

Oracle Berkeley DB: Performance Metrics and Benchmarks

ORACLE WHITE PAPER | JANUARY 2017





Introduction

This white paper presents performance measurements for Oracle Berkeley DB under various conditions on several major operating systems and platforms. Performance estimates presented here are intended to help you understand what to expect from Berkeley DB in some common configurations on commodity hardware running widely used operating systems. Your application performance depends on your data, data access patterns, cache size, other configuration parameters, operating system and hardware. Benchmarks are almost never representative of the performance of a particular application. They do, however, provide guidelines and help to set basic operational expectations.

Oracle Berkeley DB is a database engine designed for any application requiring efficient and robust embedded data storage. Berkeley DB is a complete data management library, including support for concurrent access, ACID transactions and recovery, and replication for fault tolerance and high availability. Berkeley DB is designed to support any structured information model and has been incorporated into relational, object, network and hierarchical systems. Due to its scalability and minimal overhead, it has also been used in a number of reliable messaging, data streaming, caching, and data archival systems.

Platforms and Testing Methodology

All performance tests were run on the following commodity hardware and operating system combinations, without customized configuration, tuning or any other performance enhancing techniques. All tests were run on clean installations in multi-user mode systems that were otherwise quiescent. For the workloads presented, these are the results that can be obtained under actual use. Performance is often the first critical factor when selecting a database. This paper measures the performance of Oracle Berkeley DB across a wide set of operating systems, hardware platforms, and software configurations.

The following platforms represent common operating systems and hardware platforms chosen by software and systems developers. It is by no means an exhaustive list, as Berkeley DB runs on a much larger range of platforms. The version of Berkeley DB tested was 4.3.

- » Linux – SuSE Linux 9.1 running on an AMD Athlon™ 64 processor 3200+ at 1GHz system with 1GB of RAM. Western Digital Caviar SE 120GB EIDE (7200RPM) with 8MB cache RAM was connected via an SIS 5513 controller formatted with the EXT3 file system with journaling enabled. The compiler was GCC 3.3.3. The system reported 4046.84 BogoMIPS.
- » Solaris – Sun Solaris 10 (x64) running on a SunFire v20z with dual AMD Opteron™ processors at 2GHz with 16GB of RAM. A 80GB Seagate Cheetah (ST373453LC) SCSI hard drive was connected via an integrated SCSI controller on a PCI-X bus and formatted with the UFS filesystem.
- » BSD Unix – FreeBSD 5.3p4 running on an Intel™ Pentium™ 4 CPU at 3.00GHz with 2GB of RAM. An integrated Intel ICH5 SATA150 controller on a PCI bus controlled a Barracuda 7200.7 Plus (ST3200822A) Ultra ATA/100 (7200RPM) hard drive formatted with the UFS file system. The compiler was GCC 3.4.2.

- » Mac OS/X – Mac OS/X version 10.3.9 running on a dual Power Mac G4 CPUs at 1.25GHz with 2GB of DDR SDRAM. IBM Deskstar 120GXP 120GB UATA100 hard drive was connected via an integrated ATA bus formatted with the HFS+ file system with journaling enabled. The compiler was GCC 3.3.

Note that the systems vary greatly. This is intentional. It is not the goal of this paper to demonstrate the relative performance of the profiled systems when performing similar tasks, but rather to demonstrate basic performance metrics on a wide variety of systems so that you can estimate how your system might perform when using Berkeley DB.

Any transactional database claiming durability (the 'D' in ACID) will necessarily have to interact with some form of stable storage or replicate the transaction to other systems. The most common case is to use a hard disk drive. Hard drive technology, bus subsystems, and memory interconnects vary widely, as do their performance. For this set of tests, we used commodity disk technology – that is, slow disks with average subsystems. If you add a fast disk subsystem with the fastest disks, you'll see much higher throughput with Berkeley DB, just as you would with any disk-bound application. Additional hardware options can also affect performance. For instance, using a disk subsystem with a battery-backed memory write cache will improve throughput. There are other, faster storage systems, some based entirely on memory, but they are less common. These hardware-based techniques affect all databases in a similar manner and so are irrelevant for this study. This paper explores the performance curve of the Berkeley DB software, not the surrounding hardware.

The test code can be downloaded at:

<http://www.oracle.com/technetwork/database/database-technologies/berkeleydb/downloads/perf-186166.zip>

Performance

Throughput – records read or written in a fixed time interval – is the most common measure of database performance. Each test in this benchmark uses fixed-size records containing keys that are eight bytes long and a data portion of thirty-two bytes. Pages are 32KB, the buffer for bulk transfers was set to 4MB, the log buffer is 8MB, and the buffer cache is 32MB. Source code for the benchmark is roughly 350 lines of C code and may be downloaded from the URL above.

Berkeley DB supports concurrent database access. On multi-processor systems, multi-threaded or multi-process Berkeley DB applications can take advantage of additional CPUs. Many application workloads will benefit from such concurrency, even on single-processor systems. Regardless of the system tested, the effects of concurrency are not measured in the following tests. The tests are single-threaded and single-process. On a multi-processor system, they will use no more than one of the CPUs at a time.

This first test measures throughput as operations per second using Berkeley DB Data Store (DS). Since this is a non-transactional, in-memory test, performance is dependent on the operating system, memory, and CPU, and not the underlying I/O or disk system for the read tests. When writing using DS, the data is cached in memory and flushed when buffers are filled or when the application explicitly calls for data to be flushed to disk. These tests are written such that writes occur in memory and are not flushed to disk.

DS (ops/sec)	Linux	Solaris	BSD	Mac OS/X
Single-record read	1,002,200	1,008,580	1,108,920	524,360
Single-record write	766,304	550,748	614,116	317,141

Berkeley DB provides optimized access to sequential data allowing for faster overall throughput. If you read large amounts of data sequentially from your databases, using the bulk access APIs will improve your read performance.

DS (ops/sec)	Linux	Solaris	BSD	Mac OS/X
Bulk Read	13,501,800	12,114,500	19,299,700	5,615,280

Most applications using Berkeley DB will require transactional access to data as a result of concurrent access or simply to maintain a reliable repository of data. This is accomplished by using the Berkeley DB Transactional Data Store (TDS) features. In this first test of TDS we enable the locking, logging and transaction subsystems. Although read-only transactions do not perform logging, performing reads in transactions will introduce locking overhead. That overhead is reflected in the following test results.

TDS (ops/sec)	Linux	Solaris	BSD	Mac OS/X
Single-record Read	466,623	328,809	331,523	171,549

The transactional system can be configured to run entirely in memory. If the process fails, the system is powered down, or some other unforeseen failure occurs, the data will be lost. Essentially, durability is sacrificed for speed. Replicated systems running entirely in core memory using Berkeley DB High Availability (HA) can provide full ACID transactional semantics, including durability, as long as one replica remains active, and will perform much faster than systems storing data to disk. The following measures only the in-memory speed of the transactional system without replication.

TDS (ops/sec)	Linux	Solaris	BSD	Mac OS/X
Single-record Write	239,824	143,466	111,671	82,256

When designing your system, your requirements will determine what, if any, transactional guarantees you can relax. The following tests examine the performance of each tradeoff from the most relaxed to the most stringent. Applications requiring the "ACI" (atomic, consistent, and isolatable) characteristics of transactions but willing to risk the "D" (durability) can configure Berkeley DB such that it does not synchronize transaction logs to disk on commit. In this case, the transaction log buffers are maintained in memory and only written to stable storage when Berkeley DB determines that they are full, rather than as a component of the transaction. If the application, operating system, or hardware should fail, the data in those buffers that is not yet written to disk is lost. This "TDS no-sync" case is measured below.

TDS (ops/sec)	Linux	Solaris	BSD	Mac OS/X
Single-record Write	141,109	107,626	70,528	68,591

The performance is nearly identical to that of the "TDS in-memory logs" case. In neither case is disk access required in order to complete a transaction. The difference is that for the "TDS no-sync" case, the data is eventually written to disk.

Next, we consider a configuration that will guarantee no loss of data after an application crash, as long as the operating system does not also crash. TDS is again configured with locking, logging and transactions all enabled. However, in this case Berkeley DB requests that the operating system write to disk, but not explicitly flush those writes to disk at each transaction commit. The operating system will acknowledge the write and is likely to buffer writes for some amount of time before flushing the data to disk. In this mode, TDS is sacrificing some durability, but less than before.

TDS write, no-sync	Linux	Solaris	BSD	Mac OS/X
Single-record Write	90,774	85,101	62,087	34,888


Notice the significant drop in performance, as each transaction commit requires a call into the operating system and copies data from the application's memory to the operating system's write buffers. If we enable fully synchronous writes, throughput will be limited by your disk's ability to perform such writes. The resulting numbers are, therefore, completely determined by your disk system and not the software.

Note that it is possible to get spectacular numbers in this mode, because some disks report writes as having completed once the data has been written into the disk's cache, instead of when the data is actually on the disk. This performance boost comes at the price of durability guarantees, and we advise running critical applications on disk subsystems that do not respond to synchronous write requests until the data is persistently on the disk. Berkeley DB can optionally group commit operations, allowing higher throughput when doing synchronous transactions.

In conclusion, Berkeley DB allows application designers to make transaction durability tradeoffs when building a high throughput system, based on their design goals and the characteristics of the data involved. With an understanding of these and other tuning parameters found in Berkeley DB and in your deployment system it is possible to exceed these measurements.

ASPECTS OF BERKELEY DB NOT MEASURED

- » Out of cache performance – The impact of storage to disk is so great that when measuring database performance out of cache, you are essentially measuring the underlying operating system and associated hardware. The time spent waiting on disk accesses takes up the vast majority of the overall access time. As a result, the tests have been designed to work entirely in cache except for those times when it is necessary to write to disk for transactional guarantees. Systems with faster disks or disk controllers with built-in RAM-based cache will perform better.
- » Recovery time – Transactional databases must be able to recover to a known state at startup. Berkeley DB, like most transactional database systems, will recover in time proportional to the amount of data written to the log since the last checkpoint. The additional cost of taking the checkpoints is not reflected in these measurements.
- » Deadlock discovery and resolution time – Transactional databases can deadlock. Berkeley DB offers many different methods for deadlock detection and resolution. These benchmarks do not measure the cost of deadlock processing as they are written to be deadlock-free.
- » Startup time – A database's cache subsystem is one of the most performance-critical aspects of its design. There are two states for a cache: empty and warm. We do not measure the cost of warming the cache. All benchmark programs first warmed the cache and then began collecting timings.
- » Effects of concurrency – Multi-threaded applications often get better performance than single-threaded applications, because whenever any single thread blocks for I/O, others can use the CPU to do useful work. This is especially true in Berkeley DB, since the database library includes support for a feature called "group commit", which allows one active thread to do work on behalf of other threads when one or more transactions commit. Concurrency also introduces opportunities for contention for locks, buffers or other resources. In general, multi-



threaded applications see performance improvements until contention among too many threads causes performance to fall off.

- » Access Methods – Berkeley DB provides a few different access methods. Each method provides a different performance characteristic. The measurements in this paper use the btree access method because it is the fastest access method for a wide range of workloads and data sizes.
- » Replication – The Berkeley DB High Availability subsystem is a single master, multi-replica system that supports fail-over and master election using a Paxos-compliant election algorithm. It replicates transaction logs among systems. The transport of these logs is application-specific. To replicate database information, Berkeley DB High Availability uses a transport layer provided by the application to relay coordination messages and log segments among replicas. Any form of message transport can be employed. The application designer may choose among hardware-specific transport, TCP/IP or any other communications infrastructure. For this reason, measurement of the High Availability system is beyond the scope of this paper.
- » Comparing Oracle Berkeley DB to Relational Database Management Systems

Client/Server architecture increases complexity and slows down performance because applications must cross a process boundary and often a network link to access data. SQL query processing is extra code that a relational engine must execute at runtime. Parsing, planning, optimizing and executing a query introduce runtime overhead to each data access. Oracle offers a wide range of SQL-capable database management products for applications that need the expressive power and flexibility of SQL. For those that need fast, reliable, embeddable data management engine without a human database administrator, however, Berkeley DB is often a more attractive choice. Berkeley DB delivers the most direct route to data and avoids cross-process copies, query parsing and planning overhead found in relational systems.

Naturally, head-to-head comparisons depend heavily on the application and its data management requirements.

Conclusions

Oracle Berkeley DB, as a component of your application, can be configured in a number of different ways to meet specific application requirements for speed, scale, and reliability. This paper has taken a look at a few of the critical aspects of any database system and shown that Berkeley DB has benefited from years of research, design and careful tuning.

Berkeley DB provides outstanding speed, efficiency and flexibility for a wide range of applications. Berkeley DB offers extraordinary programming flexibility for applications ranging from mobile phones to desktop applications, large servers, networking gear and major websites. Download Berkeley DB today at

<http://www.oracle.com/database/berkeley-db.html> , and run the performance tests on your platform.







Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

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

Oracle Berkeley DB: Performance Metrics and Benchmarks
January 2017
Author: Brey/Schulman/Maik