Nonstop database
HiRDB Version 9

HITACHI
Inspire the Next

HiRDB Version 9

HiRDB is an evolving database for continuing business services.

Ensuring Non-stop Business - HiRDB

Hitachi is an IT vendor representing Japan. Hitachi covers a wide range of business fields and continues to support the infrastructures of society in Japan and around the world.
"MADE IN JAPAN" hardware has worldwide recognition for its high precision, quality, and performance. For example, the well-known Hitachi disk array system is esteemed for its world-leading technology.
Hitachi's technology also covers more than just hardware. Hitachi has dealt with systems that support the lifelines of society for many years, and therefore has developed solid technology, always keeping in mind that our customers' businesses should never be interrupted. Hitachi's engineering mindset and technology have been inherited in the research and development of servers, storage, and database systems that support system foundations, based on a comprehensive plan from establishing an information infrastructure to providing constant stable non-stop operation and disaster management.
This movement led to the **Hi**ghly-scalable **R**elational **D**ata**b**ase - **HiRDB** Version 9 (hereafter referred to as HiRDB). To protect IT infrastructures, HiRDB incorporates Hitachi's commitment to reliability and high technology.

HiRDB technology - inheriting Hitachi's database resources

Since the initial development of database technology, Hitachi has developed their own original mainframe databases using network data, hierarchical data, and relational data models. Since then, for more than 30 years, Hitachi has widely supported nationwide mission-critical information systems such as bank account online systems.
HiRDB is an open-platform relational database, inheriting mainframe database technologies with a proven track record. HiRDB is the embodiment of the various high-availability and high-reliability technologies that Hitachi has developed over the decades.

History of Hitachi Database Development

Year	Hitachi Activities
1974	Released the **PDM** and **PDMII** database products using the network data model. Released the **ADM** database product using the hierarchical data model.
1984	Released the first RDBMS product called **RDB1**.
1988 to 1991	Released the large-scale RDBMS XDM/RD and structured XDM/SD database products.
	Developed the built-in database processor **IDP** containing the operation function for RDB, and the **RDSP** filtering processor.
1994	Released the **HiRDB** RDBMS using the **XDM/RD** technology that supports open systems.
1995 to 1997	Released **HiRDB V2** through **HiRDB V4** in succession, with improved functions and performance.
1999	Released **HiRDB V5** that supports ORDBMS.
2001	Released **HiRDB V6** enhanced with XML and JAVA™ technologies to meet the demand of Internet Business and Marketing.
2003	Released **HiRDB V7** with disaster recovery, the high-speed failover function available 24 hours a day, 7 days a week, a security function, and storage system support.
2006	Released **HiRDB V8** for support of the information integration aspect of SOA-based service integration.
2008	**HiRDB V8.4** obtained ISO/IEC 15408 EAL4+ALC_FLR.1 certification.
2010	Released **HiRDB V9**, a high-performance and high-reliability database that supports the expansion of business through cloud computing.
2011	Released **HiRDB V9.1**, which supports the grid batch linkage facility for improving parallel processing capabilities.

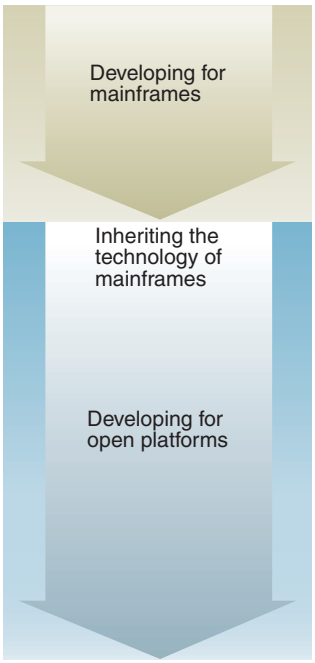

*ORDBMS: Object Relational Database Management System

HiRDB reliability - supporting society's infrastructures

In Japan, HiRDB is well received and running in many important mission-critical information systems supporting the infrastructures of society.

Public administration
- Taxation management databases
- Pension collection and payment control databases
- Immigration control databases
- Registry databases

Banks
- Account system online databases

Securities
- Securities transaction systems

Power
- Electricity bill collection databases

Telephone
- Telephone bill databases

Transportation
- Ticketing and seat reservation databases of railroad companies
- Aircraft maintenance information databases for airlines

Medical
- Electronic medical chart databases

...as well as many other databases

HiRDB continues to stably run and fulfill the crucial responsibilities of these critical systems, 24 hours a day, 7 days a week.

HiRDB evolution - incorporating the latest technology

HiRDB is compliant with the ISO SQL92, ISO SQL99, and ISO SQL2003 international standard interfaces. It also uses parallel processing technologies such as a Shared Nothing architecture and parallel servers to ensure advanced scalability. And with an enhanced data warehousing platform and interface supporting XML and Java, it has evolved even further.
In the latest version of HiRDB, batch processing has been dramatically accelerated. This has been achieved by linking with the grid batch system, and distributing batch processing across servers to enable fast, concurrent execution. Furthermore, by having the database data reside in memory, disk I/O processes have been eliminated during batch processing, which has also contributed to further acceleration. In addition, the native XML search functionality has been enhanced to enable full-text searches for XML data as well as searches for variant notation and synonyms. It also focuses on non-stop database functionality, to support IT infrastructure, and has been enhanced with failure recovery functionality including world-class high-speed cluster-switching functionality that can be performed in dozens of seconds, and disaster recovery functionality with zero data loss thanks to linkage with Hitachi's renowned disk array systems. Last but not least, it also includes autonomic functionality to assist with security and daily operation.

Hitachi will continue innovating and incorporating technology into HiRDB. Our customers can safely entrust their systems to the highly reliable, high-quality database - HiRDB.

Providing Even Greater Extensibility

Enabling high extensibility through a Shared Nothing architecture

HiRDB adopts a Shared Nothing architecture able to maintain high performance even for increasing amounts of data, just by adding new servers.

IT systems face increasing amounts of data handled as a result of business growth. This means that the processing performance demanded for IT systems is growing significantly. Since HiRDB enables processing performance to grow proportionally to the number of servers, it enables high performance to be maintained even when the amount of data increases. This is due to its "Shared Nothing" architecture, in which resources are never shared across multiple servers. Since there is no competition over each server resource, it can be applied to even large-scale systems of more than 10 servers.
Systems can start on a small scale, and then be developed into large-scale systems, enabling operation to continue seamlessly to suit the scope of business.

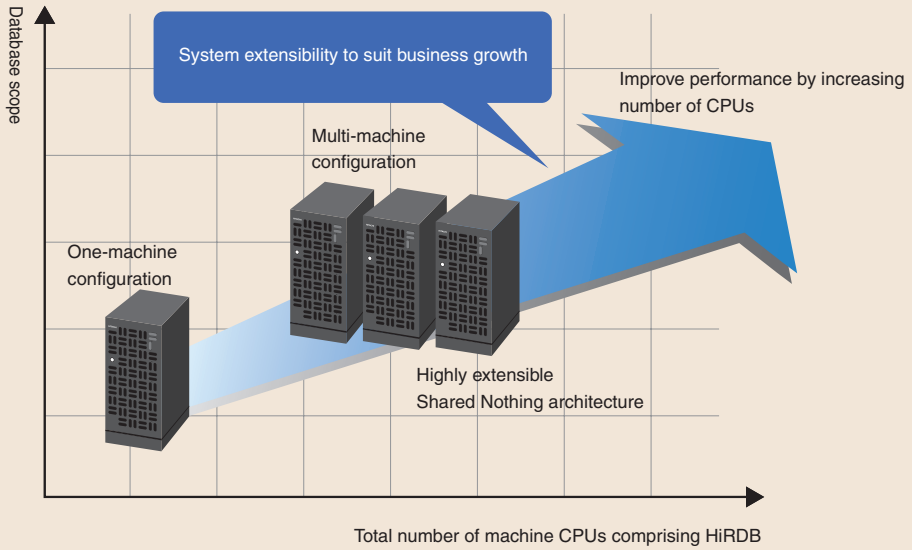

Accelerate performance

Accelerate batch processing extending the execution time available to online transaction processing

The amount of data processed by business systems increases every year. This has resulted in a greater amount of time being taken up by batch processing to handle the data, as well as leading to reduced online operation time. HiRDB carries out batch processing quickly to extend the time available for online operation.

HiRDB can logically split up a table, and store the logical subsets on multiple servers across different locations-such as various business facilities-to process them in a distributed and parallel way. HiRDB can link with the grid batch system to make sets out of the data that has been stored on each of the respective servers and the batches that process that data. This makes it possible to quickly process the batches on the different servers concurrently.

In addition, HiRDB has in-memory data processing capabilities that allows it to execute batch processing in memory. This further accelerates batch processing by eliminating the disk I/O processes that can create performance bottlenecks.

In this way, HiRDB can process even large amounts of data in a shorter time period, which means online business operation time can be extended.

* Note: The optional products HiRDB Accelerator and uCosminexus Grid Processing Server are required. Only a Japanese-language version is available for these products.

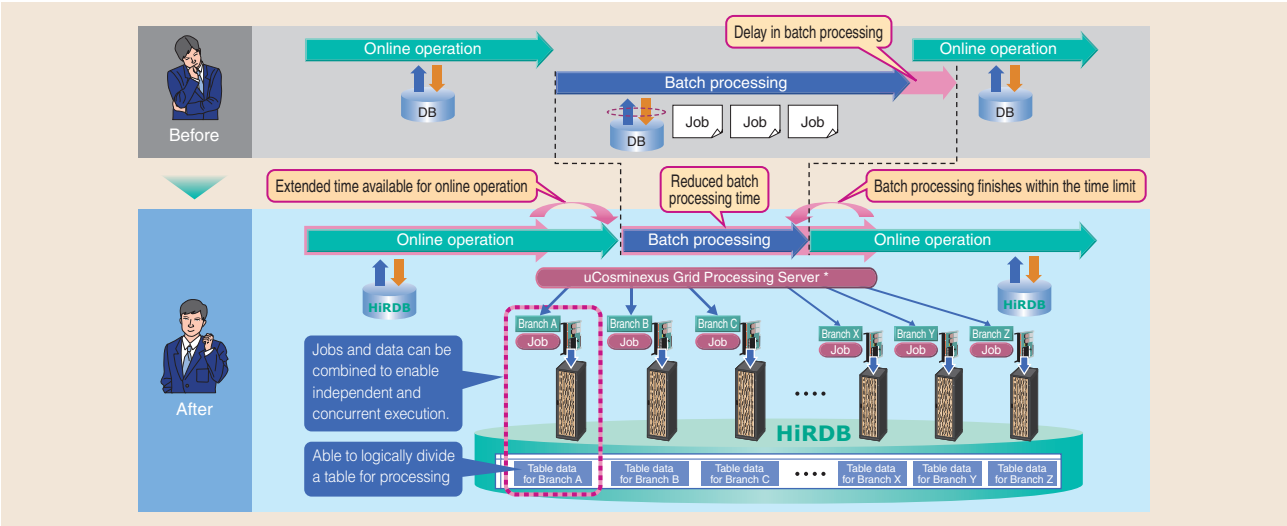

*: uCosminexus Grid Processing Server is a distributed batch job execution system that is designed to facilitate the construction and operation of systems for decentralizing batch processing.

Native XML support

Taking advantage of XML's flexibility to enable rapid addition and modification of operations

XML is highly flexible, and widely used as an inter-operational interface and content exchange format. With XML, even when the data structure used for inter-operational interfaces or data structure for content exchange no longer needs to be changed, system changes can be kept to a minimum, to minimize impact on operations.

HiRDB supports XML natively. The HiRDB XML format can be used to store XML data as-is in tables, without redefinition in the event that the XML data structure changes, taking maximum advantage of the flexibility of XML. HiRDB also supports most of the functionality for SQL/XML, an SQL standard that uses XQuery, a standard XML query language.

* Note: The related product HiRDB XML Extension is required. Only a Japanese-language version is available for this product.

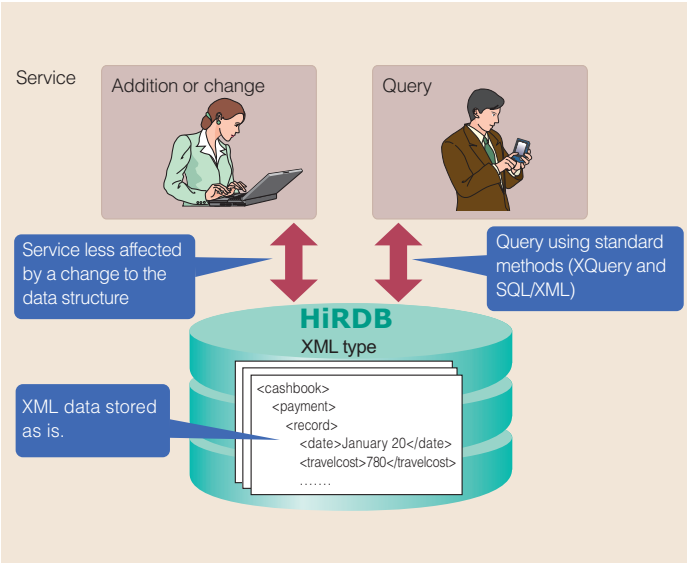

Providing Even Higher Availability

Pursuing higher availability with linkage to storage

Batch processing and reorganization are possible 24 hours a day, everyday, without affecting online performance.

HiRDB can execute batch processing, such as data aggregation and mass-data storage, in parallel, with minimum influence on online service performance, thus supporting 24 hours a day, 7 days a week, non-stop service.
HiRDB can commit the database at a desired time to create a moment in which all transactions are completed. In that moment, a consistent replica database can be created from the master database using the SAN high-speed volume copy for the Hitachi disk array system and other software mirroring functions. This enables online backup, batch processing, and reorganization that is independent of the master database. Since the master database and replica database use different volumes, there is very little reduction of online performance.

*SAN: Storage Area Network

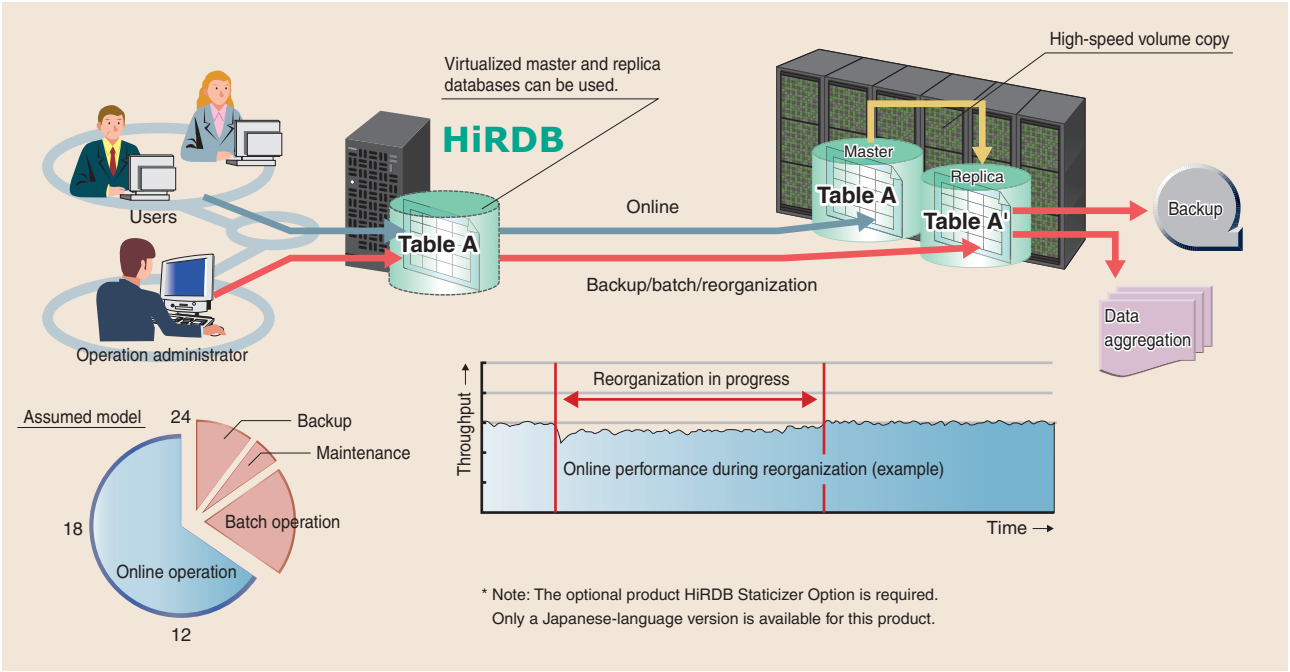

* Note: The optional product HiRDB Staticizer Option is required.
Only a Japanese-language version is available for this product.

Example of using online reorganization

You can temporarily switch the database used for the online service to a replica database, and reorganize the master database without interrupting the online service. During reorganization, there is very little reduction of online performance. Differential data in the update is automatically applied to the replica database.

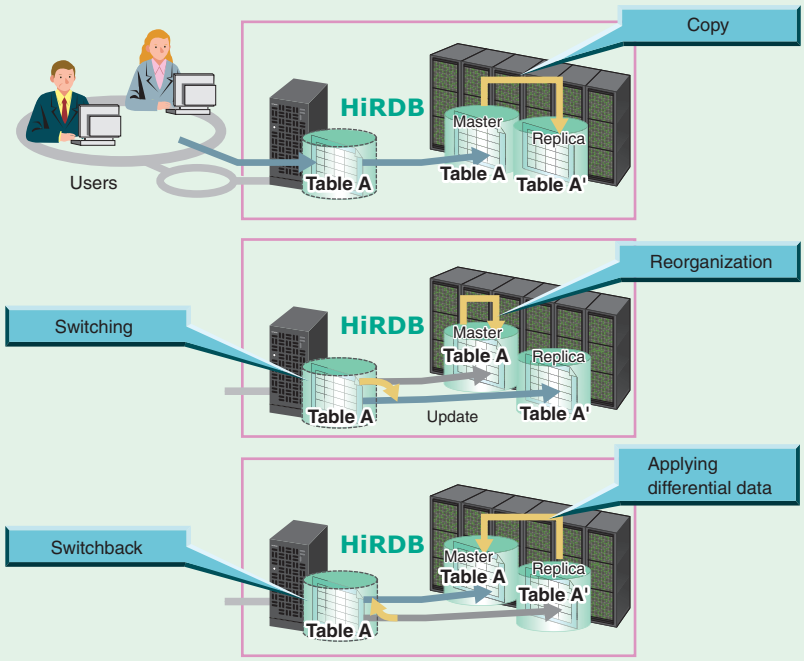

Scene 1: From 21:00

- Copy the master database and create a replica database.

Scene 2: From 21:30

- Switch the online service to the replica database on which the service continues.
- Reorganize the master database.

Scene 3: From 22:00

- Apply the differential data of the update of the replica database to the master database, and then switch the online service back to the master database on which the service continues.

Providing Even Higher Reliability

Sophisticated fault-tolerance technology

Complete fault-tolerance enables a highly reliable non-stop system.

Even if a failure occurs in a database or server, the online service does not stop.

HiRDB linked to the cluster software enables the world's fastest level of stable failover in a few dozen seconds.

New transactions can be queued during system switchover so that the end user does not even notice the failure.

New transactions that are accepted during system switchover of a few dozen seconds can be held until the service restarts. This enables service to continue without requiring end user intervention.

Even if a failure occurs, the workload can be distributed to the remaining running servers to enable both high performance and high availability.

HiRDB uses a Shared Nothing architecture that can improve parallel processing capability in a linear manner by increasing the number of servers. HiRDB also supports Active-Active Cluster (for system switchover that requires no standby server) in which, if a failure occurs, the workload can be distributed to the remaining running servers. This reduces idle assets and enables both the high reliability and high availability.

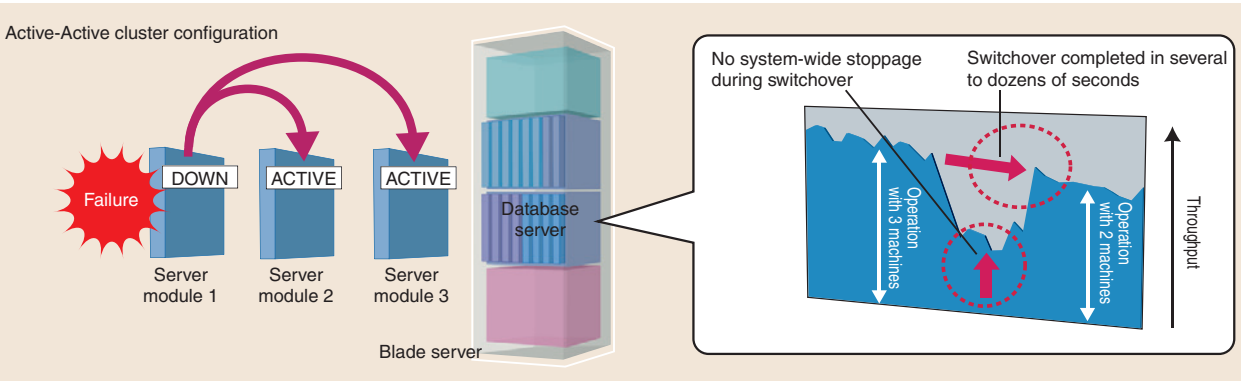

* Note: The optional products HiRDB Advanced High Availability and HA Monitor are required. Only a Japanese-language version is available for HA Monitor.

No data loss! Advanced disaster recovery

Even if a wide-area disaster occurs, the transaction results are guaranteed to ensure online service performance.

With HiRDB disaster recovery, the database is sent asynchronously to the remote site, so that only the system logs (update information) are transferred synchronously. This special "real-time SAN replication" functionality can minimize impact to online services running on the main site, for operation in which performance is barely degraded at all. HiRDB's own checkpoint methodology is also used to guarantee restart points on the remote site. This enables performance to be enabled with zero transaction defects, so that systems can be set up to guarantee valid transaction results.

Verification environment

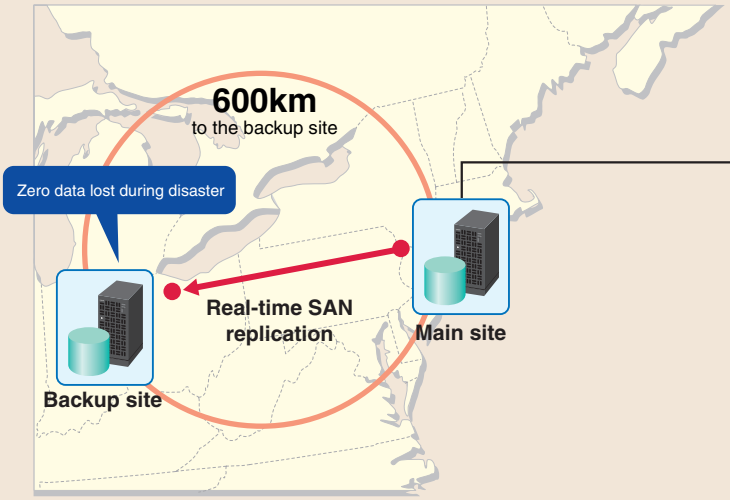

Relationship between throughput and distance on the main site

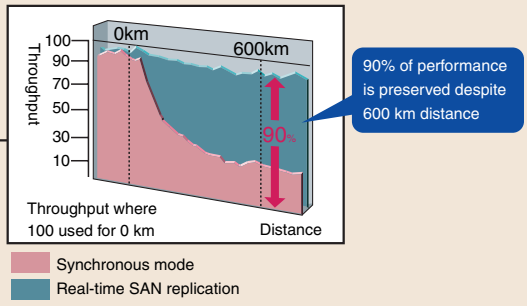

Differences between the synchronization methods

Data transfer method	Main site online functionality	Data loss during failure
Synchronous mode	C Degrades depending on the distance	A Zero data loss
Real Time SAN Replication	B Does not degrade	A Zero data loss

* Note: This R&D was partly supported by e-Society Leading Project (Development of foundation software: Storage Fusion, Tokyo Univ. and Hitachi, Ltd.) of the Ministry of Education, Culture, Sports, Science and Technology.

* Real-time SAN replication refers to the hybrid methods and log synchronization methods

Robust security through WORM functionality and audit trails

Protect important data and verify the validity of data.

HiRDB supports an audit trail function that can record an audit log containing the connection authentication to HiRDB and the access log for specified database resources. This function can find and track any invalid operation.
HiRDB also has the Write Once Read Many (WORM) function that only permits adding data to the specified table. This function prohibits even the table owner or database administrator from deleting or updating data. This protects important data from falsification or invalid access, thus improving system security.

ISO/IEC15408 certification
HiRDB has been certified as EAL4+ALC_FLR.1 -conformant, based on the ISO/IEC15408 Information Technology Security Assurance Requirements.
- HiRDB/Single Server Version 8 08-04 and HiRDB/Parallel Server Version 8 08-04

* ISO/IEC 15408 is an international standard for objectively evaluating security functionalities implemented in IT products. HiRDB, which has obtained this certification, allows you to use a database system that provides superior security.

Supported platforms

- Red Hat Enterprise Linux
- Windows ® XP
- Windows Vista ®
- Windows ® 7
- Windows Server ® 2003
- Windows Server ® 2008
- AIX
- HP-UX
- Solaris

*: Requires a 64-bit environment for operation.

SQL support

- ISO SQL92 (EntrySQL-compliant)
- ISO SQL99 (object-oriented support)
- ISO SQL2003 (support for SQL/XML linkage with XQuery)
- Integrity constraint
- Blocking transfer facility
- SQL Statements for manipulation using array variables
- Holdable cursor
- SQL optimization option
- Stored procedure/stored transaction
- Trigger
- Row value expression
- Case expression (conditional value expression)
- Outer join
- Derived table as table reference

etc.

Programming interface

- JDBC (JDBC 2.0)
- ODBC (ODBC 3.5)
- ADO.NET(ADO.NET 2.0)
- ADO
- OLE DB(OLE DB 1.1)
- DAO
- RDO
- Embedded SQL (supports C, C++, Java ™, COBOL85, and COBOL2002)

■ Information service ■

Global site: **http://www.hitachi.co.jp/hirdb-e/**

Japanese site: **http://www.hitachi.co.jp/hirdb/**

Hitachi, Ltd.

Hitachi, Ltd., Software Division
5030 Totsuka-cho, Totsuka-ku, Yokohama-shi, Kanagawa-ken, 244-8555 Japan
E-mail: WWW-mdc@itg.hitachi.co.jp

Mar., 2011