

eXtremeDB® Cluster


Clustering real-time database system distributes processing across multiple nodes to improve speed, scalability and reliability.

The first clustering DBMS designed for embedded software as well as high performance desktop, workstation and server-based applications.

eXtremeDB, the real-time embedded database for devices that are eXtremely innovative

Overview

eXtremeDB Cluster is McObject's distributed DBMS that manages databases across multiple hardware nodes, enabling two or more servers to share the workload. As the first clustering database solution designed for embedded systems and real-time enterprise applications, eXtremeDB Cluster delivers a cost-effective and powerful distributed database for systems including carrier grade telecom/networking equipment, capital markets applications, Web services and hosted Software-as-a-Service (SaaS) platforms.



eXtremeDB's in-process (rather than client/server) database architecture integrates the DBMS within the application process, resulting in lower complexity.

eXtremeDB Cluster dramatically increases available net processing power, reduces system expansion costs, and improves scalability and reliability. In McObject's benchmarks, eXtremeDB Cluster delivered an astonishing 161% throughput improvement when scaling to four nodes from one node.

Cost-Effective & Reliable

In an eXtremeDB Cluster deployment, every database instance serves as a "master." Any process on any node can update its local database, and the eXtremeDB Cluster software replicates the changes to other nodes. This more evenly balances the workload, compared to High Availability solutions consisting of a single master database and one or more read-only replicas.

The hardware for each eXtremeDB Cluster node can be a low-cost (i.e. "commodity") server, so that the system expands cost-effectively. Multiple platforms are supported within a single cluster.

Distributing the system across multiple hosts ensures continuous availability in the event of a failure on one node. eXtremeDB Cluster's "shared nothing" architecture eliminates reliance on a shared SAN or other storage resource. eXtremeDB Cluster supports the same ACID transactions offered by the non-clustering eXtremeDB editions, making it an attractive choice for applications that demand integrity of distributed data.

Lower Complexity Clustering Solution

Ask developers and IT managers about their clustering solutions' biggest drawback, and they'll likely say "complexity." Clustering technology based on relational database management systems (RDBMSs) can entail separate sets of nodes for clients, database servers and SQL processing, as well as dedicated server and client modules to address system management tasks. Assembling all these pieces to work seamlessly together often requires extensive consulting engagements.

Proven Building Blocks

McObject developed its clustering solution on a proven foundation of eXtremeDB technology:

In contrast, McObject's streamlined architecture integrates these functions on the nodes where eXtremeDB Cluster resides, making for a greatly simplified deployment. Designed to operate in field-based systems, eXtremeDB Cluster operates without the "care and feeding" of specialists such as the armies of RDBMS database administrators (DBAs) found in many companies. This elimination of setup and operational expenses contributes to eXtremeDB Cluster's very attractive total cost-of-ownership (TCO).

- **In-Memory Database System** - eXtremeDB's core in-memory design cuts I/O, caching and other overhead inherent in disk-based DBMSs.
- **Multi-Version Concurrency Control (MVCC)** – eliminates database locking to boost scalability and speed.
- **64-bit support** - eXtremeDB-64 is proven in very large deployments, including managing a 1.17 TB, 15.54 billion row database in memory, with results exceeding 87 million query transactions per second.

- **Hybrid storage** – *eXtremeDB* Fusion enables persistent (disk- or flash-based) storage for selected record types.
- The ***eXtremeDB* High Availability** edition contributes key pieces to *eXtremeDB* Cluster:
 - A two-phase commit protocol
 - Replication *protocol* and *interface* layers to insulate developers from underlying complexity

Developer Tools

eXtremeDB includes powerful development features to optimize applications for speed, footprint and other desired characteristics.

eXtremeDB's **native, project-specific API** for development in C/C++ ensures that each database operation in the API reflects the type of data being manipulated (type safety). Optional ***eXtremeSQL*** supports the widely used SQL standard as well as ODBC and JDBC.

McObject offers full source code, to give an in-depth understanding of *eXtremeDB* within an application.

eXtremeDB supports complex data types including structures, arrays, vectors and BLOBs.

eXtremeDB provides extremely efficient indexing for queries. Rather than storing duplicate data, indexes contain only a reference to data, keeping memory requirements to an absolute minimum. Supported indexes include:

- Hash indexes for exact match searches
- Tree indexes for pattern match, range retrieval and sorting
- R-tree indexes for geospatial searches
- KD-tree for spatial and Query-By-Example (QBE)
- Patricia trie indexes for networking & telecom
- Object-identifiers and references, for direct access
- Custom indexes

For application debugging, the *eXtremeDB* runtime includes **progressive error detection and consistency features**.

For development, the *eXtremeDB* runtime implements many verification traps and consistency checks. Then, when the application is debugged and consistently passes verification tests, developers can employ the optimized *eXtremeDB* runtime with fewer checks, to restore valuable clock cycles.

Additional *eXtremeDB* Features

eXtremeDB's many extras help developers and application end-users get the most from the database.

- **Small Footprint:** *eXtremeDB*'s small code size and efficient use of CPU cycles minimize hardware demands on cluster nodes, lowering cost-of-ownership.
- **Open Replication.** Data Relay technology facilitates selective data sharing between real-time systems based on *eXtremeDB*, and external systems such as enterprise DBMSs.
- **Custom Collations.** Specify the character sorting sequence (collation) for text, including collations supporting more than one language.
- **Event Notifications.** Notifies an application when something “of interest” in the database changes. Supports synchronous and asynchronous modes.
- **Security Features.** Page-level Cyclic Redundancy Check (CRC) detects unauthorized changes, while RC4 encryption blocks both tampering and unauthorized access.
- **Pattern Search.** Use wildcards to search tree index entries for single and multiple character matches.
- **Database Striping/Mirroring.** *eXtremeDB* Fusion can exploit multi-disk (solid state or spinning) configurations with its support for RAID-like data striping and data mirroring.
- **Cache Prioritization.** *eXtremeDB* Fusion enables applications to influence how long certain pages of persistent tables remain in cache, to accelerate retrieval in time-sensitive tasks.
- **Binary Schema Evolution.** *eXtremeDB* can save a database as a binary image and then restore it with a changed schema, for quick and efficient design changes.