2 Universal Database Enterprise Server Edition (ESE). DB2 Universal Database ESE adds industry-leading database management capabilities, enabling data warehouses and other integration and general application solutions that require local storage. Together, the database management capabilities and information integration functions form a robust platform for information management needs. DB2 Information Integrator Advanced Edition

can be deployed on UNIX, Windows and Intel Linux environments. This edition is licensed by the number of processors and data source connections.

DB2 Information Integrator Advanced Edition Unlimited provides all of the federated data access, data replication functions, and full-use DB2 Universal Database ESE capabilities of DB2 Information Integrator Advanced Edition. Plus, there is no charge for connections, eliminating the need to track and acquire DB2 Information Integrator Connector licenses. DB2 Information Integrator Advanced Edition Unlimited can be deployed on UNIX, Windows and Intel Linux environments. Licensing of this edition is based on the number of processors.

DB2 Information Integrator Developer

Edition offers a low-cost package for a single application developer to design, build, prototype and test applications intended to run on DB2 Information Integrator. This comprehensive offering for developers contains all the functionality of DB2 Information Integrator Advanced Edition, a developer toolkit for extending distributed access to custom data sources, WebSphere and DB2 development tools. The software in this package cannot be used for production systems. DB2 Information Integrator Developer Edition is licensed by the number of users.

Operating system requirements:

AIX

- AIX® (32-bit) 4.3.3 with maintenance level 9 or later
- AIX (32-bit & 64-bit) 5.1.0 with maintenance level 2 or later
- AIX (32-bit & 64-bit) 5.2.0

HP-UX 11i or later (32-bit and 64-bit)

Linux on Intel, AMD systems

Solaris 7, 8 and 9 (32-bit and 64-bit)

Windows

- Windows NT® 4.0 with SP™ 6a or later
- Windows 2000 Professional
- Windows XP 32-bit (supported only for development and test use)

This chart summarizes the key new functions available with DB2 Information Integrator as compared to DB2 DataJoiner, DB2 Relational Connect, DB2 Life Sciences Data Connect and DB2, V8.1:

Function	DataJoiner	DB2 UDB 7.2 & RL/LSCD	DB2 UDB 8.1	DB2 II 8.1
Federated reach:				
DB2/Informix Read/Write	●	● (R/O)	● (1PC ¹)	● (1PC)
Oracle, Sybase, SQL Server, Teradata R/W	●	● (R/O)		● (1PC)
ODBC R/W				● (1PC)
Excel, flat file, Life Sciences wrappers (R/O)		●		●
XML wrapper (R/O)		●		●
Extended search wrapper (R/O)				●
MQ Series (1PC)		●	●	●
Wrapper architecture		●		●
Wrapper toolkit for developers				●
Control Center—Discovery for nicknames				●
Caching—MQTs over nicknames			●	●
Web services provider		●	●	●
Web services consumer			●	●
Application development via standard tooling (WS Studio, MS Visual Studio)		●	●	●
Heterogeneous replication	●			●

Upgrading to DB2 Information Integrator, V8.1

IBM DB2 Information Integrator, V8.1 is the successor product for IBM DB2 Relational Connect, IBM DB2 Life Sciences Data Connect and IBM DB2 DataJoiner.

There are significant functional benefits for customers to move to DB2 Information Integrator, most notably:

- *Performance improvements (e.g, a new push down analysis phase is added to the optimizer and clients can cache federated relational data in MQTs)*
- *Usability improvements (e.g, a new discovery and auto-mapping feature helps clients configure accessed sources more quickly)*
- *Additional data sources (e.g, the extension to rich content access)*

Being built on DB2 Universal Database, Version 8.1 technology, DB2 Information Integrator contains new SQL functions of the DB2 Universal Database engine. In addition, there are new functions on heterogeneous sources unique to DB2 Information Integrator.

For more information

To find out more about the technologies and products behind IBM solutions, contact your IBM marketing representative or IBM Business Partner, call 1 877 234-5678 or visit:

ibm.com/software/data/integration



© Copyright IBM Corporation 2003

IBM Corporation
Silicon Valley Laboratory
555 Bailey Avenue
San Jose, CA 95141
U.S.A.

Printed in the United States of America
09-03
All Rights Reserved

AIX, DataJoiner, DB2, DB2 Universal Database, Domino.Doc, the e-business logo, IBM, the IBM logo, ImagePlus, Informix, Lotus, Lotus Discovery Server, Lotus Notes, Notes, QuickPlace, Sametime and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries or both.

Intel is a trademark of Intel Corporation in the United States, other countries or both.

Windows and Windows NT 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.

¹ One-phase commit.

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.

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

