

# ORACLE BERKELEY DB JAVA EDITION




- Pure Java database
- Fast, indexed, Btree storage
- ACID transactions
- Highly concurrent design
- Runs in-memory or on disk
- Zero oversight administration
- Supports Java EE JTA, JCA, JMX, and JConsole
- Simple persistent object graph Direct Persistence Layer (DPL) annotation API
- Work with large persistent datasets using the java.util.collections API
- Easy Key/Value API
- Replication for high availability, fault-tolerance, fail-over, read-scalability, etc.
- Scales to hundreds of gigabytes of data
- Designed to scale up as you add cores, CPUs and even systems
- Support for Android Platform and Dalvik JVM
- A single JAR file (~0.5-2KB in size depending on configuration)

*Oracle Berkeley DB Java Edition is a high performance, transactional database in pure Java. Using Berkeley DB Java Edition you can persistent objects using the Java Collections API, as annotated graphs or as Key/Value pairs in a fast, local database. Oracle Berkeley DB Java Edition is an easy-to-use, scalable and efficient pure Java database.*

## Persistence


Berkeley DB Java Edition provides all the features of an object-to-relational mapping solution and a transactional database without the overhead. Most, if not all, Java applications need to store some amount of object data. Sometimes that data must be accessible by other software infrastructure. In these cases simply map the object data into SQL using an ORM such as EJB and let the relational database process and persist the information. Other times this design pattern (Java – ORM – RDBMS) is used when the advantage of an intermediary format and the analytical power of a query language is never realized. If you do not need the overhead of SQL processing, why use an Enterprise Java Beans (EJB)? Berkeley DB Java Edition accomplishes the goal of transactional Java object storage without translating objects into SQL thus dramatically improving your application's performance without sacrificing any functionality.

Berkeley DB Java Edition stores and optimizes access to your Java objects, it provides concurrent transactional data storage, and data is indexed for fast retrieval just as you would expect from any database. The difference is that it does all this in a small, efficient, and easy-to-manage package.



## Scale Out and Manage Failure Using Replication

Berkeley DB provides a single-master, multi-replica highly available database replication. 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. This allows for read



A SCALABLE, FAST AND TRANSACTIONAL JAVA PERSISTENCE SOLUTION

USE CASES:

- Caching
- Application data repository
- POJO persistence
- Queuing/Buffering
- Web Services
- Mobile Devices
- Integration

KEY BENEFITS:

- Very high performance
- Simple, direct, indexed storage
- Local, in-process storage
- High concurrency
- Massive scalability
- Transactional data integrity
- Automatic recovery
- Zero administration
- Fast, simple development
- Flexible deployment

RELATED PRODUCTS:

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



scalability, fail-over, hot-standby and other distributed configurations, giving you enterprise grade features in a small, embedded package.

**Performance**

Berkeley DB Java Edition’s unique log-based database structure and concurrent Btree design brings performance into parity with that of optimized ANSI C

solutions while remaining pure Java.

Berkeley DB Java Edition scales across multi-core and multi-processor architectures to efficiently use all available resources.



**Java Enterprise Edition Integration**

Berkeley DB Java Edition fits neatly into the Java Enterprise Edition (Java EE) programmer’s toolkit by supporting the Java Transaction API (JTA), Java EE Connector Architecture (JCA), JConsole API for monitoring, and Java Management Extensions (JMX) on Oracle Application Server 11g and other popular Java EE application servers.

**Mobile to Massive**

Berkeley DB Java Edition fits equally well into any stand-alone Java application, an application on a mobile device, or within a globally distributed enterprise application running within the context of a Java EE server. Berkeley DB Java Edition fills a major gap in the data storage landscape by offering an out-of-the-box solution for situations where previously the only solution was custom code.

**Simple, Effective, Fast**

Berkeley DB Java Edition is simple to use, efficient and effective Java object storage. When requirements dictate transactional persistence consider eliminating the SQL processing overhead by using Berkeley DB Java Edition.

Copyright 2010, 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. Java and the Java Powered logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries, and are used under license. Other names may be trademarks of their respective owners.