

# eXtremeDB® for HPC

## Database System for Capital Markets Applications



“eXtremeDB enabled us to reduce latency to the sub-millisecond level per order while implementing a complex risk and compliance system.”

-- NSE.IT

eXtremeDB for HPC, the low latency database management system for capital markets.

### Financial Systems: The Database Challenge

McObject developed the eXtremeDB for HPC database system to break through the financial IT data management bottleneck, in applications such as algorithmic trading, risk management and order matching. The technology leverages proven eXtremeDB strengths – including an in-memory database system (IMDS) design, multi-core optimization, maximum developer flexibility and high scalability – and adds features to address key financial data management challenges. Specialized features include:

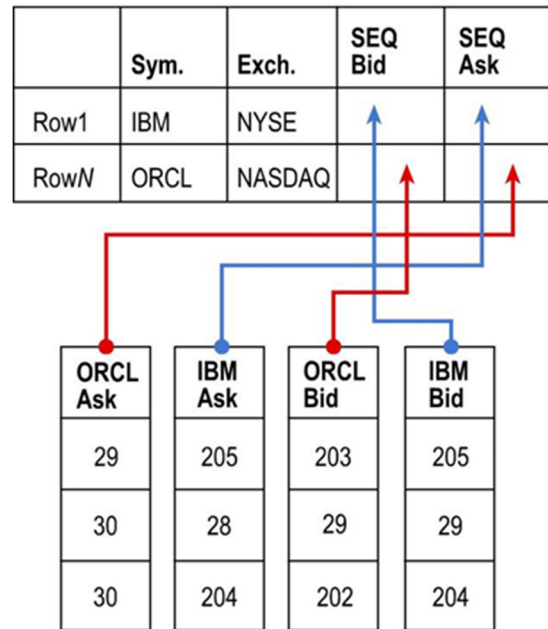
- Columnar data layout for fields of type ‘sequence’. Sequences are ideal for representing time series such as tick streams, historical quotes and other sequential data
- A library of **vector-based statistical functions** that can be pipelined to form an assembly line of operations on time series data for statistical/quantitative analysis
- A **GUI dashboard**, the xPanel, to configure, monitor, optimize databases, including transaction throughput, memory consumption and other key metrics

These combine with eXtremeDB’s wealth of core features, and many specialized options, to deliver the fastest database system with the flexibility and reliability demanded for financial applications.

### Design Goal: Minimizing Latency

eXtremeDB for HPC’s core in-memory database system (IMDS) eliminates I/O-based latency that is hard-wired into traditional disk-based database management systems (DBMSs). It can run entirely within the application process, eliminating inter-process communication (IPC) overhead, or as a highly scalable client/server architecture.

Keeping code and data in CPU (L1/L2) cache eliminates costly data transfers to/from main memory. Column-based data layout maximizes L1/L2 cache efficiency with market data, and small code size maximizes the likelihood that the entire code path for a given operation is loaded into the cache at once.



**Traditional DBMSs bring rows of data into L1/L2 cache for processing, but financial data – such as ticks, trades and quotes – are better handled by a column-based layout that maximizes efficiency in fetching needed information. The result is higher performance: the database system benefits from avoiding fetching unneeded data when it’s stored in a row-wise organization.**

eXtremeDB supports standard SQL/ODBC/JDBC and a native C/C++ API that can sometimes deliver even faster and more predictable performance (i.e., to minimize latency spikes).

### Record-setting STAC-M3 Results

How fast is eXtremeDB for HPC? In the STAC- M3™ – an independent, audited benchmark suite considered the gold standard for assessing time-series data management solutions (tick databases) – since 2013, eXtremeDB has consistently set the fastest mean response times ever reported for the STAC-M3™ benchmarks, utilizing lower cost hardware than the competition. Get details at [www.mcobject.com/comparison](http://www.mcobject.com/comparison).

## Powerful Horizontal Scalability

*eXtremeDB* delivers the power of distributed database processing via two sophisticated features:

**Sharding** is the horizontal distribution of data across multiple database instances that collectively represent a single logical database. These instances can be distributed across multiple drives connected to a single server to take advantage of multiple I/O channels and multiple CPUs/CPU cores, or across multiple servers. A sophisticated **distributed query engine** shields application logic from the need to know the awareness of the network topology.

**High Availability:** Master/slave replication, with automatic failover, to ensure continuous database operation even in the face of hardware or software failure. Replication is either synchronous (2-safe; strong consistency, lower replication speed) or asynchronous (1-safe; weaker consistency, faster replication). Replica databases are read-only, so query loads can be further distributed.

## Superior Vertical Scalability

*eXtremeDB* delivers vertical scalability to manage Big Data:


**64-bit support** enables in-memory databases scaling to terabyte-plus sizes.

**Hybrid storage** enables selective persistent (disk- or flash-based) storage. Persistent data size is limited only by space on the media. This is ideal for handling live quote data and historical data within a single database.

An optional **multi-version concurrency control (MVCC)** transaction manager eliminates “pessimistic” locking to accelerate multi-threaded applications running on multi-core hardware.

## Other Key Features

- Transactional (supports ACID properties)
- Implemented in C/C++
- Event notifications
- LUA stored procedures
- Security features: cyclic redundancy check (CRC) & AES encryption
- Source code and object code licenses available
- Querying methods: B-tree indexes, Hash, KD-tree, custom indexes and more
- Database striping/mirroring
- Successful track record (proven in automated trading, risk management, order matching, ticker plant, analytics, order execution and more)
- Multiple APIs: standard SQL/ODBC/JDBC & native C/C++, Java and C# (.NET)
- Native C/C++ API is type-safe. Data typing errors are caught during compilation



*eXtremeDB* outperforms competing database solutions in sorting, storing and retrieving market data. This lends a competitive advantage to capital markets IT.