


IBM UniVerse



Highlights

- **Scalable and cost-effective open client/server architecture for both small and large deployments**
- **Flexible development environment with support for industry standards and multiple platforms**
- **Comprehensive toolkit for enhanced developer productivity and systems maintenance.**

Conceive, build, deploy — with UniVerse

To realize the reliability and scalability that result from today's open, relational database technology, companies need well-integrated information systems that tie together databases, business intelligence applications and e-commerce solutions. The IBM UniVerse® extended relational database solution is such a system, offering the capabilities you need to produce, enhance and deploy high-performance business applications in open-systems environments. Standards-compliant and flexible,

UniVerse provides data access, storage and management capabilities across Microsoft® Windows NT®, Linux® and UNIX® platforms.

UniVerse is the core of a complete development solution that offers a full suite of tools and services for:

- *Business intelligence, data marts and data warehouses*
- *Distributed database and enterprise-connectivity solutions*
- *E-commerce solutions.*

A versatile solution, UniVerse can be easily adapted to a wide range of situations. Its scalable, cost-effective open client/server architecture is suited for both large and small organizations. UniVerse provides the core functionality developers need to quickly build and deploy OLTP applications. It is an ideal database for client/server and Web-based application development because its nested relational data model and ability to dynamically and automatically

Keep mission-critical applications running smoothly

resize files help minimize database maintenance. UniVerse supports a wide range of hardware platforms and industry-standard interfaces. It also maximizes interoperability through support of both industry and de facto standards, including fully integrated ANSI-standard structured query language (SQL), ODBC, ActiveX and Java™ database interfaces.

Rapid development of line-of-business applications

Designed to be developer friendly, UniVerse features a native, high-performance, record-oriented, non-SQL programming paradigm. When used in conjunction with the UniVerse stored procedure database environment, it provides a rapid and real-world approach to application development that minimizes development time and maximizes operational performance.

Multiplatform support delivers high performance and scalability

Using symmetric multiprocessor (SMP)-based hardware, UniVerse scales from the smallest workgroup to thousands of concurrent users. UniVerse solutions for both UNIX and Microsoft Windows NT are economically sized and priced.

Tuned for resource efficiency

Automatic and dynamic tablespace allocation to fit the size of the database and tables provides optimum performance and minimizes disk usage for maximum efficiency. This capability also streamlines traditional database monitoring, reconfiguration and maintenance tasks.

New in UniVerse Version 9.6

UniVerse Version 9.6 offers new features for server-to-server interoperability, enhancements to the development environment, configurable Century Pivot date as well as core engine enhancements.

- **Server-to-server interoperability.** *To help you function efficiently in today's Web environment, UniVerse gives you the ability to communicate with and even drive Web-based applications from your UniVerse application. Using IBM's common middleware infrastructure, UniVerse now offers CallHTTP, a Basic client interface to any HTTP server that supports level 1.1 HTTP protocol, including URL encoding, multipart formats and redirection. CallHTTP includes API functions to provide application logging.*

The UniVerse Basic Sockets API provides the ability to interact with an application running on another machine through the socket interface. The Sockets API enables you to write distributed Basic applications by supporting both server and client interfaces that can cooperate on tasks through efficient, easy-to-implement socket communication.

- **Development environment.** *UniDebugger is the graphical integrated development environment (IDE) that interfaces with UniVerse Basic and its RAID debugger, providing a Windows®-based workspace with multiple windows to show the source code, application, breakpoints, watched variables and any messages from the debugging session. The editor provides full-screen, cut-and-paste and keyword color-coding. The debugging interface allows you to set break and watch points, execute line by line, or step into and over subroutines.*

- **Century Pivot Date configuration.** UniVerse Version 9.6 offers a new configuration option to modify the interpretation of two-digit years upon input. The standard default interprets years 00-29 as 2000-2029 and years 30-99 as 1930-1999. With the new Century Pivot Date configurable option you can change the default interpretation to more closely match user needs or industry standards.
- **Core engine enhancements.** Major core engine enhancements in UniVerse Version 9.6 improve usability and stability. These include a 32-bit replacement for UV/Term and Dynamic Connect, which provides full support for device licensing.

Administration

- Online backup by record, file, table and database
- Warm-start recovery
- High-availability disk subsystems
- GUI administration tool.

Large database support

- Horizontal table partitioning
- Multiple segment partition tables
- High-speed bulk-data loader
- Parallel index build.

SQL implementation

- ANSISQL 1989 level 2 compliant, with integrity enhancements
- ANSISQL 1992 entry-level compliant
- ODBC 2.0 core level API
- Robust SQL extensions
- ANSI declarative entity and referential integrity constraints.

Concurrency control

- Contention-free queries
- Table- and row-level locking
- ANSI 1992 Isolation levels (1-4)
- Implicit/explicit lock management.

Database characteristics

- Variable-length columns and rows
- Automatic tablespace management
- Clustered tables and indexes
- Automatic unique indexes by primary key
- Hash indexes
- NF2 extensions
- B-tree secondary indexing
- Calculated column definitions
- Cost-based/rule-based optimization.

National language support

- Unicode standard support
- Simultaneous accommodation of single and multibyte characters
- Per session control of character mapping and national conventions
- Support for national conventions, including date/time, monetary, collate, casing and numeric representation.

Transaction management

- Automatic roll forward and roll back
- Journaling (pre-image)
- Selectable ANSI 1992 Isolation levels (1-4)
- Deadlock detection and resolution
- Transaction logging
- Check-pointing and warm-start recovery.

SQL procedures

- Shared, compiled form
- Easy-to-use Basic
- Nested transactions
- Appropriate isolation levels
- Callable from all development interfaces
- Multiple, parallel procedure execution
- Interactive debugging facilities
- RPC support.

Database triggers

- Fired once per row after any file I/O operation, including INSERT, UPDATE and DELETE
- Support for BEFORE, AFTER, ENABLE and DISABLE syntax
- Uses the UniVerse Basic stored procedure language
- Modeled on ANSI SQL3 specification.

Distributed database

- Full location transparency
- Transparent multisite distributed queries
- RPC support
- TCP/IP networking.

Data replication

- *Publication and subscription methodology*
- *Automatic multisite, multitable replication*
- *Continuous or snapshot options*
- *Hot standby mode.*

Programmatic ODBC gateways

- *Supports any ODBC 2.0-compliant driver*
- *Available for Windows NT and UNIX ODBC driver managers.*

Interfaces

- *Windows ODBC drivers*
- *OLEDB Provider*
- *Object-oriented native interface*
- *ActiveX control (OCX)*
- *100% Pure Java API.*

Development tools

- *Visual Basic, Visual InterDev, Basic, C++*
- *Delphi*
- *Any ActiveX- or ODBC- or OLE.*

Database-enabled tool

- *Any Java IDE.*

Additional tools for UniVerse

With UniVerse, a wide range of additional development tools are also available to developers:

- *RedBack[®]—multithreaded, rules-based application server for building and delivering e-commerce solutions for OLTP*
- *MITS—native on-line analytical processing (OLAP) tool optimized for use with multi-valued databases for gathering, managing, distributing and analyzing data for improved strategic and tactical decision-making*
- *SB+—a 4GL development environment for host-based and client/server applications*
- *SBClient—the SB+ client component of a client/server solution that brings GUI features to host-based applications*
- *wIntegrate—a GUI application revitalization and desktop integration tool.*

Hardware and software requirements

The UniVerse extended relational database solution runs on major uniprocessor and SMP-based hardware platforms. It can run on Microsoft Windows NT/2000 or UNIX. UniVerse requires a suitable mass-media device such as a CD-ROM drive.

With UniVerse you have an advanced and cost-effective database management system that exhibits outstanding performance, increases developer productivity and simplifies system administration, making it the ideal platform for application development and deployment.

For more information

Please contact your IBM marketing representative or an IBM Business Partner, or call 1-800 331 1763 within the U.S. Also, visit our Web site at ibm.com/software/data/u2



© Copyright IBM Corporation 2001

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

Printed in the United States of America
11-01
All Rights Reserved

The e-business logo, IBM, the IBM logo, RedBack and UniVerse are trademarks of International Business Machines Corporation in the United States, other countries or both.

Microsoft, Windows and Windows NT 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.

Linux is a registered trademark of Linus Torvalds.

All other products or product names are trademarks or registered trademarks of their respective owners.

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.



GC27-1569-00