

eXtremeDB™ Fusion

Hybrid in-memory/on-disk database system for maximum performance and data durability.

"eXtremeDB simplifies development and testing, especially in situations where the database must coordinate multiple processes."

-- Tyco Thermal Controls

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

Overview

In-memory database systems (IMDSs) offer superior performance and the possibility of very small RAM, CPU and storage demands. IMDSs boost speed by eliminating file system I/O, multiple data copies, and redundant processes, such as caching. This streamlined design can also dramatically reduce system footprint.

In contrast, on-disk databases cache frequently requested data in memory, for faster access, but write database inserts, updates and deletes through the cache to persistent storage. Byte-for-byte, disk storage can cost less than memory, and require less physical space: RAM chips can't yet approach the density of a micro-drive, for instance. So for small form-factor devices with large storage needs, such "spinning memory" can be better.

eXtremeDB Fusion provides the best of both worlds, marrying in-memory database technology with the traditional disk-based database system. The result is a hybrid database for resource-constrained and high performance systems that affords developers the ultimate in flexibility.

McObject's eXtremeDB

Since its introduction, McObject's eXtremeDB has set the standard for small footprint, in-memory embedded database systems, offering benefits including:

- **Tiny code size** of approximately 100K or less
- Blazing speed: **micro-second transactions** even on modest hardware
- C/C++ developers benefit from a **type-safe, intuitive API** with extensive checking to speed development
- Optional **SQL** and **XML** interfaces
- **Java Native Interface (JNI)** affords Java developers the ease of working with "plain old Java objects" (POJOs)
- **High Availability Edition**, with asynchronous (1-safe) or synchronous (2-safe) replication, for applications requiring complete fault tolerance
- **Available source code**, for porting to new platforms and highest degree of control over development
- **64-bit edition** scales beyond 1TB in-memory data
- **Multi-version concurrency control (MVCC)** transaction manager and advanced memory management fully leverage multi-threaded, multi-core systems

eXtremeDB Fusion: Best of Both Worlds

eXtremeDB Fusion enables the developer to combine in-memory *and* on-disk paradigms in a single database system. Specifying that data will be stored in memory (transient), or on disk (persistent), requires a simple database schema declaration, as shown below.

```
transient class classname {
    [fields]
};

persistent class classname {
    [fields]
};
```

The resulting system retains in-memory strengths (speed, footprint, etc.), yet leverages the potential cost savings and durability of an on-disk database.

Key On-Disk Database Features

eXtremeDB Fusion's on-disk features are uniquely configurable, including:

- Three transaction logging policies – Undo, Redo and No Logging – to meet the target system's footprint, performance and durability needs
- Synchronous or asynchronous transaction logging
- Developers can specify the maximum database size, which is especially important when the 'disk' is actually a flash memory file system
- Database cache can be saved and re-used across sessions – for example, so a user can resume some activity when a device is switched back on
- The database can exist in one file, to simplify maintenance, limit I/O and reduce size
- Logical Database Devices feature can spread a database across multiple disks, including in a RAID, with the database striped across RAID disks
- Or, pages can be written simultaneously to multiple RAID disks for perpetual backup

With these tools, the developer fine-tunes the database according to the speed, footprint and other requirements of the target system. eXtremeDB Fusion puts the developer in charge.

Highly efficient indexing

For transient classes, rather than storing duplicate data, *eXtremeDB* Fusion's diverse indexes contain only a reference to data, minimizing memory requirements. 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 network, telecom
- Object-identifier references, for direct access
- Custom indexes

Additional Features

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

- **HTML database browser/editor.** Retrieve database and class statistics and lists of classes; generate schema in the form of a Data Definition Language (DDL) file
- **XML Extensions.** Generates interfaces to create or update an object in the database from the content of an XML document, export an object as an XML document, and to generate an XML schema
- **Remote procedure call mechanism (MCORPC).** Framework enables remote processes to read/update an *eXtremeDB* in-memory or persistent database
- **Database calculator.** Collect information needed to choose ideal page size and to optimize schema designs, storage layout and performance
- **Pattern search.** Use wildcards to search tree index entries for single and multiple character matches:

Supported Platforms

Embedded Platforms:

- VxWorks 5.5, 6.x
- VxWorks 653 RTOS (for avionics)
- INTEGRITY OS
- QNX 6.x
- Various Real-Time Linux distributions
- Lynx OS
- RTXC Quadros, RTXC 3.2
- Microsoft Windows Embedded
- eCos
- Nucleus
- Bare bones boards (no operating system required)

Development environments

- gnu toolchain (gcc 2.95 and higher)
- Tornado 2.0 and 2.2 (GNU and Diab compilers)
- QNX Momentics IDE (C, C++, Embedded C++)
- Metrowerks CodeWarrior IDE (various platforms)
- GreenHills Multi
- Microsoft Visual Studio (C/C++, .NET)

Server and Desktop Platforms:

- Sun Solaris 8, 9 and 10
- HP-UX 11.x
- Linux distributions
- Classic Windows platforms (98/NT/2000/XP/Vista)

Database Specifications

Maximum in memory database and/or cache size

32-bit: 3 gigabytes

64-bit: 18 exabytes

Maximum database size: file system limit

Maximum classes per database: 32,767

Maximum indexes per database: 32,767

Maximum fields per class: 32,767

Maximum fields per index: 32,767

Maximum elements per vector: 32,767

Memory requirements: 100K or less

Maximum database connections: configurable

Maximum open databases: configurable

Supported Data types

- 1, 2, 4, 8-byte signed/unsigned integers
- float, double
- date, time
- char (fixed length)
- string (variable length)
- rect(angle)
- Unicode
- boolean (array of bits)
- enum
- fixed-size array
- variable-length vector
- structs (embedded to any depth)
- autoid (auto-increment)
- user-defined object-id and references