

Todo List

Table of contents

1 high.....	2
2 medium.....	2
3 low.....	2

1. high

- **[src]** Add diff synchronization strategy alternative that leverages a last modified timestamp column. This will minimize data needed to be queried, as well as speed up row comparisons. # PF
- **[src]** Create a feasible synchronization strategy for large databases. I envision a buddy schema that intelligently records inserted/modified/deleted primary key values per table. Synchronization will only need to insert/update/delete the recorded rows. # PF
- **[src]** Add optional "protect-commit" mode that detects and recovers from partial commits incurred as a result of a node crash. # PF

2. medium

- **[src]** Create distributable balancer implementation for round-robin and load. Allow distributable configurations to opt for local balancer implementations. # PF
- **[src]** Add more robust handling of jgroups partitions. # PF
- **[]** Create integration test suite using H2 or some full-featured, highly standard compliant embedded database. # PF
- **[doc]** Update performance graphs with results from latest version. # PF

3. low

- **[src]** Support for using HA-JDBC in conjunction with native database replication capabilities. HA-JDBC would add high-availability to cluster through multiple master support while hiding write-to-master/read-from-slave details from the application. # PF
- **[src]** Investigate the potential of SQLState codes as a replacement for secondary isAlive() queries when determining how to handle failures. Also consider using Connection.isValid(int) with Retrotranslator plugin implementation to support Java 1.4 + 1.5. # PF