

ORACLE BERKELEY DB PRODUCT FAMILY

OVERVIEW



CONSIDER ORACLE BERKELEY DB WHEN:

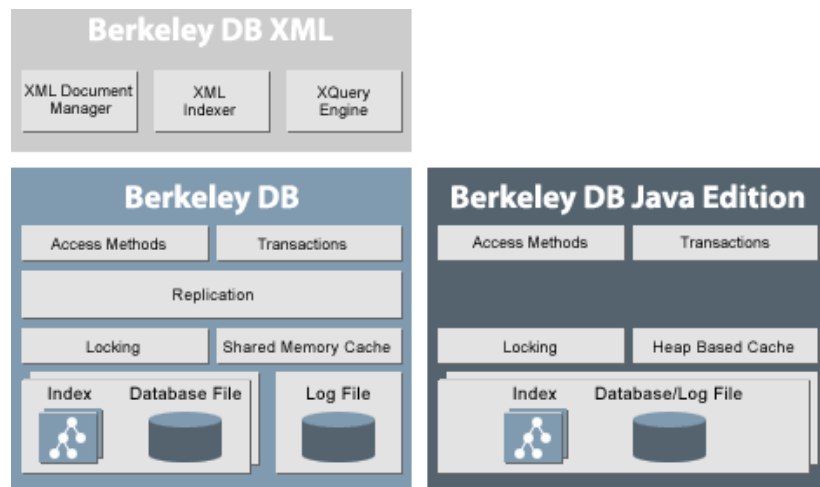
- Performance matters
- Unattended operation is a requirement
- Access patterns are predictable (static queries over dynamic data)
- Local, in-process storage is preferred
- Failure recovery must be automatic
- High availability and fault tolerance are critical
- Cost of ownership is a key factor
- Access to source code is important

The Oracle Berkeley DB product family consists of Berkeley DB, Berkeley DB Java Edition and Berkeley DB XML. All three are high performance, embeddable database engines for developers building mission-critical systems. They deliver superior performance, scalability and availability for applications that must store data locally and run unattended without human administration.

Overview

The Oracle Berkeley DB family of high performance, open source, embeddable databases allows developers to embed in their applications a fast, transactional database engine with industrial grade reliability, scalability and availability. For Independent Software Vendors, device and equipment manufacturers, and enterprises building applications for internal use, the Oracle Berkeley DB family of products provides fast, local persistence with zero administration.

Your customers and end-users will experience an application that simply works, reliably manages data, scales under extreme load, and requires no ongoing database administration. Your development team can focus on your application and be confident that Berkeley DB will manage your application's data.



Oracle Berkeley DB Product Family

The Oracle Berkeley DB products link directly into your application, eliminating the performance penalty of client-server architectures. Berkeley DB provides simple function-call interfaces, allowing programmers to store and retrieve information quickly and easily. Data is stored in the application's native format, eliminating the need for translation or mapping.

ORACLE BERKELEY DB PRODUCT FAMILY

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

ENTERPRISE-GRADE
PERFORMANCE,
SCALABILITY AND
AVAILABILITY IN AN
EMBEDDABLE PACKAGE

Berkeley DB supports full ACID transactions and recovery for data integrity, multiple processes and multi-threading for high concurrency, and replication for high scalability and availability. It is lightweight, designed for high performance, scales up to carrier and enterprise class applications and scales down to mobile devices and other constrained environments.

Berkeley DB is very flexible and configurable, giving the application developer control over how resources are allocated, the amount of memory dedicated to caching records, the degree of concurrency, support for recovery and more. It includes full source code for easier porting, integration, debugging and optimization.

It is also important to understand what Berkeley DB is not. It is not a database server that handles network requests. It is not an SQL engine that executes queries. It is not a relational or object-oriented database management system. It is an embeddable database engine that has been designed to be simple, fast, small and reliable.

Berkeley DB supports all modern operating systems including Linux, UNIX, Windows, Mac OS X, and RTOS such as VxWorks and QNX. It supports a variety of programming languages, including C, C++, Java, Perl, Python, PHP, Tcl, Ruby and others.

Dual License

The Berkeley DB products are available under a dual license model, which offers great advantages for customers. An open source license permits customers to use Berkeley DB at no charge under the condition that if a customer uses Berkeley DB in an application they distribute to a third party, the complete source code for the application must be available. A commercial software license is also available and permits the customer to distribute their application without releasing their source code.

Applications That Use Berkeley DB

Berkeley DB is used within a wide variety of open source and proprietary products for telecommunications, networking, storage, security, enterprise infrastructure, financial services and many other markets. There are over 200 million copies of Berkeley DB in deployment. Open source projects that use Berkeley DB include every version of Linux, BSD UNIX, Sendmail, OpenLDAP, Apache, and many other open source projects. Proprietary products that use Berkeley DB include directory servers from Sun and Hitachi; messaging servers from Openwave and LogicaCMG; switches, routers and gateways from Cisco, Motorola, Lucent, and Alcatel; storage products from EMC and HP; security products from RSA Security and Symantec; and Web applications at Amazon.com, Google and AOL.

CUSTOMER BENEFITS OF THE DUAL LICENSE MODEL

- Huge user community
- Very high code quality
- Easier integration, debugging, porting and tuning
- Easy to download, evaluate and develop
- Legal assurances, warranties and indemnification
- Technical support
- Roadmap upgrades
- Single vendor to hold accountable

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.