

HA-JDBC: High-Availability JDBC

Table of contents

1 Overview.....	2
2 Features.....	2
3 Dependent Libraries.....	2
4 Related Software.....	2

1. Overview

HA-JDBC is a JDBC proxy that provides light-weight, transparent, fault tolerant clustering capability to any underlying JDBC driver.

2. Features

- Supports any database accessible via JDBC.
- High-availability/Fault Tolerance - An HA-JDBC database cluster can lose a node without failing/corrupting open transactions.
- Live activation/deactivation allows for maintenance/upgrading of a database node without loss of service.
- Improves performance of concurrent read-access by distributing load across individual nodes.
- Supports full JDBC 3.0 and 4.0 feature set.
- Out-of-the-box database-independent strategies for synchronizing a failed cluster node.
- Exposes JMX management interface to allow administration of databases and clusters.
- Ability to add/subtract database nodes to/from a cluster at runtime.
- Can be configured to auto-activate failed database nodes during scheduled off-peak times.
- Open source (LGPL).

3. Dependent Libraries

- [JGroups](#) - reliable multicast communication framework
- [Quartz](#) - enterprise job scheduler
- [JiBX](#) - an XML binding framework
- [SLF4J](#) - a simple facade for various logging APIs.

4. Related Software

- [Sequoia](#) - Database-independent clustering middleware
- [PGCluster](#) - Multi-master replication system for PostgreSQL.
- [MySQL Cluster](#) - Fault tolerant database architecture using NDB storage engine
- [H2 Clustering](#) - a simple clustering / high availability mechanism for H2