

# Red Brick Warehouse:

## A Read-Mostly RDBMS for Open SMP Platforms

Phillip M. Fernandez  
Red Brick Systems  
485 Alberto Way  
Los Gatos, CA 95032  
philf@redbrick.com

### Overview

The Red Brick Warehouse™ is a commercial Relational Database Management System designed specifically for query, decision support, and data warehouse applications. Red Brick Warehouse is a software-only system providing ANSI SQL support in an open client/server environment.

Red Brick Warehouse is distinguished from traditional RDBMS products by an architecture optimized to deliver high performance in read-mostly, high-intensity query applications. In these applications, the workload is heavily biased toward complex SQL SELECT operations that read but do not update the database.

The average unit of work is very large, and typically involves multi-table joins, aggregation, duplicate elimination, and sorting. Multi-user concurrency is moderate, with typical systems supporting 50 to 500 concurrent user sessions.

Query databases are often very large, with tables ranging from 100 million to many billion rows and occupying 50 Gigabytes to 2 Terabytes. Databases are populated by massive bulk-load operations on an hourly, daily, or weekly cycle. Time-series and historical data are maintained for months or years.

Red Brick Warehouse makes use of parallel processing as well as other specialized algorithms to achieve outstanding performance and scalability on cost-effective hardware platforms.

### Execution Environment

Red Brick Warehouse is designed to run on a wide variety of “commodity” open-systems symmetric multiprocessing platforms from vendors such as Hewlett-Packard, Sun, IBM, Sequent, and NCR. Red Brick Warehouse currently runs under Unix, but is directly portable to other operating system environments.

Current systems from these vendors include 2 to 32 or more processors, and have the performance, I/O bandwidth, and storage capacity required for very large applications, yet may be deployed for well under \$1 million.

### Parallel Processing

Red Brick Warehouse includes several parallel algorithms designed to match I/O concurrency and processor concurrency to the needs of large unit-of-work query and data loading requests.

Each table in a Red Brick Warehouse database is managed by a Parallel Table Server subsystem. Row data for a table are segmented by value or hash to dozens or hundreds of physical disk drives as necessary for performance and capacity. Table server subsystems use a process-per-disk architecture to achieve high I/O concurrency and use buffering and scheduling to minimize seeks and to maximize I/O transfer rates.

Red Brick Warehouse uses a process-per-user session. Because the average unit of work is large, this concurrency approach achieves excellent machine utilization and good overall multi-user performance. For computationally-intensive queries, finer-grained intra-query parallelism is used to achieve higher performance.

Red Brick Warehouse uses a proprietary multi-dimensional join index to quickly perform single-pass joins of 2 to 10 tables. A parallel version of this algorithm horizontally partitions a single STARJoin™ operation across 2 to 32 processors, and achieves near linear speedup as additional processors are applied.

Some query operations require relation scans, which can be horizontally partitioned and combined with parallel I/O for linear speedup as processors are added.

In multi-user environments oriented around a few very large tables, many users may need to scan the same table. A unique coordinated table scan capability allows single disk read requests to be shared among multiple scan operations executing concurrently on multiple processors, resulting in substantially sub-linear performance degradation as additional concurrent users are added.

Bulk data loading is also performed in parallel using a pipelined approach. Two to 16 processors may be applied to a single load stream, with individual processors performing concurrent tape I/O, disk I/O, data conversion, referential integrity checking, and index building.