



## IBM InfoSphere Warehouse 9.5

---

### Highlights

---

- **Helps simplify enterprise data warehouse development, deployment and maintenance with a comprehensive, integrated solution**
- **Helps streamline business intelligence processes and enhance flexibility with powerful analytic components**
- **Boosts data warehouse performance and improves productivity with advanced IBM DB2 9.5 optimization features**

### Getting to a single version of the truth

Access to timely, accurate information is critical as enterprises strive to boost competitiveness and innovation.

Limited by inflexible data warehouse and business intelligence solutions, IT teams are often challenged to meet business users' information needs.

As a result, isolated datamarts and data warehouses—often containing conflicting data—have proliferated.

Worse yet, most business intelligence solutions built on data warehouses use components from various vendors, creating integration and enterprise standardization challenges. Downsides can include frustrated end users, high development and maintenance costs, and poorly performing solutions that do not scale. IT teams must be able to support business requirements for a single version of the truth, while containing costs, accommodating ever-increasing data volumes and dealing with the constantly changing needs of different business constituents.

IBM InfoSphere™ Warehouse 9.5 software is a comprehensive solution for integrated dynamic warehousing that provides extended capabilities

from design to analytics and visualization. Based on open standards and supported by the powerful IBM DB2® 9.5 platform, InfoSphere Warehouse provides the tools and infrastructure to help data warehouse architects and administrators efficiently design, deploy and maintain an enterprise data warehouse.

### Implementing a comprehensive enterprise data warehouse solution

InfoSphere Warehouse is an ideal platform for dynamic warehousing solutions, providing a powerful range of capabilities that go beyond traditional warehouses, including:

#### **A powerful DB2 database foundation.**

With its massively scalable, shared-nothing architecture, DB2 9.5 provides high performance for mixed-workload query processing against both relational and native XML data. Advanced features such as data partitioning, row compression, multidimensional clustering (MDC), materialized query tables (MQTs) and online analytical processing (OLAP) capabilities make DB2 9.5 a powerful engine for dynamic warehousing.

**Modeling and design tools.** InfoSphere Warehouse software includes tools, based on IBM Rational® Data Architect software, that enable you to design, model and reverse-engineer physical database schemas. These tools provide a range of capabilities, including physical data modeling, model validation, compliance and impact analysis, and change management.

**Embedded data movement and transformation.** An included SQL Warehousing Tool enables you to prepare and populate the data warehouse structures required for data mining, multidimensional analytics and embedded analytics. Data flows, control flows and transformations can be built using a graphical interface and deployed within the warehouse.

**Data partitioning.** InfoSphere Warehouse provides advanced capabilities for data partitioning, giving you multiple ways to distribute data across servers for large-scale parallelism and linear scalability. The shared-nothing architecture of DB2 9.5 helps ensure that performance won't degrade as your warehouse grows. And, because InfoSphere Warehouse can physically cluster data on multiple dimensions, order data by value range and limit input/output (I/O) to relevant data partitions, it can reduce the work needed to resolve many queries.

## Software packages at a glance

Features for InfoSphere Warehouse	Base Edition	Enterprise Edition
IBM DB2 Enterprise Server Edition 9.5 software <sup>†</sup>	✓	✓
Warehouse administration and control tools	✓	✓
Modeling and design tools	✓	✓
Embedded data movement and transformation	✓	✓
Cubing services for OLAP	✓	✓
Data partitioning	✓	✓
Extreme workload management		✓
Deep compression		✓
Data mining, modeling and scoring		✓
Unstructured analytics		✓
In-line analytics and visualization		✓
IBM Cognos 8 Business Intelligence Starter Edition software	✓	✓

<sup>†</sup>InfoSphere Warehouse also supports IBM DB2 8.2 software.

**Deep compression.** With InfoSphere Warehouse, you can leverage storage optimization technology in DB2 9.5 that can significantly reduce disk space requirements and improve query performance. Tests conducted using data from the industry standard Transaction Processing Performance Council TPC-H data warehousing benchmark demonstrated between 45 and 69 percent savings in disk space.\*

**Extreme workload management.** InfoSphere Warehouse workload management capabilities can enable realtime delivery of business insights without compromising performance. User activity strain on traditional server solutions can inhibit realtime delivery of information to a broad set of users

and applications. With the advanced workload management capabilities in InfoSphere Warehouse, you can prioritize queries coming from different users and applications and control the number of underlying resources dedicated to those processes, better supporting service level agreements for key constituents.

**No copy analytics.** Embedded analytics capabilities deliver a set of sophisticated, yet easy-to-use tools within the data warehouse. Advanced cubing for OLAP, unstructured analytics, in-line analytics, and data mining and visualization capabilities help you tap unutilized sources of information about your business. These capabilities form a robust analytics solution that can help deliver valuable business intelligence to a wide pool of users.

## Supported operating systems and Web browsers

Operating systems	
InfoSphere Warehouse database and application server	InfoSphere Warehouse client
IBM AIX® V5.3, V6.1 platform HP-UX (Intel® Itanium® processor-based system) Red Hat Enterprise Linux® 4 and 5 SUSE Linux Enterprise Server 9 and 10 Sun Microsystems Solaris 9 and 10 Microsoft Windows® Server 2003 and 2008	Microsoft Windows XP and Microsoft Windows Vista Red Hat Enterprise Linux 4 (32-bit) SUSE Linux Enterprise Server 9 (32-bit)
Web browsers	
Microsoft Internet Explorer 6, 7 Mozilla Firefox 1.5, 2.0 Mozilla 1.7.8 (for IBM WebSphere® Application Server software only)	

**Cubing services for OLAP.** The InfoSphere Warehouse cubing services for OLAP feature enables multidimensional data analysis without extracting the data from your warehouse. InfoSphere Warehouse 9.5 includes native support for Microsoft® Excel PivotTable Service software, enabling ad hoc analyses or the delivery of standard spreadsheet reporting—all while working within the Microsoft Excel application. In addition, cubing services cubes are now first-class data providers to the IBM Cognos® 8 Business Intelligence platform. The entire suite of Cognos clients and applications can leverage these powerful warehouse-based data cubes.

### **Data mining, modeling and scoring.**

Unlike solutions that require users to extract the data from the warehouse, independently analyze it and then send the results back to the warehouse, InfoSphere Warehouse provides embedded data mining, modeling and scoring capabilities.

This means you're working with current data and delivering analytics in realtime. InfoSphere Warehouse supports standard data mining model algorithms such as clustering, associations, classification and prediction; additional algorithms may be imported in industry standard Predictive Model Markup Language (PMML) format.

**Unstructured analytics.** Most business intelligence solutions can't access the majority of information captured across the organization, such as call center notes, customer feedback, and freeform text fields, along with documents and Web pages. InfoSphere Warehouse supports the analysis of previously untapped unstructured data, helping to provide additional insights into customer and product issues.

### **In-line analytics and visualization.**

InfoSphere Warehouse includes features for visual analysis of data mining results, along with display and navigation of multidimensional analysis. These

capabilities are delivered as rich presentation components—rather than via a stand-alone application—which can be easily customized and embedded directly into Web-based applications for use by a broader range of users.

### **Leveraging powerful analytics tools to increase flexibility and process efficiency**

InfoSphere Warehouse provides a set of tools that help simplify data warehouse and analytics development and deployment. These interfaces enable you to design the warehouse and populate data structures, as well as perform analytics and manage data mining and multidimensional cubing through common interfaces.

The tools within InfoSphere Warehouse are built on Eclipse technology and provide the design-time environment for building, testing and debugging the elements of a dynamic warehouse. A familiar graphical development environment allows users to drag and drop, and cut and paste intuitive warehouse operators and entities. Easy integration of common code management repositories, such as IBM Rational ClearCase® software and concurrent versions system (CVS) tools, supports collaborative team development. The Web-based administration console enables database administrators and architects to deploy and manage InfoSphere Warehouse operations in a production run-time environment.



To help users understand and use the new environment quickly, InfoSphere Warehouse includes a sample database environment and tutorials for the development and deployment of solutions.

### **Boosting data warehouse performance with DB2 9.5 software features**

In addition to advanced compression capabilities, DB2 9.5 offers MQTs and MDC to boost warehouse performance. An MQT is a presummarized, preaggregated table that stores query results as data. The DB2 software optimizer transparently redirects queries from base tables to matching MQTs, helping to improve the performance of complex aggregate queries. Regular tables and MQTs can be organized into multidimensional clusters. This helps optimize disk I/O, reduces system maintenance and improves analytic query performance.

### **Introducing IBM Cognos 8 Business Intelligence Starter Edition software**

All InfoSphere Warehouse editions include Cognos 8 Business Intelligence Starter Edition software, which helps end users learn to perform enterprise

reporting against the data warehouse. Built on the powerful Cognos 8 Business Intelligence platform, the starter edition provides standard relational reporting and query capabilities, right out of the box.

### **Why IBM?**

InfoSphere Warehouse helps enable dynamic data warehousing solutions that can integrate with your existing investments and scale to hundreds of terabytes. And high-availability features support 24x7 availability and quick recovery from hardware or software failures. Moreover, IBM offers worldwide support for InfoSphere Warehouse software implementations, so virtually all of the necessary tools and resources are available to help make your enterprise data warehouse a reality.

### **For more information**

To learn more about IBM InfoSphere Warehouse software, contact your IBM sales representative or IBM Business Partner, or visit:

[ibm.com/bi](http://ibm.com/bi)

[ibm.com/software/data/infosphere/warehouse](http://ibm.com/software/data/infosphere/warehouse)

© Copyright IBM Corporation 2008

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

Produced in the United States of America  
December 2008  
All Rights Reserved

IBM, the IBM logo, [ibm.com](http://ibm.com), and InfoSphere are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft and Windows are trademarks of Microsoft 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.

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. The information contained in this documentation is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this documentation, 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 documentation or any other documentation. Nothing contained in this documentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

Each IBM customer is responsible for ensuring its own compliance with legal requirements. IBM customers are responsible for ensuring their own compliance with legal requirements. It is the customer's sole responsibility to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

\* Chang, Lee, Bissoon, "Row compression in DB2 9," *IBM*, <http://www.ibm.com/developerworks/db2/library/long/dm-0610chang/#download>.