

eXtremeDB™

High Availability Edition

Fault-tolerant in-memory database for embedded systems that cannot afford to fail.

“eXtremeDB High Availability will provide the highest degree of reliability for the Apache’s mission-critical systems.”

-- The Boeing Company

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

Overview

eXtremeDB High Availability (HA) is a fault-tolerant version of the eXtremeDB embedded database.

Designed to power embedded systems that cannot afford to fail, eXtremeDB-HA delivers the highest degree of reliability, along with unsurpassed performance and an ultra-small footprint.

Many fault-tolerant systems use “lazy” replication schemes that write new information to backup database instances asynchronously (after changes to the master database commit). This approach entails greater latency, slower failover, and the risk of lost transactions, all of which are unacceptable in many time-critical embedded systems. In contrast, the eXtremeDB-HA runtime uses *time-cognizant eager replication* in which related updates occur within a single transaction; strict processing deadlines ensure on-time delivery of data from master to replica sites, and failover procedures are extremely short. Typical hardware configurations include:

- Multiple processes or threads within the same hardware instance
- Two or more boards in a chassis with a high-speed bus for communication
- Separate address spaces on boards connected via industry standard communication media and protocols such as RS-485/RS-232, Control Area Network (CAN) or Ethernet
- Two or more controllers connected via proprietary communication media and/or protocols
- Separate computers on a LAN

McObject’s eXtremeDB-HA eager replication is based on a rugged, time-cognizant two-phase commit protocol that ensures changes to the master database and identical replica databases succeed or fail together. A high availability control interface exported by the eXtremeDB-HA runtime provides the means for the application to configure, establish, maintain and terminate eXtremeDB-HA connections. The time-cognizant HA transport protocol enables communication as well as detection of timeout situations.

High Availability Application Framework

The High Availability Application Framework source code included with eXtremeDB-HA provides an abstraction mechanism to allow master and replica applications to exchange data and messages via user-defined communication channels during transaction processing. By implementing communication channels outside the eXtremeDB core library, the HA Application Framework stays independent of application and eXtremeDB database code, providing flexibility to test and switch between communication channels at compile-time.

To facilitate faster development and deployment of fault-tolerant database solutions, the High Availability Application Framework provides re-usable code that demonstrates how to incorporate the database into applications. This includes working examples of communication channels built over various transports (currently TCP/IP, UDP/IP, Named Pipes and Qnet™) or a user-defined channel. In addition, the framework offers a HA database-enabled application prototype that can be configured to use any of the channels. The prototype also enables developers to see eXtremeDB-HA in action immediately after installation.

Key Features

- Survives software and hardware failure without losing data or ceasing operation
- Rugged time-cognizant two-phase commit protocol
- Easy-to-use C and C++ APIs
- Communication media and protocol independent
- Support for hot standby
- Replicas are always synchronized with the master
- Load balancing via distributing queries to replicas
- Support for multiple replicas; new replicas can join at any time
- Prototype demonstrates eXtremeDB-HA immediately after installation