



# Firebird in 2011/2012: Development Review

Dmitry Yemanov  
<mailto:dimitr@firebirdsql.org>

Firebird Project  
<http://www.firebirdsql.org/>

# Packages Released in 2011

## **Firebird 2.1.4**

March 2011

96 bugs fixed

4 improvements, new port (Renesas SH)

## **Firebird 2.5.1**

October 2011

140 bugs fixed

16 improvements, 1 new feature

New ports: Linux/M68K, Linux/HPPA, Linux/Alpha

# What's New in v2.5.1

## Key enhancements

SQLSTATE as a PSQL context variable  
(like GDSCODE/SQLCODE)

Record compression ratio estimations

Write operations for GTTs in RO transactions

Performance improvements for GTTs

GBAK restore performance

Canceling waits inside the lock manager

Persistent statement IDs in the MON\$ tables

# Maintenance Releases in 2012

## Firebird 2.0.7

24 bugfixes up to date

4 months overdue, expected in Q1 2012

## Firebird 2.1.5

32 bugfixes and 2 improvements up to date

No delay, expected in Q2-Q3 2012

## Firebird 2.5.2

13 bugfixes up to date

No delay, expected in Q2-Q3 2012

# Firebird 3.0 : Introduction

## Development summary

Started in October 2009

433 total tickets assigned

379 of them are already resolved

More than 2000 changes committed

Code cleanup, only ODS 12 is supported

New layout of binaries

# Firebird 3.0 : New Features

## Implemented SQL features

External procedures / functions / triggers

PSQL functions

PSQL packages

Sub-procedures and sub-functions

DDL triggers

Window (analytic) functions

Identity columns

Scrollable PSQL cursors

## Firebird 3.0 : New Features

### Implemented SQL features (continued)

Standard compliant BOOLEAN data type

ALTER statements for changing NULL / NOT NULL

ALTER DATABASE for changing the default charset

MERGE enhancements (multiple MATCHED clauses, RETURNING clause, DELETE command)

CONTINUE statement

Parametrized exceptions

# Firebird 3.0 : New Features

## Security

Pluggable authentication

Per database user management

Over-the-wire encryption

Database encryption

DDL permissions

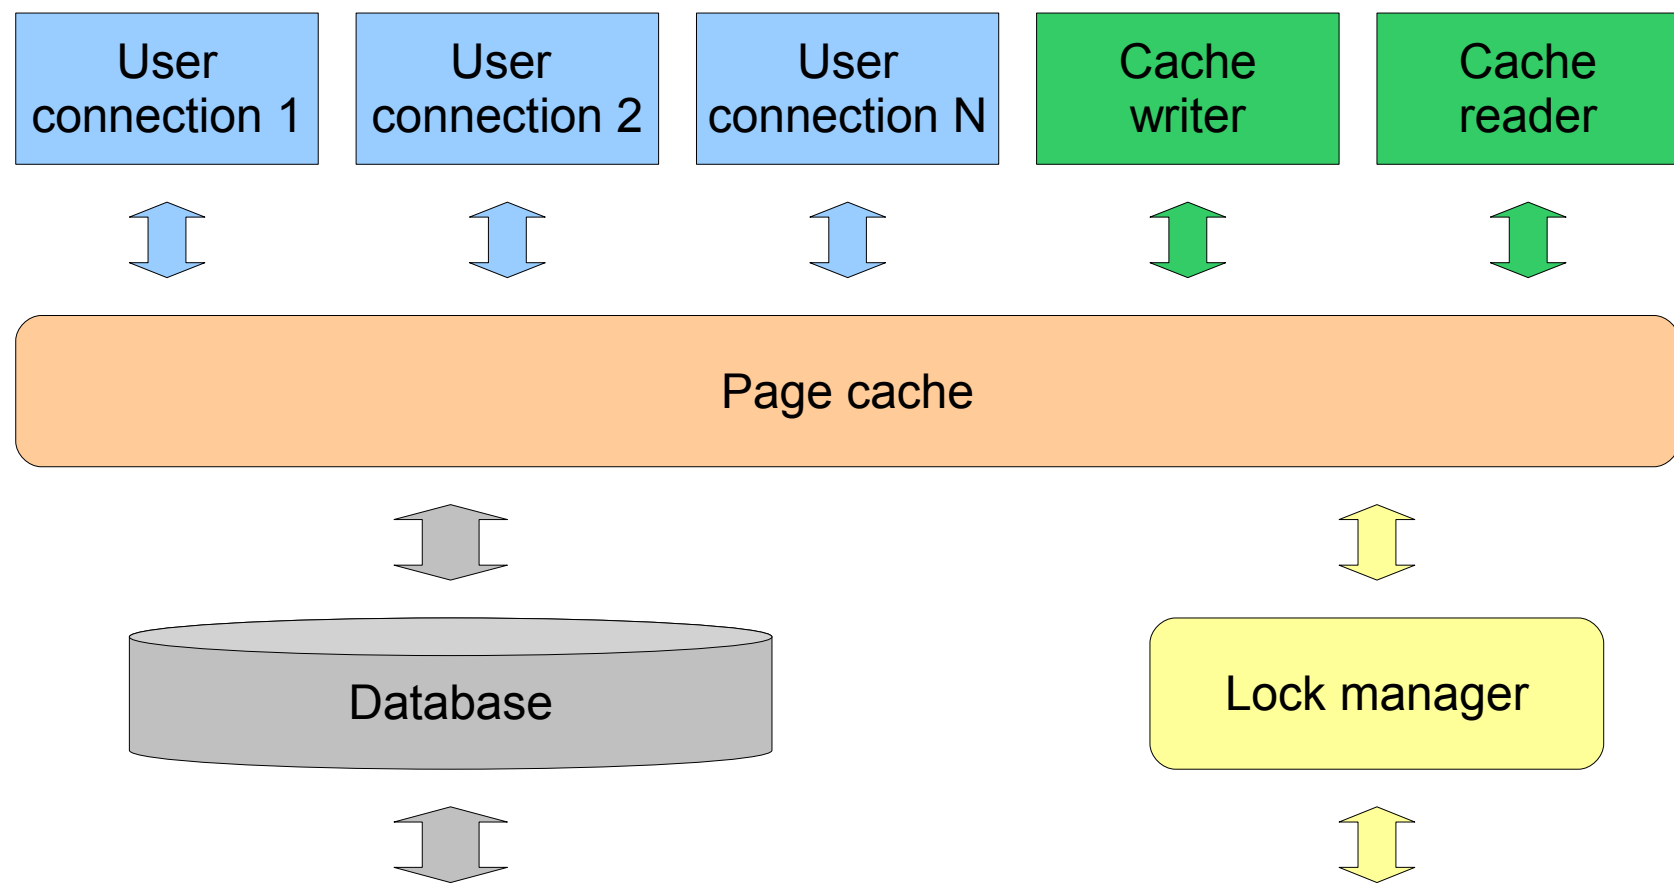
Aggregated privileges for roles

User groups (implicit roles)



# Firebird 3.0 : Shared Cache

## Basic architecture



# Firebird 3.0 : Shared Cache

## Configuring the architecture

SharedCache – whether multiple connections of a single process can share the database page cache

SharedDatabase – whether database pages can be shared among multiple page caches

## Matrix of working modes

	SharedDatabase = 0	SharedDatabase = 1
SharedCache = 0	Singleton	Classic / SuperClassic
SharedCache = 1	SuperServer	Special (*)

# Firebird 3.0 : Shared Cache

## Synchronization

Latches and locks

Database file open mode (SharedDatabase)

Cache start up (no locks)






New connection (SharedCache)

Acquire "lazy" locks

Two level synchronization


# Firebird 3.0 : Shared Cache

## Concurrency test: 1GB database, 10 terminals



# Firebird 3.0

## MySQL: single thread performance



## Firebird 3.0 : ODS 12

### “Primary” and “secondary” data pages

New flag for both pointer pages and data pages

Primary page contains some “head” version of the record version chain

Secondary page contains only back versions, fragments or small blobs inside

Full table scans (plan NATURAL) skips secondary data pages

## Firebird 3.0 : ODS 12

### “Swept” data pages

New flag for both pointer pages and data pages

Swept data page means that all records stored there have no back versions and they are known to be created by already committed transactions

Set by the sweeper / garbage collector, cleared when new version is created on the page

Sweep process skips swept data pages

## Firebird 3.0 : ODS 12

### SCN inventory page

New page type

Global array of SCNs for the every database page

SCN inventory pages are distributed evenly across the database file

Slots are updated when some page gets a new SCN assigned

Nothing is changed for a level 0 backup

Incremental (level N) backups read only pages modified since the last (level N – 1) backup



# Firebird 3.0 : Optimizer

## Detailed plan output

An alternative (optional) way to analyze the plan

All operations are reported in the hierarchical view

Some additional details are available

Cardinality and cost estimations will be reported

Run-time information

## Firebird 3.0 : Optimizer

### Materialized sub-queries

Underlying stream is sequentially read and cached inside a record buffer

Pipeline nature of the record buffer, both sequential and random access are supported

Storage is provided by the temporary space manager and is dynamically balanced between memory and disk (see TempSpaceLimit in firebird.conf)

Useful for invariant sub-queries, hash joins, etc

# Firebird 3.0 : Optimizer

## Hash joins

Larger input stream becomes an outer one, smaller input stream becomes an inner one

Inner stream is read in advance row by row, hash value is calculated for the join keys, row is stored in the buffer corresponding to the given hash group

Outer stream is read sequentially, hash value is calculated for the join keys and probed against the hash table

For a positive match, collisions are compared by the binary comparison of the join keys

# Firebird 3.0 : Optimizer

## Miscellaneous

Semi-materialized nested loop joins

Reworked FULL OUTER JOIN implementation that is able to utilize indices

Support for hash and merge algorithms for LEFT OUTER joins

Conditional retrievals, e.g. WHERE (COL = :VAL OR :VAL IS NULL)

Constant predicates can be evaluated beforehand

Hashing is considered for aggregates

# Firebird 3.0 : Optimizer

## Table statistics

Number of data pages, average fill ratio

Number of rows, average row size

## Index statistics

Tree depth

Number of leaf pages, average fill ratio

Number of nodes, average node size

Clustering factor

# Firebird 3.0 : Optimizer

## Column statistics

Number of distinct values

Number of NULLs

## Histogram

Endpoint boundary

Endpoint value

# Firebird 3.0 : Monitoring

## General

More details about page I/O

Temporary space I/O statistics

Statement statistics: parses / compiles / executes

Misc extensions

## Timing statistics

Total elapsed time

CPU time: user mode and kernel (system) mode

Waiting time: I/O, lock waits, latch waits, etc

## Firebird 3.0 : API

### Extending the length limits

API: SQL statements and their BLR counterparts

API: input/output messages

Internal: intermediate (non-persistent) records

Internal: sorting records

Not about the ODS records!





**Questions?**