

# ORACLE BERKELEY DB XML

## OVERVIEW



- Embeddable XML database engine
- Support for XQuery and XPath
- Flexible indexing for optimizing search and retrieval
- Flexible storage control – nodes or whole documents
- Partial document updates
- Schema support, including validation
- Runs in process with the application
- Concurrent access by multiple threads or processes
- Full ACID transactions
- Replication for high availability
- Zero administration

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

### Embedded XML Data Management

Oracle Berkeley DB XML offers software developers a fast, reliable and scalable XQuery engine with sophisticated XML content storage in a packaged that can be embedded inside an application. Berkeley DB XML is equally at home managing content at nodes within a Service-Oriented Architecture (SOA), storing data for a PHP Web application, and managing massive volume of genomic or geographic data requiring complex query analysis.

XML is a widely used format for publishing and exchanging data among systems, in particular for Web Services and SOAs. As the volume of XML data increases, it has become increasingly critical that applications and SOA components be able to store, manage, and retrieve XML without the overhead of translating it into other formats. Berkeley DB XML allows developers to do just that quickly and easily within their applications, without resorting to an external storage system or expensive and time consuming setup and administration. Berkeley DB XML runs directly inside the application that uses it, so no separate database server is required. Berkeley DB XML stores XML documents in XML, and supports the XQuery standard for querying data.

### Enterprise-class Storage

Because it is built on top of the proven and mature Oracle Berkeley DB storage engine, Berkeley DB XML inherits all the features, performance, reliability and scalability that Berkeley DB delivers. Berkeley DB XML can store terabytes of data in a single document repository. The database can survive power outages, software errors and hardware failures without losing data. It supports simultaneous access to the document store by large numbers of users.

### Designed for Networked Applications

One of the most common uses of Berkeley DB XML is for storing documents in network-based applications. Content management systems, which store and publish XML documents over the Internet, may use Berkeley DB XML as their repository. High-performance networking devices like switches, routers and gateways use

## BERKELEY DB XML – EMBEDDED XML

BERKELEY DB XML DELIVERS ENTERPRISE-GRADE STORAGE SERVICES FOR XML-BASED APPLICATIONS.

### KEY BENEFITS:

- High performance
- Direct storage of XML without the need for translation
- Standard XQuery-based document storage
- Flexible indexing and storage model for optimizing performance
- High concurrency
- Transactional data integrity
- Automatic recovery
- Replication for high availability
- Zero administration

### RELATED PRODUCTS AND SERVICES:

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

Berkeley DB XML to manage control information and the data that flows through them.

A common requirement of these systems is a fast, local, persistent store for XML data. Low latency and high throughput are crucial. Berkeley DB XML eliminates runtime overhead by operating within the embedding application thereby providing direct memory access, rather than client/server based access, to XML content. This eliminates the round trip to and from the server during runtime, and translates into better performance and scalability.

### Flexibility

Berkeley DB XML is intended for use within an application, for applications that need to store data locally and run in an unattended fashion. The variety of applications that use it is broad; ranging from SOA components to network storage products to financial applications. Berkeley DB includes an Apache module for use as an out of the box Web-Service. PHP is also supported, greatly enhancing XML-based web sites.

Berkeley DB XML gives developers control over virtually every aspect of the deployed system. It can be configured to meet the needs of the application, and to operate within the constraints of the environment in which it is deployed.

For example, Berkeley DB XML allows the developer to declare a wide variety of indexes on the document repository. Individual nodes in a document, the presence or absence of edges in the document graph and the value of named attributes may all be stored in an index for fast retrieval at query time.

In addition, Berkeley DB XML allows the developer to decompose a document into smaller units – sub-trees – for storage. This improves update speeds, since only part of the document need be rewritten when it changes. In addition, this can improve read performance for queries that want part of, but not all of, a document.

### Embedded XML

Berkeley DB XML provides systems architects and developers with data management services that are optimized for XML documents. Its support for XML schema, and for XQuery/XPath for document manipulation, mean that Berkeley DB XML drops seamlessly into an XML-based application. Because Berkeley DB XML runs directly in the application, it provides superior performance, and does not require an administrator to keep the system running.

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.