

Oracle Berkeley Database 6.2



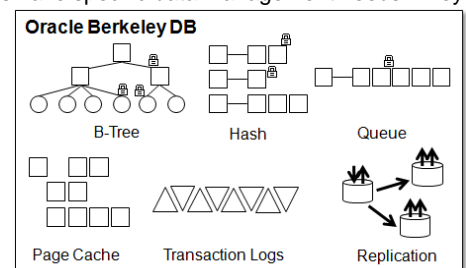
Oracle Berkeley DB is a high-performance embeddable database providing SQL, Java Object and Key/Value storage. Berkeley DB offers advanced features including transactional data storage, highly concurrent access, scalability on large SMP systems, replication for high availability, and fault tolerance in a self-contained, small footprint software library

Embedded Data Management

Embedded, mobile and edge applications have specific data management needs. They run unattended – in routers, mobile phones and software infrastructure. An embedded database must be a highly robust, secure, and resource-frugal solution with support for multiple APIs and storage options to accommodate a wide variety of challenges.

Embedded data storage can be

straight forward, simple and work every time without exception – if you choose Berkeley DB.




CONSIDER BERKELEY DB WHEN:

- Database services in a self-contained software library
- Access SQL data using the SQLite3 API, ODBC, JDBC, and ADO.NET
- Use FullText indexes to search SQL databases
- Use Rtree indexes to manage spatial data in SQL databases
- Manage data as key/value pairs or using SQL
- Btree, Queue, Recno, Hash data indexes
- Concurrent access by multiple threads or processes
- Recoverable ACID transactions with multiple isolation levels (including MVCC), nested transactions and long running transactions.
- Replication for high availability, fault-tolerance, fail-over, read-scalability,.
- Partition data based on key ranges
- Support for compression of data items
- Support for database file encryption, hardware

Berkeley DB satisfies the demands of this new generation of mobile and device applications by providing the same features – transactions, failure recovery, high concurrency, scalability and replication for high availability – expected of client/server enterprise-scale SQL databases. Berkeley DB is designed for zero oversight in deployment, automating all administrative tasks with programmatic APIs. If you do not need SQL then use the key/value API for data storage. Berkeley DB manages key/value data in any format, encoding or schema; your application implements its own optimal data storage format. Berkeley DB meets your mobile and embedded data management requirements regardless of how complex they might be.

Relational Storage, SQL Access

Berkeley DB supports SQL. Manage relational data and access it with the industry standard APIs ODBC, JDBC, ADO.NET or using the SQLite3 API. Berkeley DB's SQL implementation is, in fact, SQLite's SQL processor



acceleration on Intel

- Store data and transaction logs in memory or on disk or both
- Pure in-memory data management for ultra-low latency and high performance
- Zero manual administration, all tasks performed programmatically
- Scales to 256 TB of data in a single data file, multiple data files supported
- Efficient concurrent access on multi-core and multi-processor systems
- Library size less than 1MB
- Source code available

KEY BENEFITS

- SQL or key/value access
- Local, in-process storage
- High concurrency
- Mobile device support
- Transactional data integrity
- Automatic Recovery
- Replication for high availability
- Zero manual administration

ORACLE BERKELEY DB PRODUCT FAMILY

- Berkeley DB
- Berkeley DB Java Edition
- Berkeley DB XML

RELATED PRODUCTS


- Oracle NoSQL Database
- Oracle Database Mobile Server (DMS)

layered on top of the Berkeley DB data management library. This brings together the best small-scale SQL92 engine with the best transactional indexed data storage library. The result is a perfect database for mobile and embedded systems.

Replication

Berkeley DB provides a single-master, multi-replica highly available database configuration.

Transactional data is delivered to all replica nodes with flexible consistency policies per transaction. In the event the master replica node fails, a PAXOS-based automated fail-over election process minimizes downtime. Replication enables read scalability, fast fail-over, hot-standby and other distributed configurations, giving you enterprise grade features in a small, embedded package.




Performance

Berkeley DB is designed for high-performance, transactional data management. The in-process architecture enables speed and reduces complexity. The Berkeley DB engineering team has optimized the software for more than fifteen years with countless real-world deployments in mission-critical applications. A rich set of APIs allow developers to tailor caching, locking, logging and other crucial subsystems to deliver outstanding performance without sacrificing reliability, regardless of the runtime environment – large or small.

Synchronize Relational Data with Oracle Database

Berkeley Database helps to extend the reach of existing applications to mobile devices by supporting unparalleled performance and a robust data store on the mobile device. Oracle Database Mobile Server delivers critical bi-directional data synchronization capability to mobile devices, while providing a centralized backend interface for managing mobile deployments.



Flexible, Yet Predictable


Because requirements vary so widely, a database system must offer developers very fine-grained control over configuration, policies and resource allocation. At the same time, modern mobile and ISV software will always require concurrency, reliable recovery, and ease of adoption. Berkeley DB delivers on these requirements while providing predictable performance regardless of database size. Berkeley DB is the only database able to meet your needs without compromise.



CONTACT US

For more information about [insert product name], visit oracle.com or call +1.800.ORACLE1 to speak to an Oracle representative.

CONNECT WITH US



blogs.oracle.com/oracle

facebook.com/oracle

twitter.com/oracle

oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0116 MS

