

ORACLE BERKELEY DB

OVERVIEW



- Embeddable database engine
- Stores data in application-native format
- Runs in-process with the application
- Fast indexed and sequential retrieval (Btree, Queue, Recno, Hash)
- Concurrent access by multiple threads or processes
- Full ACID transactions
- Replication for high availability
- In memory or on disk operation
- Zero administration
- Scales to 256 terabytes of data in a single table
- Footprint as small as 400KB

Oracle Berkeley DB is a high-performance embeddable database engine for developers of mission-critical systems. It runs directly in the application that uses it, so no separate server is required, and no human administration is needed. Berkeley DB offers advanced data management services, including concurrent access, transactions and replication for high availability and fault tolerance.

Embedded Data Management

Applications today have moved beyond the data center. They run at the “edge,” in the network infrastructure – routers, switches and mobile networks – and on end-user devices at work, at home and while mobile. Like the data center applications that preceded them, these applications often need fast, scalable and reliable storage services for the data on which they operate. Unlike those data center applications, however, these new applications must manage themselves, running with local data and without a systems administrator.

Oracle Berkeley DB was designed to satisfy the demands of this new generation of applications. It offers the same operating guarantees – transactions, failure recovery, high concurrency, scalability and replication for high availability – that customers expect from mission-critical, enterprise databases. The Berkeley DB products, however, link directly into the address space of the application that uses them. Berkeley DB stores and retrieves data in the format that the application uses, so no translation to an external relational format is required. Berkeley DB provides an easy-to-use programmatic API for data access, eliminating the overhead of SQL processing at runtime. Most importantly, Berkeley DB is invisible – once it is built into an application, the end user does not need to worry about installing, configuring or maintaining a database system.

Berkeley DB offers developers fast, reliable and scalable storage services for their application without the overhead of using a separate relational database server.

Performance

Berkeley DB is designed from the ground up for high-performance embedded database management. Its in-process architecture and its use of a straightforward direct-access API instead of ad-hoc query processing languages mean that data stays in the application at runtime, and fewer instructions are executed.

In addition, the Berkeley DB engineering team has optimized the software over more than a decade of real-world deployments in mission-critical applications. The caching, locking, logging and other crucial subsystems have been tuned to deliver outstanding performance without sacrificing reliability.

BERKELEY DB – ENTERPRISE EMBEDDED

BERKELEY DB DELIVERS ENTERPRISE-GRADE PERFORMANCE, SCALABILITY AND RELIABILITY AND IN AN EMBEDDABLE PACKAGE.

KEY BENEFITS:

- Very high performance
- Simple, direct, indexed storage
- Local, in-process storage
- High concurrency
- Massive scalability
- Transactional data integrity
- Automatic recovery
- Replication for high availability
- Zero administration

RELATED PRODUCTS AND SERVICES:

Berkeley DB is a part of the Berkeley DB embeddable database family, which includes Berkeley DB XML and Berkeley DB Java Edition.

Flexibility

There are enormous differences among applications that use Berkeley DB. Networking equipment is different than a high-speed cache on a Web server, and both are very different than a mobile phone. Yet the same Berkeley DB product is used within all of these applications. Because requirements vary so widely, the database system must offer developers very fine-grained control over configuration, policies and resource allocation, as required by the application.

At design time and at runtime, developers can configure Berkeley DB to use resources properly, and to enforce the policies that make sense for a given application. Single- or multi-user, with or without transactions, on a handheld device or a multi-processor server, Berkeley DB can be easily controlled to meet the application's requirements. This flexibility allows Berkeley DB to be used in situations that would otherwise call for a custom or home-grown solution.

Scalability

Berkeley DB scales up extremely well. It can manage multi-terabyte tables with single records as large as four gigabytes. It can support a large number of simultaneous users. Systems designers can choose between conventional write-ahead logging or multi-version concurrency control to reduce contention and improve performance. Berkeley DB makes efficient use of large-system resources, including very large physical memory and disk systems.

In addition, Berkeley DB scales down well. The code is less than half a megabyte in common deployments. Developers can configure it for operation in constrained environments such as mobile phones, with limited memory and CPU.

Reliability

Berkeley DB provides carrier-class reliability in an embeddable package. The transaction system guarantees that Berkeley DB can recover from software and hardware failures without losing data. A sophisticated replication system provides high availability and scalability for applications which may run on multiple blades in a chassis, servers in a datacenter or machines distributed around the world. In addition, database upgrades can be performed without bringing down the entire replicated system.

Enterprise Embedded

Berkeley DB provides sophisticated, reliable, and configurable data storage services to system architects and developers with the performance, scalability, and reliability everywhere the new generation of networked applications demands it.

Copyright 2006, Oracle. 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 is it 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, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.