



Firebird 2.5 Release Notes

Helen Borrie (Collator/Editor)

30 September 2011 - Document v.0251_08 - for Firebird 2.5.1 Sub-release

Firebird 2.5 Release Notes

30 September 2011 - Document v.0251_08 - for Firebird 2.5.1 Sub-release
Helen Borrie (Collator/Editor)

Table of Contents

1. General Notes	1
Bug Reporting	1
Documentation	2
2. New in Firebird 2.5	3
Firebird 2.5.1 Release (September 2011)	3
Firebird 2.5 Release (October 2010)	4
Other New Features	4
Administrative Enhancements	4
Other SQL Language Additions and Enhancements	4
Data-handling Enhancements	5
API Additions	5
International Language Support	5
3. Changes in the Firebird Engine	6
New Threading Architecture	6
“Superclassic”	7
Thread-safe Client Library	8
Improvements	8
Immediate Detection of Disconnected Clients on Classic	8
Optimizations	9
Cache Size Limit Increased for 64-bit Servers	9
Default Database Location	10
DLL Loading for Windows Embedded Engine	10
Large External Table Support Enabled	10
Statistics Now Work Properly with 64-bit Values	11
UDFs Safeguard	11
Diagnostics	11
Metadata Improvements	12
4. Changes to the Firebird API and ODS	13
ODS (On-Disk Structure) Changes	13
New ODS Number	13
Maximum Page Size	13
Maximum Number of Page Buffers in Cache	13
API (Application Programming Interface) Extensions	13
BLOB Conversions Enabled in the API Functions	13
Connection Strings & Character Sets	14
Support for SQLSTATE Completion Codes	15
“Efficient Unprepare”	16
Cancel Operation Function	16
Shutdown Functions	18
Tighter Control Over Header-level Changes	22
New Trace Services for Applications	23
Other Services API Additions	25
New Trace API	28
5. Reserved Words and Changes	29
Clean-up of Reserved Words	29
Newly Reserved Words	29
Keywords Added as Non-reserved	29
6. Configuration Parameter Additions and Changes	30

AuditTraceConfigFile	30
Parameters Affecting Filesystem Cache Usage	30
FileSystemCacheSize	30
FileSystemCacheThreshold	31
MaxFileSystemCache	31
ConnectionTimeout	32
Authentication	32
Changes in V.2.5	32
MaxUserTraceLogSize	32
OldSetClauseSemantics	33
RemoteAuxPort For Classic and Superclassic	33
Use Hostname for RemoteBindAddress	33
RemoteFileOpenAbility	34
7. Administrative Features	35
New RDB\$ADMIN System Role	35
Multiple Databases and Superusers	35
System “Superusers”	36
Escalating RDB\$ADMIN Scope for User Management	37
Trace and Audit Services	38
Overview of Features	38
The System Audit Session	38
User Trace Sessions	39
Trace Scope on Windows	42
Use Cases	42
Trace Plug-in Facilities	42
Monitoring Improvements	43
Extended Access for Ordinary Users	43
New MON\$ Metadata for ODS 11.2 Databases	43
Usage Notes	44
8. Security Hardening	46
Windows Platforms	46
No SYSDBA Auto-mapping (Windows)	46
9. Data Definition Language (DDL)	47
Quick Links	47
Visibility of Procedure Definition Changes on Classic	47
CREATE/ALTER/DROP USER	48
Syntaxes for Altering Views	49
Extensions for CREATE VIEW	50
ALTER Mechanism for Computed Columns	51
Extensions for SQL Permissions	51
Default COLLATION Attribute for a Database	53
ALTER CHARACTER SET Command	54
Evolution of CREATE DATABASE	55
10. Data Manipulation Language (DML)	58
Quick Links	58
RegEx Search Support using SIMILAR TO	58
Hex Literal Support	63
New UUID Conversion Functions	64
SOME_COL = ? OR ? IS NULL Predication	65
Extension to LIST() Function	66
Extension to DATEADD and DATEDIFF() Functions	66
BIN_NOT() Function Added	66

Write to Temporary Tables in a Read-Only Database	66
Optimizer Improvements	67
Other Improvements	67
11. Procedural SQL (PSQL)	69
Quick Links	69
Autonomous Transactions	69
Borrow Database Column Type for a PSQL Variable	70
New Extensions to EXECUTE STATEMENT	71
Context Issues	72
External Queries from PSQL	73
EXECUTE STATEMENT with Dynamic Parameters	74
Exception Handling	76
Examples Using EXECUTE STATEMENT	77
Other PSQL Improvements	80
Subqueries as PSQL Expressions	80
SQLSTATE as Context Variable	80
12. International Language Support (INTL)	81
Default COLLATION Attribute for a Database	81
ALTER CHARACTER SET Command	81
Connection Strings & Character Sets	81
Other Improvements	81
Introducer Syntax Usage	81
Malformed UNICODE_FSS Characters Disallowed	82
Repair Switches for Malformed Strings	82
Numeric Sort Attributes	82
Character Sets and Collations	83
13. Command-line Utilities	84
fbtracemgr	84
Action Switches	84
Parameters	84
Examples using fbtracemgr	85
V.2.5.1 Improvements	85
Retrieve Password from a File or Prompt	85
New -fetch_password Switch	85
gsec	86
Mapping Switch for Windows Administrators	86
Command-line Help for gsec	87
fbsvcmgr	87
gbak	88
Repair Switches for Malformed Strings	88
Preserve Character Set Default Collation	88
Improve Insertion Performance for Restore	88
nBackup	88
New Switches Added	88
I/O Resource Load Reduced on POSIX	89
isql	90
SQLSTATE instead of SQLCODE	90
Improvement for Exponential Number Output	90
SHOW COLLATIONS in isql Help	90
SET ROWCOUNT Statement Added	90
gpre (Precompiler)	91
Some Updates	91

gstat	91
14. Installation Notes	92
Linux (POSIX)	92
Installation Scripts Cleanup	92
Dedicated Firebird Switches for 'configure'	92
Detection of Path to Firebird's Binaries	93
Windows	93
Deployment Structure for Embedded	93
Managing MSCV8 Assemblies	94
15. Compatibility Issues	96
Effects of Unicode Metadata	96
Configuration Parameters Removed	96
SQL Language Changes	96
Reserved Words	97
Execution Results	97
Utilities	97
fb_lock_print	97
API Changes	98
Rejection of Inconsistent TPB Options	98
Addition of SQL_NULL Constant	98
Security Hardening	98
No SYSDBA Auto-mapping (Windows)	99
Default Authentication Method (Windows)	99
Access Path for Services	99
Known Platform Issues	99
MacOSX	100
16. Platform Ports	101
HPPA	101
Linux/HPPA	101
Linux/Alpha	101
IBM eServer z-Series	101
Linux/s390 (32-bit)	101
Linux/s390x (64-bit)	102
Linux/sh4 (Renesas SH)	102
HP-UX	102
Lock Table Improvement for HP-UX	102
Port for Very Old Windows 32-bit Platforms	102
17. Bugs Fixed	104
Firebird 2.5.1 Release	104
Core Engine/API	104
Server Crashes	111
Data Manipulation Language	113
Command-line Utilities	115
Database Monitoring/Administration	117
Services Manager	117
POSIX-Only Bugs	117
Windows-Only Bugs	119
Remote Interface/API	119
Firebird 2.5.0 Release	121
Core Engine/API	121
Server Crashes	122
Command-line Utilities	122

Old Bugs Fixed	122
Core Engine/API	122
Server/Client Crashes	141
Database Monitoring/Administration	143
Data Manipulation Language	144
Command-line Utilities	145
Services Manager	150
Remote Interface/API	150
Security	152
International Language Support	153
POSIX-specific	155
Windows-specific	157
MacOSX-specific	158
Miscellaneous Bugs	158
18. Firebird 2.5 Project Teams	159
Appendix A: SQLSTATE	161
SQLSTATE Codes & Messages	161
Appendix B: Licence Notice	169

List of Tables

10.1. Character class identifiers	60
18.1. Firebird Development Teams	159

Chapter 1

General Notes

As well as a substantial number of bug fixes, this sub-release takes in a few minor improvements and optimizations that missed the initial release. In summary:

- Significant for MacOSX 10.7 users was a bug that caused attempts to start Superserver and Superclassic to fail. The bug is described in [Tracker ticket CORE-3589](#) and also in the [Bug Fixes list](#) for this sub-release.
- An SQLSTATE context variable is now available in PSQL, parallel to the the context variables GDSCODE and SQLCODE that are used in WHEN blocks to test for error conditions.
- A couple of improvements were made, that should help performance when global temporary tables (GTTs) are being used:
 - The “undoing”, on rolling back the transaction, of changes made to GTTs that were created with the ON COMMIT DELETE ROWS option was an unnecessary overhead and has been bypassed.
 - Garbage collection in global temporary tables was being delayed unnecessarily by active transactions in other attachments. That bottleneck is gone.
- A needed optimization in the temporary space manager with regard to small chunk allocations has been implemented.
- The Lock Manager has been provided with the capability to cancel waiting, avoiding a condition whereby a transaction in WAIT mode could wait interminably for the end of another transaction that could not be achieved by either a **DELETE FROM MON\$xxx** or an **fb_cancel_operation** request.
- The query optimizer now estimates the actual record compression ratio, enabling better guesses about the number of stored records in tables.
- Some minor improvements to the remote interface:
 - Any unused bytes of varchar values in the message buffer are now set to zero.
 - Set the SO_KEEPALIVE option on the client TCP socket
- The MON\$STATEMENT_ID value constant is now kept constant among monitoring snapshots.
- Ports of this sub-release were done for Linux/HPPA and Linux/alpha platforms.

Bug Reporting

- If you think you have discovered a new bug in this release, please make a point of reading the instructions for bug reporting in the article [How to Report Bugs Effectively](#), at the Firebird Project website.
- If you think a bug fix hasn't worked, or has caused a regression, please locate the original bug report in the Tracker, reopen it if necessary, and follow the instructions below.

Follow these guidelines as you attempt to analyse your bug:

1. Write detailed bug reports, supplying the exact build number of your Firebird kit. Also provide details of the OS platform. Include reproducible test data in your report and post it to our [Tracker](#).
2. You are warmly encouraged to make yourself known as a field-tester of this pre-release by subscribing to the [field-testers' list](#) and posting the best possible bug description you can.
3. If you want to start a discussion thread about a bug or an implementation, please do so by subscribing to the [firebird-devel list](#). In that forum you might also see feedback about any tracker ticket you post regarding this alpha.

Documentation

You will find all of the README documents referred to in these notes—as well as many others not referred to—in the doc sub-directory of your Firebird 2.5 installation.

An automated "Release Notes" page in the Tracker provides lists and links for all of the Tracker tickets associated with this and other pre-release versions. [Use this link](#).

--The Firebird Project

Chapter 2

New in Firebird 2.5

The primary goal for Firebird 2.5 was to establish the basics for a new threading architecture that is almost entirely common to the Superserver, Classic and Embedded models, taking in lower level synchronization and thread safety generally.

Firebird 2.5.1 Release (September 2011)

The months that have passed since the initial release have seen numerous Tracker issues tackled to address things that have been reported broken, bent or compromised in v.2.5. As well as a long [list of bug fixes](#), a few minor improvements have been made. In summary:

- The SQLSTATE code has been made available as a PSQL context variable, for use in WHEN .. exception blocks, in the same manner as GDSCODE and SQLCODE
- Now it is possible to write to global temporary tables in a read-only database
- Diagnostics for internal trace errors were improved
- The *fbtracemgr* utility will now do a periodic flush to output
- Performance of *gbak restore* at the data insertion stages has improved
- Conversions between BLOBs and other data types can now be effected in the API functions
- The Services API now supports the “metadata-only” restore
- A “silent install” switch has been implemented for **make install** on POSIX
- The unused bytes of VARCHAR values in the message buffer are now set to zero
- The actual record compression ratio is now estimated in the optimizer
- The MON\$STATEMENT_ID value now stays constant among monitoring snapshots
- The SO_KEEPALIVE option on the client TCP socket will now be set, as a measure to guard against aggressive timing out of sockets by newer Windows systems
- Lock Manager can now cancel waits that become interminable
- A platform port of v.2.5.1 for HPPA has been done for both Linux and Alpha

Firebird 2.5 Release (October 2010)

Although SQL enhancements are not a primary objective of this release, for the first time, user management becomes accessible through SQL CREATE/ALTER/DROP USER statements and syntaxes for ALTER VIEW and CREATE OR ALTER VIEW are implemented. PSQL improvements include the introduction of autonomous transactions and ability to query another database via EXECUTE STATEMENT.

Other New Features

Other new features and improvements in this release include:

Administrative Enhancements

- System audit tracing and user trace sessions via the Services API, making it possible to monitor and analyse everything going on in a database in real time
- New system role RDB\$ADMIN in the ODS 11.2 database allows SYSDBA to transfer its privileges to another user on a per-database basis
- More information in the monitoring tables
- Asynchronous cancellation of connections
- Capability for ordinary users to monitor any of their own attachments as well as CURRENT_CONNECTION

Other SQL Language Additions and Enhancements

- Regular expression support using the SIMILAR TO predicate
- ALTER COLUMN for computed columns
- Autonomous transactions within a PSQL module (stored procedure, trigger or dynamically executable PSQL block)
- Enhanced access to stored procedures in view definitions
- Optional GRANTED BY (or, alternatively, AS) for GRANT and REVOKE statements, enabling the grantor to be a user other than the CURRENT_USER (the default).
- REVOKE ALL syntax to dispose of all privileges for a user or role at once
- Support for WHERE SOME_COL = ? OR ? IS NULL predications
- Removal of “reserved” status for all but a handful of keywords that are not reserved in the SQL standard

Data-handling Enhancements

- New built-in functions for converting UUID CHAR(16) OCTETS strings to RFC4122-compliant format and vice versa
- Ability to pass 32-bit and 64-bit integers as hexadecimal in numeric literal and X-prefixed binary string literal formats

API Additions

- Statements now return an SQL-2003 standard 5-alphanumeric SQLSTATE completion code

Tip

In the v.2.5.1 sub-release, SQLSTATE was added to PSQL as a context variable for use in WHEN .. exception blocks, in a manner similar to SQLCODE and GDSCODE.

- New constant DSQL_unprepare available for use with isc_dsql_free_statement for efficient unpreparing of statements

International Language Support

- Default COLLATE clause for CREATE DATABASE
- Ability to change the default COLLATE for a used character set
- GBAK restore switches FIX_FSS_DATA and FIX_FSS_METADATA to restore legacy databases with UNICODE_FSS data and/or metadata correctly without resorting to scripts and manual techniques
- Accent-insensitive collation for Unicode

Chapter 3

Changes in the Firebird Engine

The primary objective of this release was to refactor Firebird's threading architecture to take advantage of the symmetric multiprocessing (SMP) capabilities of multiprocessor hardware. This has a noticeable effect on the scalability of Superserver when multiple databases are being accessed simultaneously but its major effect is the emergence of the architectural “Superclassic” model that will underpin the fine-grained multi-threading under development for Firebird 3.

New Threading Architecture

Dmitry Yemanov
Vladyslav Khorsun
Alex Peshkov
also -
Nickolay Samofatov
Roman Simakov

For Superserver, the new architecture will be most obvious in two ways:

1. In a multiple database environment, Superserver threads for each database are allotted evenly to available processors.

Note

The default *CpuAffinity* setting still binds SuperServer to a single processor only. In order to take advantage of this improvement when working with multiple databases, this setting should be changed in `firebird.conf`.

2. A slight improvement in scaling should be apparent for single database usage on SMP hardware

It is with Classic that the effects are most evident:

1. Classic Server can now be multi-threaded. The one worker thread per process model remains but now it is possible to use additional threads for parallel tasks such as asynchronous shutdown, sweep, inter-process communications with the lock manager and more.
2. On POSIX, services in Classic also run in threads now, rather than in forked processes as previously.

Note

For Windows Classic, services became threadable in v.2.1.

3. The embedded libraries—*libfbembed.so* on POSIX and *fbembed.dll* on Windows—are now multi-thread-capable and thread-safe, so they can be used in multi-threaded applications.
4. Testing suggests that the performance of Classic in this version will be significantly faster than previous Classic versions.

“Superclassic”

This multi-threaded mode for Classic has been dubbed “Superclassic” for its capability to handle multiple worker threads—dedicated or pooled—inside a single server process. It shares all the usual Classic features, with a few differences:

- Safe, full shutdown of the server engine is possible on any platform
- Under some TPC conditions, it can outperform Classic—by about 15-20%
- It uses fewer kernel resources (although not less memory)
- When a Superclassic process crashes, it takes all its connections with it
- Recognised limitations in the Services API for the Classic server, such as the inability to retrieve the list of attachments/active users, do not apply to SuperClassic.
- On POSIX, Superclassic does not require *[x]inetd*.

Embedded Server Notes

- The embedded server in the Windows library, *fbembed.dll*, now uses Superclassic, not Superserver as previously, thus unifying its model with that of local connection to Superclassic on POSIX. The database file-lock that previously restricted connections to a single application space is replaced by a global lock table that allows simultaneous access to the same database from different embedded server modules. This facilitates concurrent debugging of applications and use of native utility tools like *gbak*, *gstat* and so on.
- A single attachment handle can now be shared by simultaneous threads. (Tracker reference [CORE-2498](#), A. dos Santos Fernandes).

Usage Notes

Windows

On Windows, the same `fb_inet_server.exe` binary delivers either the Classic or the Superclassic working modes, according to switch settings. Classic is the default mode.

To use the Superclassic mode as a service, add the `-m[ulti-threaded]` switch to the `instsvc.exe` command line, as follows:


```
instsvc install -multithreaded
```

When intending to run Superclassic *as an application*, use

```
fb_inet_server -a -m
```

New Binary for POSIX

On POSIX, the new binary `fb_smp_server` is supplied for the Superclassic model. It contains the network listener, meaning it works similarly to `fbserver` with regard to attachment requests and does not require `[x]inetd`.

The multi-threaded engine used by `fb_smp_server` is the `libfbembed.so` library, in accordance with OS-RI requirements. The Classic packages also include `fbguard` (the Guardian) which, in a Superclassic installation, starts `fb_smp_server`, rather than `fbserver` as it does when the Superserver model is installed with Guardian.

Important

Do not try to use `fbguard` when running the traditional Classic server.

Thread-safe Client Library

Dmitry Yemanov
Vladyslav Khorsun
Alex Peshkov

Tracker reference [CORE-707](#).

The client libraries, including the embedded one, can now be used in multi-threaded applications without any application-level synchronization.

Improvements

Improvements implemented include:

Immediate Detection of Disconnected Clients on Classic

Vladyslav Khorsun

The Classic server now detects immediately when a Classic process has been broken by a client disconnection. Its response is to terminate any pending activity, roll back the active transaction and close the network connection.

Tracker reference [CORE-818](#).

Optimizations

Important optimizations include:

Data Retrieval

Dmitry Yemanov

An optimization improves data retrieval performance for tables from which no fields are accessed. This applies, for example to the **SELECT COUNT(*)** style of query.

Tracker reference [CORE-1598](#).

BLOB Memory Usage

Adriano dos Santos Fernandes

An optimization avoids memory consumption of <page size> bytes for each temporary BLOB created during assignment.

Tracker reference [CORE-1658](#).

Performance Improvement for Updates

V. Khorsun

The aim of this improvement was to reduce the amount of precedence writing that the engine undertakes when performing its “careful write” procedure for updates. The existing procedure had a noticeable effect on the performance of mass updates, especially for “updates-in-place”, where the same records are updated more than once in the same transaction. In the worst case, a page could be written to disk for every single new record version created during an update.

See the [Tracker reference \(CORE-2672\)](#) for a simplified technical description of what was involved internally.

The solution addresses the potentially time-consuming process involved in protecting the write operations from causing circular references between the pages on which new record versions and back versions are placed to maintain correct precedence.

Cache Size Limit Increased for 64-bit Servers

V. Khorsun

Tracker reference [CORE-1687](#)

Previous 64-bit Firebird server versions could not benefit from 64-bit address space and be configured for more than 2 GB (16K * 128 K) of database cache. The problem has been rectified in this version. On 64-bit Firebird, if the resources are available, it is now possible to configure cache large enough to accommodate a database of 5-10 GB completely in RAM .

Although Firebird's caching can get a lot of help from the filesystem cache, it is a feature that could be important for high-throughput systems whose load is mainly reads. The theoretical upper limit for caches on x64 Firebird servers is now $2^{31} - 1$ (2,147,483,647) pages.

Default Database Location

A. Peshkov

Tracker reference [CORE-1643](#)

The configuration parameter *DatabaseAccess* now has more “meaning” attached to it. In the absence of any other indication, the first location defined in the “Restrict” list for *DatabaseAccess* is taken by the engine as the default location for creating a new database and for locating a database where the connection parameters do not specify either an alias or the full path specification.

The seek logic is similar to that use for finding external tables from the Restrict list supplied to the *External-FileAccess* parameter, viz.,

1. All directories in the Restrict list are searched first.
2. If the database specified is not found:
 - If CREATE DATABASE is involved, then the first location in the Restrict list is used.
 - Otherwise, the attach fails in the expected fashion.

NOTE :: Current Working Directory

This feature does not suppress the use of the current working directory as the implicit location of the specified database file for direct local connections. The Y-valve handles the path resolution in these cases, just as it ever did.

For a stand-alone server working via the remote subsystem, trying to connect using the database file name with no path, although unlikely, is not recommended, since there is not really any way to be certain where the database would “land”. On Windows, for example, under these conditions the current working directory would be `%system%`.

DLL Loading for Windows Embedded Engine

Adriano dos Santos Fernandes

The root determination mechanism for the Windows embedded engine has been changed to avoid common problems that occur when an installation of the application structure encounters “DLL Hell”. Previously, the implicit root directory was the directory containing the user application's main executable file. Now it is the directory where the renamed *fbembed.dll* library is located.

Tracker reference [CORE-1814](#).

Large External Table Support Enabled

Vlad Khorsun

Previous Firebird versions used 32-bit I/O when working with external tables, limiting the size of the external file to < 2 GB. The mechanism has been enhanced to use 64-bit I/O on filesystems that support it, effectively eliminating the 2 GB limit.

Tracker reference [CORE-2492](#).

Statistics Now Work Properly with 64-bit Values

V. Khorsun
A. Peshkov

Tracker reference [CORE-2619](#)

Memory and other statistics did not work properly 64-bit values in older Firebird versions. The issue had two parts:

- a. To make the engine use 64-bit integers for statistics, the internals of **AtomicCounter** were changed.
- b. The *isql* and *qli* tools had to be taught to work with 64-bit values.

Incompatibility with Older Clients

To enable the 32-bit tools to work correctly with a 64-bit server, it was necessary to introduce some new internal API functions (**struct perf64** and **perf64_xxx**) and change *isql* and *qli* to use them. This means that the *isql* and *qli* programs in V.2.5 are not compatible with older Firebird clients.

UDFs Safeguard

Adriano dos Santos Fernandes

Tracker reference [CORE-1937](#).

When a string UDF is written to return a pointer not allocated by the same runtime as the Firebird server is accessing, the presence of the `FREE_IT` keyword in its declaration corrupts memory and crashes the server. As a safeguard against such dysfunctional UDFs, the engine now

1. detects such UDFs and throws an exception
2. depends on the presence of the updated *ib_util* library in the path for all server models, including embedded

Diagnostics

Transaction Diagnostics

Claudio Valderrama

Better diagnostics and error reporting when TPB contents are malformed. The new TPB validation logic now rejects:

- explicitly conflicting options within the same category, e.g., **{WAIT}** and **{NOWAIT}** specified together, or **{READ COMMITTED}** and **{SNAPSHOT}**, or **{READ ONLY}** and **{WRITE}**
- options making no sense, e.g. **[NO] RECORD VERSION** specified for a **SNAPSHOT** isolation mode
- incorrect order of table reservation options, e.g. **{PROTECTED READ <TABLE>}** instead of **{READ <TABLE> PROTECTED}**

Tracker reference [CORE-1600](#).

Access Privilege Error Messages

Alex Peshkov

Both table and column names are now reported when access privilege exceptions occur for a column.

Tracker reference [CORE-1234](#).

Message Improvement

V. Khorsun

Tracker reference [CORE-2587](#)

The diagnostic message when the engine cannot create shared memory that has already been mapped by another engine process in another Windows session is now a bit more user-friendly. It used to say:

The requested operation cannot be performed on a file with a user-mapped section open.

Now, it says:

Database is probably already opened by another engine instance in another Windows session.

Metadata Improvements

Preserve Character Set Default Collation

Adriano dos Santos Fernandes

An improvement allows the current value of RDB\$DEFAULT_COLLATE_NAME in the system table RDB\$CHARACTER_SETS to survive the backup/restore cycle. The mechanism for such customisation is the new [ALTER CHARACTER SET](#) command.

Tracker reference [CORE-789](#).

Chapter 4

Changes to the Firebird API and ODS

ODS (On-Disk Structure) Changes

On-disk structure (ODS) changes include the following:

New ODS Number

Firebird 2.5 creates databases with an ODS (On-Disk Structure) version of 11.2

Maximum Page Size

The maximum page size remains 16 KB (16384 bytes).

Maximum Number of Page Buffers in Cache

The maximum number of pages that can be configured for the database cache depends on whether the database is running under 64-bit or 32-bit Firebird:

- 64-bit :: $2^{31} - 1$ (2,147,483,647) pages
- 32-bit :: 128,000 pages, i.e., unchanged from V.2.1

API (Application Programming Interface) Extensions

Additions to the Firebird API include.-

BLOB Conversions Enabled in the API Functions

A. dos Santos Fernandes

Tracker reference [CORE-3446](#).

Conversion either way between BLOBs and other types are enabled in the API functions (XSQLVAR or blr messages).

- BLOBs can now be moved to or from different types in the execute and fetch calls.
- For input parameters, it allows a parameter to be put as a string without the need to create and fill a BLOB from the client side.
- For output (execute or fetch), it will be helpful for an application that knows its data and can describe a BLOB as a string with its maximum length.

Connection Strings & Character Sets

A. dos Santos Fernandes

Previous versions had no way to interoperate with the character set(s) used by the operating system and its filesystem. Firebird 2.5 has been made “environmentally aware” with regard to the file names of databases, other files and string parameters generally, when accessed through and/or passed in API connection requests. This change significantly improves Firebird's ability to accept and work with file names and other parameters containing characters that are not in the ASCII subset.

Only DPB Connections Support this Feature

In the current implementation, only connections made through the DPB (database parameter block) support this feature. It is not supported for Services API (*isc_spb**) functions.

isc_dpb_utf8_filename

The new connection option **isc_dpb_utf8_filename** has been introduced, to enable Firebird to be specifically informed that the file name or other character item being passed is in the UTF8 (UTF-8) character set. If the option is not used, the character set defaults to the codepage of the operating system.

Client-Server Compatibility

New client, older server

If the client is V.2.5 or newer and it is connecting to a pre-V.2.5 remote server, using the **isc_dpb_utf8_filename** option causes the *client* to convert the file name from UTF-8 to the *client codepage* before passing it to the server. It removes the **isc_dpb_utf8_filename** option from the DPB.

Compatibility is assured when the same codepage is being used on both the the client and server stations.

New client, new server, without isc_dpb_utf8_filename

If the client is V.2.5 or newer and it is connecting to a V.2.5 or newer remote server without using the **isc_dpb_utf8_filename**, the client converts the file name from the OS codepage to UTF-8 and inserts the **isc_dpb_utf8_filename** option into the DPB.

The file name received on the server is not subject to any special treatment. However, unlike older clients, the V.2.5 client may convert the file name automatically and insert the **isc_dpb_utf8_filename** option into the DPB. Compatibility is guaranteed, regardless, when the host and client are using the same code page.

New client, new server, with isc_dpb_utf8_filename

Whenever the **isc_dpb_utf8_filename** option is used, the client passes the unmodified filename to the server. The client thus always passes a UTF-8 file name to the server along with the **isc_dpb_utf8_filename** option.

Code Page Conversions

On Windows the code page used for conversions is Windows ANSI. On all other platforms, UTF-8 is used.

The operating system codepage and UTF-8 may not be the best choice for file names. For example, if you had a script or other text file for processing in *isql* or some other script-running tool that used another connection character set, it would not be possible to edit the file correctly using multiple character sets (code pages).

There is a solution: the *Unicode code point*. If used correctly, it enables correct interpretation of a character even if the client is older than V.2.5.

Using Unicode Code Points

Any Unicode character may now be encoded on the connection string file name as though it were an ASCII character. It is accomplished by using the symbol # as a prefix for a Unicode code point number (in hexadecimal format, similar to U+XXXX notation).

Write it as #XXXX with X being 0-9, a-f, A-F.

If one of the characters happens to be the literal #, you could either “double” the hash character (##) or use the code point number for it, #0023.

Note

The hash character is interpreted at the server with these new semantics, even if the client is older than v2.5.

Support for SQLSTATE Completion Codes

W. Oliver
D. Yemanov

Tracker reference [CORE-1761](#).

A new client-side API function, **fb_sqlstate()** is available to convert the status vector item for an error into the corresponding SQL-2003 standard 5-alphanumeric SQLSTATE.

- The SQLSTATE code represents the concatenation of a 2-character SQL CLASS and a 3-character SQL SUBCLASS.
- Statements now return an SQLSTATE completion code.
- The *isql* utility now prints the SQLSTATE diagnostic for errors instead of the SQLCODE one
- The SQLCODE diagnostic is deprecated—meaning it will disappear in a future release
- (v.2.5.1) Exception handling syntax in PSQL, of the style **WHEN SQLSTATE**, has been added.

Deprecated SQLCODE

Although the SQLCODE is deprecated and use of the SQLSTATE is preferred, it remains in Firebird for the time being. The *isc_sqlcode()* API function is still supported, as is the **WHEN SQLCODE** exception handling.

Appendix A: **SQLSTATE** provides a list of all SQLSTATE codes in use in this release, along with the corresponding message texts.

“Efficient Unprepare”

W. Oliver
D. Yemanov

Tracker reference [CORE-1741](#).

The new option *DSQL_unprepare* (numeric value 4) for the API routine *isc_dsql_free_statement()* allows the DSQL statement handle to survive the “unpreparing” of the statement.

Previously, the *isc_dsql_free_statement()* function supported only *DSQL_close* (for closing a named cursor) and *DSQL_drop* (which frees the statement handle).

The API addition is:

```
#define DSQL_close 1
#define DSQL_drop 2
#define DSQL_unprepare 4
```

Cancel Operation Function

Alex Peshkov

New *fb_cancel_operation()* API call, allowing cancellation of the current activity being performed by some kind of blocking API call in the given connection.

Syntax

```
ISC_STATUS fb_cancel_operation(ISC_STATUS* status_vector,
                               isc_db_handle* db_handle,
                               ISC_USHORT option);
```

Parameters

status vector (*ISC_STATUS* status_vector*)

A regular status vector pointer structure.

db_handle (*pointer to a isc_db_handle*)

A regular, valid database handle. It identifies the attachment.

option (*unsigned short: symbol*)

Determines the action to be performed. The option symbols are:

-

fb_cancel_raise: cancels any activity related to the **db_handle** specified in the second parameter. The effect will be that, as soon as possible, the engine will try to stop the running request and return an exception to the caller via the status vector of the interrupted API call.

“..as soon as possible” will be, under normal conditions, at the next rescheduling point.

Example

```

Thread1:                                Thread2:
-----                                -----

isc_dsql_execute(status, ....)          fb_cancel_operation(cancel_status, ...)
.....                                  cancel_status[1] = 0;
status[1] == isc_cancelled;

```

- *fb_cancel_disable*: disables execution of *fb_cancel_raise* requests for the specified attachment. It can be useful when your program is executing critical operations, such as cleanup, for example.
- *fb_cancel_enable*: re-enables delivery of a cancel execution that was previously disabled. The 'cancel' state is effective by default, being initialized when the attachment is created.
- *fb_cancel_abort*: forcibly close client side of connection. Useful if you need to close a connection urgently. All active transactions will be rolled back by the server. 'Success' is always returned to the application. **Use with care !**

Usage

The cycle of *fb_cancel_disable* and *fb_cancel_enable* requests may be repeated as often as necessary. If the engine is already in the requested state there is no exception: it is simply a no-op.

Usually *fb_cancel_raise* is called when you need to stop a long-running request. It is called from a separate thread, not from the signal handler, because it is *not* async signal safe.

Pay attention to the asynchronous nature of this API call!

Another aspect of asynchronous execution is that, at the end of API call, the attachment's activity might be cancelled or it might not. The latter is always a possibility. The asynchronicity also means that returned status vector will almost always return *FB_SUCCESS*. Exceptions, though, are possible: a network packet error, for example.

Example

```

Thread A:
fb_cancel_operation(isc_status, &DB, fb_cancel_enable);
isc_dsql_execute_immediate(isc_status, &DB, &TR, 0, "long running statement", 3, NULL);
// waits for API call to finish...

Thread B:
fb_cancel_operation(local_status, &DB, fb_cancel_raise);

Thread A:
if (isc_status[1])
isc_print_status(isc_status); // will print "operation was cancelled"

```

Shutdown Functions

Alex Peshkov

This release exposes a variety of API functions for instigating server shutdowns of various types from client applications.

Two Interrelated *fb_shutdown** Functions

This release exposes two *fb_shutdown** functions that may be useful for embedded server applications: *fb_shutdown()* and *fb_shutdown_callback*.

Prototypes

```
typedef int (*FB_SHUTDOWN_CALLBACK)(const int reason, const int mask, void* arg);

int fb_shutdown(unsigned int timeout,
               const int reason);

ISC_STATUS fb_shutdown_callback(ISC_STATUS* status_vector,
                               FB_SHUTDOWN_CALLBACK callback_function,
                               const int mask,
                               void* arg);
```

fb_shutdown()

fb_shutdown() performs a smart shutdown of various Firebird subsystems (yValve, engine, redirector). It was primarily designed for use by the internal engine, since it is only applicable to the current process. It is exposed by the API for its possible usefulness to user applications in the embedded server environment.

Currently operational only for the embedded engine, this function terminates all the current activity, rolls back active transactions, disconnects active attachments and shuts down the embedded engine instance gracefully.

Important for Application Developers

fb_shutdown() does not perform a shutdown of a remote server to which your application might be concurrently attached. In fact, all of the Firebird client libraries—including the one in embedded—call it automatically at `exit()`, as long as the client is attached to at least one database or service.

Hence, it should never be called by a client in the context of a remote attachment.

Parameters

fb_shutdown() takes two parameters:

1. timeout in milliseconds
2. reason for shutdown

The reason codes (**const int reason**), which are negative, are listed in `ibase.h`: refer to constants starting with **fb_shutrsn**.

Note

When calling **fb_shutdown()** from your program, you must pass the value as *positive*, for it will be passed as an argument to **fb_shutdown_callback()** by way of your **callback_function**, the routine where you would code the appropriate actions.

Return Values

- A return value of zero means shutdown was successful
- A non-zero value means some errors occurred during the shutdown. Details will be written to `firebird.log`.

fb_shutdown_callback()

fb_shutdown_callback() sets up the callback function that is to be called during shutdown. It is a call that almost always returns successfully, although there are cases, such as an out-of-memory condition, which could cause it to return an error.

Parameters

fb_shutdown_callback() takes four parameters:

status vector (ISC_STATUS status_vector)*

A regular status vector pointer structure.

pointer to callback function (FB_SHUTDOWN_CALLBACK callback_function)

This points to the callback function you have written to perform the actions (if any) to be taken when the callback occurs.

Your callback function can take three parameters. The first and second parameters help to determine what action is to be taken in your callback:

1. reason for shutdown

Two shutdown reasons are of especial interest:

- `fb_shutrsn_exit_called`: Firebird is closing due to `exit()` or unloaded client/embedded library
- `fb_shutrsn_signal`, applies only to POSIX: a SIGINT or SIGTERM signal was caught

Note

Firebird code always uses negative reasons. Users are expected to use positive values when calling **fb_shutdown()** themselves.

2. actual value of the mask with which it was called

The purpose of this parameter to help determine whether the callback was invoked before or after engine shutdown.

3. argument passed to `fb_shutdown_callback()` by the user application

Can be used for any purpose you like and may be NULL.

Return Value from the Callback Function

If the callback function returns zero, it means it performed its job successfully. A non-zero return value is interpreted according to the call mask (see next parameter topic, below):

- For *fb_shut_postproviders* calls, it means some errors occurred and it will result in a non-zero value being returned from **fb_shutdown()**. It is the responsibility of the callback function to notify the world of the exact reasons for the error condition being returned.
- For *fb_shut_preproviders* calls, it means that shutdown will not be performed.

Tip

It is *NOT* a good idea to return non-zero if the shutdown is due to `exit()` having been called ! ;-)

call mask (const int mask)

Can have the following symbolic values:

- *fb_shut_preproviders*: callback function will be called before shutting down engine
- *fb_shut_postproviders*: callback function will be called after shutting down engine
- An ORed combination of them, to have the same function called in either case

Values for call mask

fb_shut_confirmation

Engine queries: Is everyone ready to shut down?

fb_shut_preproviders

Actions to be done before providers are closed

fb_shut_postproviders

Actions to be done when providers are already closed

fb_shut_finish

Final cleanup

Returning non-zero for *fb_shut_confirmation* (although not *fb_shut_preproviders*) means that shutdown will not be performed.

argument (void arg)*

This is the argument to be passed to **callback_function**.

Using the *fb_shutdown* Functions

Following is a sample of using the shutdown and shutdown callback feature to prevent your program from being terminated if someone presses Ctrl-C while it has database attachments.

```
#include <ibase.h>

// callback function for shutdown
static int ignoreCtrlC(const int reason, const int, void*)
{
    return reason == fb_shutrsn_signal ? 1 : 0;
}
```

```
}  
  
int main(int argc, char *argv[])  
{  
    ISC_STATUS_ARRAY status;  
    if (fb_shutdown_callback(status, ignoreCtrlC, fb_shut_confirmation, 0))  
    {  
        isc_print_status(status);  
        return 1;  
    }  
    // your code continues ...  
}
```

New isc_spb_prp_* Constants for Shutdown

The new database shutdown modes can now be set using calls to the Services API. A number of new **isc_spb_prp_*** constants are available as arguments.

isc_spb_prp_shutdown_mode and isc_spb_prp_online_mode

These arguments are used for shutting down a database and bringing it back on-line, respectively. Each carries a single-byte parameter to set the new shutdown mode, exactly in accord with the *gfx -shut* settings:

- `isc_spb_prp_sm_normal`
- `isc_spb_prp_sm_multi`
- `isc_spb_prp_sm_single`
- `isc_spb_prp_sm_full`

The shutdown request also requires the *type of shutdown* to be specified, viz., one of

- `isc_spb_prp_force_shutdown`
- `isc_spb_prp_attachments_shutdown`
- `isc_spb_prp_transactions_shutdown`

Each takes a 4-byte integer parameter, specifying the timeout for the shutdown operation requested.

Note

The older-style parameters are also supported and should be used to enter the default shutdown (currently 'multi') and online ('normal') modes.

Usage Examples

Following are a few examples of using the new parameters with the *fbsvmgr* utility. For simplicity, it is assumed that login has already been established. Each example, though broken to fit the page-width, is a single line command.

Shutdown database to single-user maintenance mode:

```
fbsvcmgr service_mgr action_properties dbname employee
prp_shutdown_mode prp_sm_single prp_force_shutdown 0
```

Next, enable multi-user maintenance:

```
fbsvcmgr service_mgr action_properties dbname employee
prp_online_mode prp_sm_multi
```

Now go into full shutdown mode, disabling new attachments for the next 60 seconds:

```
fbsvcmgr service_mgr action_properties dbname employee
prp_shutdown_mode prp_sm_full prp_attachments_shutdown 60
```

Return to normal state:

```
fbsvcmgr service_mgr action_properties dbname employee
prp_online_mode prp_sm_normal
```

Tighter Control Over Header-level Changes

Alex Peshkov

Several DPB parameters have been made inaccessible to ordinary users, closing some dangerous loopholes. In some cases, they are settings that would alter the database header settings and potentially cause corruptions if not performed under administrator control; in others, they initiate operations that are otherwise restricted to the SYSDBA. They are.-

- `isc_dpb_shutdown` and `isc_dpb_online`
- `isc_dpb_gbak_attach`, `isc_dpb_gfix_attach` and `isc_dpb_gstat_attach`
- `isc_dpb_verify`
- `isc_dpb_no_db_triggers`
- `isc_dpb_set_db_sql_dialect`
- `isc_dpb_sweep_interval`
- `isc_dpb_force_write`
- `isc_dpb_no_reserve`
- `isc_dpb_set_db_readonly`
- `isc_dpb_set_page_buffers` (on Superserver)

The parameter `isc_dpb_set_page_buffers` can still be used by ordinary users on Classic and it will set the buffer size temporarily for that user and that session only. When used by the SYSDBA on either Superserver or Classic, it will change the buffer count in the database header, i.e., make a permanent change to the default buffer size.

Important Note for Developers and Users of Data Access Drivers and Tools

This change will affect any of the listed DPB parameters that have been explicitly set, either by including them in the DPB implementation by default property values or by enabling them in tools and applications that access databases as ordinary users. For example, a Delphi application that included 'RESERVE PAGE SPACE=TRUE' and 'FORCED WRITES=TRUE' in its database Params property, which caused no problems when the application connected to Firebird 1.x, 2.0.1, 2.0.3, 2.04 or 2.1.0/2.1.1, now rejects a connection by a non-SYSDBA user with ISC ERROR CODE 335544788, "Unable to perform operation. You must be either SYSDBA or owner of the database."

New Trace Services for Applications

Vlad Khorsun

Five new services relating to the management of the new user trace sessions have been added to the Services Manager, each with its corresponding Services API action function.

isc_action_svc_trace_start

Starts a user trace session

Parameter(s)

`isc_spb_trc_name` : trace session name, string, optional
`isc_spb_trc_cfg` : trace session configuration, string, mandatory

The mandatory parameter is a string encompassing the text for the desired configuration. A template file named `fbtrace.conf` is provided in Firebird's root directory as a guide to the contents of this string.

Note

1. Unlike system audit sessions, a user session does not read the configuration from a file. It will be the responsibility of the application developer to devise a mechanism for storing configurations locally at the client and retrieving them for run-time use.
2. Superfluous white space in the string is fine: it will simply be ignored.

Output

- A text message reporting the status of the operation, EITHER:

```
Can not start trace session. There are no trace plugins loaded
```

OR

```
Trace session ID NNN started
```

- In the second case, the results of the trace session in text format follow.

isc_action_svc_trace_stop

Stops a designated trace session

Parameter(s)

`isc_spb_trc_id` : trace session ID, integer, mandatory

Output

A text message providing the result (status) of the request:

- Trace session ID NNN stopped
- No permissions to stop other user trace session
- Trace session ID NNN not found

isc_action_svc_trace_suspend

Suspends a designated trace session

Parameter(s)

`isc_spb_trc_id` : trace session ID, integer, mandatory

Output

A text message providing the result (status) of the request:

- Trace session ID NNN paused
- No permissions to change other user trace session
- Trace session ID NNN not found

isc_action_svc_trace_resume

Resumes a designated trace session that has been suspended

Parameter(s)

`isc_spb_trc_id` : trace session ID, integer, mandatory

Output

A text message providing the result (status) of the request:

- Trace session ID NNN resumed

- No permissions to change other user trace session
- Trace session ID NNN not found

isc_action_svc_trace_list

Lists existing trace sessions

No parameters

Output

A text message listing the trace sessions and their states:

- Session ID: <number>
- name: <string>. Prints the trace session name if it is not empty
- user: <string>. Prints the user name of the user that created the trace session
- date: YYYY-MM-DD HH:NN:SS, start date and time of the user session
- flags: <string>, a comma-separated set comprising some or all of the following:

active / suspend

Run state of the session.

admin

Shows *admin* if an administrator user created the session. Absent if an ordinary user created the session.

system

Shows *system* if the session was created by the Firebird engine (system audit session). Absent if an ordinary user created the session.

audit / trace

Indicates the kind of session: *audit* for an engine-created audit session or *trace* for a user trace session.

log full

Conditional, appears if it is a user trace session and the session log file is full.

Note

The output of each service can usually be obtained using a regular *isc_service_query* call with either of the *isc_info_svc_line* or *isc_info_svc_to_eof* information items.

Other Services API Additions

Alex Peshkov

Other additions to the Services API include:

Mapping for RDB\$ADMIN Role in Services API

Two tag items have been added to the services parameter block (SPB) to to enable or disable the **RDB\$ADMIN role** for a privileged operating system user when requesting access to the security database.

Note

This capability is implemented in the *gsec* utility by way of the new **-mapping** switch. Refer to the notes in the [relevant section](#) of the *Command-line Utilities* chapter.

Tag Item *isc_action_svc_set_mapping*

Enables the RDB\$ADMIN role for the appointed OS user for a service request to access *security2.fdb*.

Tag Item *isc_action_svc_drop_mapping*

Disables the RDB\$ADMIN role for the appointed OS user for a service request to access *security2.fdb*.

Parameter *isc_spb_sec_admin*

The new parameter **isc_spb_sec_admin**, is the SPB implementation of the new DDL syntax introduced to enable SYSDBA or another sufficiently privileged user to grant or revoke the RDB\$ADMIN role in the security database (*security2.fdb*) to or from an ordinary Firebird user. An ordinary user needs this role to acquire the same privileges as SYSDBA to create, alter or drop users in the security database.

isc_spb_sec_admin is of type *spb_long* with a value of either 0 (meaning REVOKE ADMIN ROLE) or a non-zero number (meaning GRANT ADMIN ROLE).

For more information, refer to the topic [CREATE/ALTER/DROP USER](#) in the chapter *Data Definition Language*.

Tag item *isc_spb_bkp_no_triggers*

This new SPB tag reflects the Services API side of the **-nodbtriggers** switch introduced in the *gbak* utility at V.2.1 to prevent database-level and transaction-level triggers from firing during backup and restore. It is intended for use as a member of the **isc_spb_options** set of optional directives that includes items like **isc_spb_bkp_ignore_limbo**, etc.

Tag item *isc_spb_res_metadata_only*

Tracker reference: [CORE-3462](#).

If you pass the **isc_spb_bkp_metadata_only** tag to the restore service option in the SPB, it will perform a metadata-only restore. Many tools, including Firebird's *fbvcmgr* do not provide for specifying a backup option in a restore request.

(v.2.5.1) To avoid confusion, the tag **isc_spb_res_metadata_only**, simply reimplementing **isc_spb_bkp_metadata_only** was added as a constant to the public header and *fbvcmgr*.

nBackup Support

Three additions were made to support nBackup actions in the Services API.

Backup and Restore

Tracker reference: [CORE-1758](#).

The nBackup utility performs two logical groups of operations: locking or unlocking a database and backing it up or restoring it. While there is no rationale for providing a service action for the lock/unlock operations—they can be requested remotely by way of an SQL language request for ALTER DATABASE—a Services API interface to the backup/restore operations is easily justified.

Backup and restore must be run on the host station and the only way to access them was by running nBackup.

The two new service actions now enabling nBackup backup and restore to be requested through the Services API are:

- `isc_action_svc_nbak` - incremental nbackup
- `isc_action_svc_nrest` - incremental database restore

The parameter items are:

- `isc_spb_nbk_level` - backup level (integer)
- `isc_spb_nbk_file` - backup file name (string)
- `isc_spb_nbk_no_triggers` - option to suppress database triggers

Usage Examples

Following are a few examples of using the new parameters with the `fbsvcmgr` utility. For simplicity, it is assumed that login has already been established. Each example, though broken to fit the page-width, is a single line command.

Create backup level 0:

```
fbsvcmgr service_mgr action_nbak dbname employee
        nbk_file e.nb0 nbk_level 0
```

Create backup level 1:

```
fbsvcmgr service_mgr action_nbak dbname employee
        nbk_file e.nb1 nbk_level 1
```

Restore database from those files:

```
fbsvcmgr service_mgr action_nrest dbname e.fdb
        nbk_file e.nb0 nbk_file e.nb1
```

Direct I/O Feature Support

isc_spb_nbk_direct on|off

This new tag enables the same action as calling **nbackup -D on|off** from the command line, to force direct I/O on or off. As with other nbackup-related tags, it is valid only with **action_nbak**.

Note

Information regarding the usage of this feature can be found in the notes about the [new nbackup switches](#) in the *Utilities* chapter.

FIX_FSS_DATA and FIX_FSS_METADATA Options Enabled in Services API

A. Peshkov

Tracker reference [CORE-2439](#)

The `FIX_FSS_DATA` and `FIX_FSS_METADATA` functions that were implemented as **gbak -restore** switches in Firebird 2.1 have been implemented in the core engine and exposed as corresponding tag constants for the **isc_action_svc_restore** structure in the Services API. Thus, developers now have a path for writing applications to automate the migration of older Firebird databases to the new on-disk structure.

The new SPB tags are **isc_spb_res_fix_fss_data** and **isc_spb_res_fix_fss_metadata**.

New Trace API

A new Trace API is under construction, providing a set of hooks which can be implemented as an external plug-in module to be called by the engine when any traced event happens. It is as yet undocumented, since it is subject to change in forthcoming sub-releases.

For a little more information about it, refer to the topic [Trace Plug-in Facilities](#) in the *Administrative Features* chapter.

Chapter 5

Reserved Words and Changes

Note

Asterisks (*) mark keywords that are reserved, or otherwise recognised by Firebird's grammar as keywords, but are not reserved words in the SQL standard.

Clean-up of Reserved Words

A. Peshkov

Tracker reference [CORE-2638](#)

The number of Firebird-specific reserved words has been reduced significantly, in order to mitigate the pain of keyword conflicts when converting other databases to Firebird. Where possible, words that are not reserved by the standard have been made non-reserved in Firebird's grammar.

A small list remains of words that are reserved in Firebird but not in the SQL standard. They are:

ADD *	DB_KEY *	GDSCODE *	INDEX *
LONG *	PLAN *	POST_EVENT *	RETURNING_VALUES *
SQLCODE *	VARIABLE *	VIEW *	

All other non-standard keywords previously reserved are now available for any reasonable purpose.

Newly Reserved Words

SIMILAR	SQLSTATE
---------	----------

Keywords Added as Non-reserved

AUTONOMOUS *	BIN_NOT *	CALLER *
CHAR_TO_UUID *	COMMON *	DATA
FIRSTNAME *	GRANTED	LASTNAME *
MIDDLENAME *	MAPPING *	OS_NAME *
SOURCE *	TWO_PHASE *	UUID_TO_CHAR *

Chapter 6

Configuration Parameter Additions and Changes

The following changes or additions to `firebird.conf` should be noted:

AuditTraceConfigFile

V. Khorsun

This parameter points to the name and location of the file that the Firebird engine is to read to determine the list of events required for the next system audit trace. By default, the value of this parameter is empty, indicating that no system audit tracing is configured.

Note

The template file `fbtrace.conf`, found in Firebird's root directory, contains the full list of available events, with format, rules and syntax for composing an audit trace configuration file.

For more information, see the topic [System Audit Session](#) in section [Trace and Audit Services] in the chapter about the new administrative features.

Parameters Affecting Filesystem Cache Usage

There are now two parameters for configuring how Firebird interacts with the filesystem cache.

FileSystemCacheSize

N. Samofatov

New in Firebird 2.5, *FileSystemCacheSize* controls the maximum amount of RAM used by a Windows file system cache on 64-bit Windows XP or a Microsoft Server 2003 host with Service Pack 1 or higher.

At the V.2.5 initial release, it has no effect on POSIX host systems.

The setting for this parameter is an integer expressing the percentage of the total physical RAM that is available to the OS. To be valid, settings must be within the range 10 (per cent) to 95 (per cent), or explicitly set to 0 to enforce the host caching settings. Numbers outside that range will assume the default, which is 30 (per cent).

As with any `firebird.conf` setting, changes will not take effect until the server process is restarted.

Windows Security Privileges

The OS user needs the `SeIncreaseQuotaPrivilege` in order to adjust the filesystem cache settings. This right is built in for users with Administrator privileges and for service accounts and it is also granted to the Firebird service account explicitly by the Windows installer.

Under other conditions, e.g., embedded, or where the Firebird server is run as an application, or in a custom service installation, the user may not have that privilege. The process startup does not fail as a result of this misconfiguration: it will write a warning to the `firebird.log` and start-up will simply proceed with the host OS settings.

FileSystemCacheThreshold

V. Khorsun

This parameter was introduced in V.2.1 as `MaxFileSystemCache`. Because its name has been changed, its description is repeated here to alert upgraders.

FileSystemCacheThreshold sets a threshold determining whether Firebird will allow the page cache to be duplicated to the filesystem cache or not. If this parameter is set to any (integer) value greater than zero, its effect depends on the current default size of the page cache: if the default page cache (in pages) is less than the value of `MaxFileSystemCache` (in pages) then filesystem caching is enabled, otherwise it is disabled.

Note

This applies both when the page cache buffer size is set implicitly by the `DefaultDBCACHEPages` setting or explicitly as a database header attribute. It applies to all platforms.

Thus,

- To disable filesystem caching always, set `FileSystemCacheThreshold` to zero
- To enable filesystem caching always, set `FileSystemCacheThreshold` an integer value that is sufficiently large to exceed the size of the database page cache. Remember that the effect of this value will be affected by subsequent changes to the page cache size.

Important

- The default setting for `FileSystemCacheThreshold` is 65536 pages, i.e. filesystem caching is enabled.
- Observe that, if the configured cache size affecting a particular database is greater than the *FileSystemCacheThreshold* then the setting for *FileSystemCacheSize* (see above) will have no effect on that database.

MaxFileSystemCache

MaxFileSystemCache, introduced in Firebird 2.1, is no longer a valid parameter.

ConnectionTimeout

D. Yemanov

On heavily loaded Windows systems, local connect (XNET) could fail due to the client timing out while waiting for the server to set the `xnet_response_event`. To help with this problem, the *ConnectionTimeout* parameter has been enhanced to affect XNET connections, in addition to TCP/IP.

Note

The caveat documented for this parameter, although still applicable to network transports, does not apply to XNET's protocol.

Authentication

A. Peshkov

On Windows server platforms, since V.2.1, *Authentication* has been used for configuring the server authentication mode if you need it to be other than the default.

The mode settings for v.2.5 are the same, viz.

- *trusted* makes use of Windows “trusted authentication”. Under the right conditions, this may be the most secure way to authenticate on Windows.
- *native* sets the traditional Firebird server authentication mode, requiring users to log in using a user name and password defined in the security database.
- *mixed* allows both.

Changes in V.2.5

- Under v.2.5, although the modes are unchanged, configuring 'mixed' or 'trusted' mode no longer confers SYSDBA privileges on Windows domain administrators automatically by default. Please [read the notes in the Administrative Features chapter](#) regarding the new RDB\$ADMIN role in ODS 11.2 databases and auto-mapping SYSDBA privileges to domain administrators.
- The default configuration has been changed from *mixed* to *native*. To enable trusted user authentication (whether *mixed* or *trusted*, it is now necessary to configure this parameter specifically.

Tracker reference [CORE-2376](#)

MaxUserTraceLogSize

V. Khorsun

Stores the maximum total size of the temporary files to be created by a user trace session using the new Trace functions in the Services API. The default limit is 10 MB. Use this parameter to raise or lower the maximum total size of the temporary files storing the output.

OldSetClauseSemantics

D. Yemanov

Before Firebird 2.5, the SET clause of the UPDATE statement assigned columns in the user-defined order, with the NEW column values being immediately accessible to the subsequent assignments. This did not conform to the SQL standard, which requires the starting value of the column to persist during execution of the statement.

Now, only the OLD column values are accessible to any assignment in the SET clause.

The *OldSetClauseSemantics* enables you to revert to the legacy behavior via the `OldSetClauseSemantics`, if required. Values are 1 for the legacy behaviour, 0 (the default) for the corrected behaviour.

Warning

- Changing this parameter affects *all* databases on your server.
- This parameter is provided as a temporary solution to resolve backward compatibility issues. It will be deprecated in future Firebird versions.

RemoteAuxPort For Classic and Superclassic

Dmitry Yemanov

Tracker entry: [CORE-2263](#)

Classic and SuperClassic servers can now be configured to listen for events on a single, designated *RemoteAux-Port* port, as SuperServer has been able to do since v.1.5.

This long-awaited improvement now enables applications that connect to databases over the Internet through a firewall or a secure tunnel to use events, regardless of the server model in use.

Use Hostname for RemoteBindAddress

Alex Peshkov

Tracker entry: [CORE-2094](#)

It is now possible to use the hostname of the host where the Firebird server is running to configure *RemoteBindAddress*, where previously, only an IP address was allowed.

Important

RemoteBindAddress can be used to “pin” user connections to a specific NIC card on the host server. Take care that the hostname specified is not associated concurrently with more than one IP address, anywhere! In particular, check the `etc/hosts` file on all stations, including the host station itself.

RemoteFileOpenAbility

Nickolay Samofatov

Tracker entry: [CORE-2263](#)

Code from Red Soft was incorporated, to make this extreme option available to Windows and allow a database to be opened on a network share, in line with the long-time ability to allow access to a database on a NFS device on POSIX.

It is offered in the interests of maintaining feature consistency across platforms. There is no associated architectural change or any implication that its use in practice is considered safer now than in the past. However, it makes it possible to shadow databases to mapped locations and to connect to a database on an external filesystem for a specific, well-tested, safe purpose. An example given was a database kept under lock-and-key on a USB device that could be plugged in to a diskless workstation for performing an occasional, isolated security task.

Warning

READ THE NOTES IN FIREBIRD.CONF BEFORE YOU CONSIDER ACTIVATING THIS!

Chapter 7

Administrative Features

Certain improvements to Firebird's administrative features will be welcomed by many.

New RDB\$ADMIN System Role

Alex Peshkov

A new pre-defined system role RDB\$ADMIN has been added for transferring SYSDBA privileges to another user. Any user, when granted the role in a particular database, acquires SYSDBA-like rights when attaching to that database with the RDB\$ADMIN role specified.

To assign it, SYSDBA should log in to that database and grant the role RDB\$ADMIN to the user, in the same way one would grant any other role to a user. After the user has been granted the role, s/he must include it in the log-in in order to access the “superuser” privileges in that database.

Important

If the user attaches with a user database role passed in the DPB (connection parameters), it will not be replaced with RDB\$ADMIN, i.e., he/she will not get SYSDBA rights.

The following example transfers SYSDBA privileges to users named User1 and Admins\ADMINS. The second user in our example is typical for a Windows system user with access enabled via trusted authentication:

```
GRANT RDB$ADMIN TO User1;  
GRANT RDB$ADMIN TO "Admins\ADMINS";
```

Multiple Databases and Superusers

It should be understood that acquiring the RDB\$ADMIN role does not make a regular user into SYSDBA. Rather, it gives that user the same privileges as SYSDBA over the objects *in the database in which the role is granted to that user*.

- If the same user needs Superuser privileges in more than one database, the RDB\$ADMIN role must be explicitly granted for that user in *each* database.
- If more than one user is to have Superuser privileges in a database then each user needs to be granted the RDB\$ADMIN role.
- One user that has acquired the RDB\$ADMIN role in the database can grant it to another user.
- It is not necessary to specify WITH ADMIN OPTION (for the privilege to grant this role to others) or WITH GRANT OPTION (for the privilege to grant permissions on objects to other users without being the owner of those objects). The ADMIN and GRANT options are implicit.

System “Superusers”

On POSIX hosts, the *root* user always had SYSDBA privileges, but the same was not possible for a domain administrator on Windows until Firebird 2.1. In V.2.1, a configuration parameter, *Authentication*, was introduced, whereby a user logged in as a Windows domain administrator could automatically gain server access with SYSDBA privileges through trusted user authentication. On POSIX, the mechanism has not changed but on Windows, the introduction of the RDB\$ADMIN role has changed the way Windows administrators acquire SYSDBA privileges.

Windows trusted user authentication is no longer available by default!

By default, the *Authentication* parameter in `firebird.conf` is configured as *native*. It must be explicitly configured to either *trusted* or *mixed* to enable trusted user authentication.

Global Admin Privileges for Windows Administrators

For a trusted domain administrator to get SYSDBA access privileges on a Firebird server that is configured for trusted user authentication, the domain administrator must acquire the RDB\$ADMIN role. The [manual method](#) described above for ordinary users will work, by granting RDB\$ADMIN to each specific administrator, database by database.

However, there is a way for the SYSDBA to configure the server so that the RDB\$ADMIN role will be mapped to administrators automatically when they log into any database, thus arriving at a situation parallel to the association of root-privileged users on POSIX systems with the SYSDBA privileges. The new ALTER ROLE statement achieves this (and only this) purpose.

ALTER ROLE Statement

To configure a database to map the RDB\$ADMIN role to administrators automatically on a Windows server that is configured to enable trusted user authentication, the SYSDBA logs in to any database and issues the following statement:

```
ALTER ROLE RDB$ADMIN
  SET AUTO ADMIN MAPPING;
```

To revert to the default setting, preventing administrators from getting SYSDBA privileges automatically, issue this statement:

```
ALTER ROLE RDB$ADMIN
  DROP AUTO ADMIN MAPPING;
```

Services API Tag Items

The same effects are supported in the Services API by the provision of two tag items: **isc_action_svc_set_mapping** to enable the automatic mapping and **isc_action_svc_drop_mapping** to disable it.

These tags are supported in the *fbsvcmgr* utility.

Escalating RDB\$ADMIN Scope for User Management

The new DDL command **ALTER USER** enables an “ordinary” user (a regular Firebird user, a non-root user on POSIX or a trusted user on a Windows system where trusted authentication is enabled) the ability to change his or her password and/or personal name elements, while logged in to any database. Superusers can also use the same command to create and drop users. For more information about this new command, refer to the topic [CREATE/ALTER/DROP USER](#) in the chapter *Data Definition Language*.

Because *security2.fdb* is created as (or should be upgraded to) an ODS 11.2 database, it has the pre-defined RDB\$ADMIN role, too. Since no user—not even SYSDBA—can log into the security database, alternative means have been provided to enable the SYSDBA or a Superuser to apply the RDB\$ADMIN role in *security2.fdb* to an ordinary user that needs the ability to create or drop users. There are three ways, each having the equivalent effect:

1. Use the optional parameter **GRANT ADMIN ROLE** with a **CREATE USER** or **ALTER USER** statement.

Notes

Notice that **GRANT ADMIN ROLE** and **REVOKE ADMIN ROLE** here are not **GRANT/REVOKE** statements but 3-keyword parameters to the **CREATE USER** and **ALTER USER** statements. There is no system role named 'ADMIN'.

Any user that acquires the RDB\$ADMIN role in a database implicitly acquires the extended privileges **WITH ADMIN OPTION** and **WITH GRANT OPTION**.

Examples

To grant the RDB\$ADMIN role to user *alex* in the security database:

```
ALTER USER alex GRANT ADMIN ROLE;
```

To revoke the RDB\$ADMIN role from user *alex* in the security database:

```
ALTER USER alex REVOKE ADMIN ROLE;
```

To drop user *alex* and destroy his privileges in all databases:

```
DROP USER alex;
```

2. Use the *gsec* utility with the new switch **-admin**. The switch takes one argument: **YES** to apply the RDB\$ADMIN role to the user, or **NO** to revoke it. For more details, refer to the topic [Granting the RDB\\$ADMIN Role to An Ordinary User](#) in the *gsec* section of the *Utilities* chapter.
3. Use the new SPB parameter **isc_spb_sec_admin** which implements the assignment of the RDB\$ADMIN role for ordinary users in *security2.fdb* via a SPB connection. It is described in more detail in the chapter *Changes to the Firebird API and ODS*, in the topic [Parameter isc_spb_sec_admin](#).

The *fbsvmgr* utility also supports the use of this parameter.

Tip

Firebird 2.5 does not allow you to set up more than one security database on a server. From V.3.0, it is intended to be possible to have separate security databases for each database. For now, you can be connected to any database on the server (even *employee.fdb*) to update its one-and-only *security2.fdb*.

In future, it will be essential to send these requests from a database that is associated with the security database that is to be affected by them.

Trace and Audit Services

Vlad Khorsun

The new trace and audit facilities in v.2.5 were initially developed from the TraceAPI contributed to us by Nickolay Samofatov that he had developed for the Red Soft Database, a commercial product based on Firebird's code.

Overview of Features

The new trace and audit facilities enable various events performed inside the engine, such as statement execution, connections, disconnections, etc., to be logged and collated for real-time analysis of the corresponding performance characteristics.

A trace takes place in the context of a *trace session*. Each trace session has its own configuration, state and output.

The Firebird engine has a fixed list of events it can trace. It can perform two different sort of traces: a *system audit* and a *user trace*. How the engine forms the list of events for a session depends on which sort of trace is requested.

Tip

Every trace session is assigned a unique session ID. When any trace session begins, the Services Manager outputs this ID as the message

```
Trace session ID nnnn started
```

where **nnnn** is the ID, of course.

The System Audit Session

A system audit session is started by the engine itself. To determine which events the session is “interested in”, it reads the contents of a *trace configuration file* as it goes to create the session.

A new parameter in *firebird.conf*, *AuditTraceConfigFile* points to the name and location of the file. There can be at most one system audit trace in progress. By default, the value of this parameter is empty, indicating that no system audit tracing is configured.

A configuration file contains list of traced events and points to the placement of the trace log(s) for each event. It is sufficiently flexible to allow different sets of events for different databases to be logged to separate log files. The template file `fbtrace.conf`, found in Firebird's root directory, contains the full list of available events, with format, rules and syntax for composing an audit trace configuration file.

Tip About the `fbtrace.conf` File

The file contains a large amount of commented text explaining the purpose and syntax of each entry. As sub-releases progress, keep an eye on new events and facilities that will be added from time to time to improve the tracing capabilities.

For example, a late pre-release enhancement enables Services events to be traced by name and targeted using include and exclude filters.

As another example, the matching algorithm for path names was improved to interpret strings in the configuration file according to the platform character set and, basing the strings on UTF-8, to apply platform rules such as treating file names as case-insensitive on Windows and case-sensitive on POSIX. (Tracker reference [CORE-2404](#), A. dos Santos Fernandes).

User Trace Sessions

A user trace session is managed by user, using some new calls to the Services API. There are five new service functions for this purpose:

- start: **`isc_action_svc_trace_start`**
- stop: **`isc_action_svc_trace_stop`**
- suspend: **`isc_action_svc_trace_suspend`**
- resume: **`isc_action_svc_trace_resume`**
- list all known trace sessions: **`isc_action_svc_trace_list`**

The syntax for the Services API calls are discussed in the topic [New Trace Functions for Applications](#) in the chapter entitled *Changes to the Firebird API and ODS*.

Workings of a User Trace Session

When a user application starts a trace session, it sets a session name (optional) and the session configuration (mandatory). The session configuration is a text file conforming to the rules and syntax modelled in the `fbtrace.conf` template that is in Firebird's root directory, apart from the lines relating to placement of the output.

Note

Such files obviously do not live on the server. It will be the job of the application developer to design a suitable mechanism for storing and retrieving texts for passing in the user trace request.

For example, the command-line `fbvcmgr` utility supports a saved-file parameter, `trc_cfg`.

The output of a user session is stored in set of temporary files, each of 1 MB. Once a file has been completely read by the application, it is automatically deleted. By default, the maximum total size of the output is limited to 10 MB. It can be changed to a smaller or larger value using the `MaxUserTraceLogSize` in `firebird.conf`.

Once the user trace session service has been started by the application, the application has to read its output, using calls to `isc_service_query()`. The service could be generating output faster than the application can read

it. If the total size of the output reaches the *MaxUserTraceLogSize* limit, the engine automatically suspends the trace session. Once the application has finished reading a file (a 1 MB part of the output) that file is deleted, capacity is returned and the engine resumes the trace session automatically.

At the point where the application decides to stop its trace session, it simply requests a *detach* from the service. Alternatively, the application can use the **isc_action_svc_trace_*** functions to suspend, resume or stop the trace session at will.

Tip

The name of the character set of the attachment is included in any corresponding trace log records, placed between **user:role** and **protocol:port**, e.g.,

A.FDB (ATT_36, SYSDBA:NONE, WIN1251, TCPv4:127.0.0.1)

If no character set is specified in the DPB, NONE is written to the attachment character set slot in the trace log record.

([CORE-3008](#))

Who Can Manage Trace Sessions?

Any user can initiate and manage a trace session. An ordinary user can request a trace only on its own connections and cannot manage trace sessions started by any other users. Administrators can manage any trace session.

Abnormal Endings

If all Firebird processes are stopped, no user trace sessions are preserved. What this means is that if a Superserver or Superclass process is shut down, any user trace sessions that were in progress, including any that were awaiting a *resume* condition, are fully stopped and a *resume* cannot restart them.

Note

This situation doesn't apply to the Classic server, of course, since each connection involves its own dedicated server instance. Thus, there is no such thing as “shutting down” a Classic server instance. No service instance can outlive the connection that instigated it.

User Trace Sample Configuration Texts

The following samples provide a reference for composing configuration texts for user trace sessions.

- a. Trace prepare, free and execution of all statements within connection 12345

```
<database mydatabase.fdb>
enabled                true
connection_id          12345
log_statement_prepare  true
log_statement_free     true
log_statement_start    true
log_statement_finish   true
```

```
time_threshold      0
</database>
```

- b. Trace all connections of given user to database mydatabase.fdb, logging executed INSERT, UPDATE and DELETE statements and nested calls to procedures and triggers, and show corresponding PLANS and performance statistics.

```
<database mydatabase.fdb>
enabled             true
include_filter      %(INSERT|UPDATE|DELETE)%
log_statement_finish true
log_procedure_finish true
log_trigger_finish  true
print_plan          true
print_perf          true
time_threshold      0
</database>
```

Command-line Requests for User Trace Services

A new command-line utility, named *fbtracemgr*, has been added for working interactively with trace services. It has its own syntax of switches and parameters, discussed in detail, with examples, in the [Utilities chapter](#).

As well, the general Services utility, *fbsvcmgr*, can be used for submitting service requests from the command line, as exemplified in the following examples.

- a. Start a user trace named “My trace” using a configuration file named `fbtrace.conf` and read its output on the screen:

```
fbsvcmgr service_mgr action_trace_start trc_name "My trace" trc_cfg fbtrace.conf
```

To stop this trace session, press Ctrl+C at the *fbsvcmgr* console prompt. (See also (e), below).

- b. List trace sessions:

```
fbsvcmgr service_mgr action_trace_list
```

- c. Suspend trace session with ID 1

```
fbsvcmgr service_mgr action_trace_suspend trc_id 1
```

- d. Resume trace session with ID 1

```
fbsvcmgr service_mgr action_trace_resume trc_id 1
```

- e. Stop trace session with ID 1

```
fbsvcmgr service_mgr action_trace_stop trc_id 1
```

Tip

List sessions (see b.) in another console, look for the ID of a session of interest and use it in the current console to stop the session.

Trace Scope on Windows

Tracker reference [CORE-2588](#)

On Windows, the trace tools exhibit shared memory conflicts if multiple engine instances in different Windows sessions are allowed to run *trace* in global namespace. Tracing scope has therefore been restricted to only those processes that are accessible from the current Windows session.

Use Cases

There are three general use cases:

1. **Constant audit of engine activity**

This is served by system audit trace. Administrator creates or edits the trace configuration file, sets its name via the *AuditTraceConfigFile* setting in *firebird.conf* and restarts Firebird. Later, the administrator could suspend, resume or stop this session without needing to restart Firebird.

Important

To make audit configuration changes known to the engine, Firebird must be restarted.

2. **On-demand interactive trace of some (or all) activity in some (or all) databases**

An application (which could be the *fbtracemgr* utility) starts a user trace session, reads its output and shows traced events to the user in real time on the screen. The user can suspend and resume the trace and, finally, stop it.

3. **Engine activity collection for a significant period of time (a few hours or perhaps even a whole day) for later analysis**

An application starts a user trace session, reading the trace output regularly and saving it to one or more files. The session must be stopped manually, by the same application or by another one. If multiple trace sessions are running, a listing can be requested in order to identify the session of interest.

Trace Plug-in Facilities

A new Trace API will provide a set of hooks which can be implemented as an external plug-in module to be called by the engine when any traced event happens. A plug-in will assume responsibility for logging such events in some custom way.

This Trace API exists in Firebird 2.5 and is in use. However, since it will be changed in forthcoming sub-releases, it is not officially published and must be regarded as unstable.

The implemented “standard” trace plug-in, `fbtrace.dll` (`.so`), resides in the `\plugins` folder of your Firebird 2.5 installation.

Monitoring Improvements

Dmitry Yemanov

Firebird 2.5 sees the enhancement of the “MON\$” database monitoring features introduced in V.2.1, with new tables delivering data about context variables and memory usage in ODS 11.2 and higher databases. Also, in these databases, it becomes possible to terminate a client connection from another connection through the MON \$ structures.

Extended Access for Ordinary Users

The original design allowed non-privileged database users to see monitoring information pertaining only to their `CURRENT_CONNECTION`. Now they can request information for any attachment that was authenticated using the same user name.

Tracker reference [CORE-2233](#).

Notes

1. For an application architecture that entails a middleware tier logging in multiple times concurrently with the same user name on behalf of different end users, consideration should be given to the impact on performance and privacy of exposing the monitoring features to the end users.
2. The same extension was implemented in V.2.1.2.

New MON\$ Metadata for ODS 11.2 Databases

Note

For the ODS 11.1 metadata please refer to the V.2.1 documentation.

Character Set Change for MON\$ Metadata

The system domain `RDB$FILE_NAME2`, that is used to define those columns in the MON\$ tables that pertain to file specifications has been altered from `CHARACTER SET NONE` to `CHARACTER SET UNICODE_FSS`. The columns currently affected are `MON$DATABASE_NAME`, `MON$ATTACHMENT_NAME` and `MON$REMOTE_PROCESS`. This change makes the affected data consistent with the updated v.2.5 handling of filespec and other character parameter items in the DPB.

(Tracker entry [CORE-2551](#), A. dos Santos Fernandes)

`MON$MEMORY_USAGE` (current memory usage)

- `MON$STAT_ID` (statistics ID)

- MON\$STAT_GROUP (statistics group)
 - 0: database
 - 1: attachment
 - 2: transaction
 - 3: statement
 - 4: call
- MON\$MEMORY_USED (number of bytes currently in use)

High-level memory allocations performed by the engine from its pools.
Can be useful for tracing memory leaks and for investigating unusual memory consumption and the attachments, procedures, etc. that might be responsible for it.
- MON\$MEMORY_ALLOCATED (number of bytes currently allocated at the OS level)

Low-level memory allocations performed by the Firebird memory manager.
These are bytes actually allocated by the operating system, so it enables the physical memory consumption to be monitored.

Note

Not all records have non-zero values. On the whole, only MON\$DATABASE and memory-bound objects point to non-zero “allocated” values. Small allocations are not allocated at this level, being redirected to the database memory pool instead.

- MON\$MAX_MEMORY_USED (maximum number of bytes used by this object)
- MON\$MAX_MEMORY_ALLOCATED (maximum number of bytes allocated from the operating system by this object)

MON\$CONTEXT_VARIABLES (known context variables)

- MON\$ATTACHMENT_ID (attachment ID)

Contains a valid ID only for session-level context variables.
Transaction-level variables have this field set to NULL.
- MON\$TRANSACTION_ID (transaction ID)

Contains a valid ID only for transaction-level context variables.
Session-level variables have this field set to NULL.
- MON\$VARIABLE_NAME (name of context variable)
- MON\$VARIABLE_VALUE (value of context variable)

Memory Usage in MON\$STATEMENTS and MON\$STATE

Memory usage statistics in MON\$STATEMENTS and MON\$STATE represent actual CPU consumption.

Tracker reference: [CORE-1583](#)

Usage Notes

Examples

“Top 10” statements ranked according to their memory usage:

```
SELECT FIRST 10
  STMT.MON$ATTACHMENT_ID,
  STMT.MON$SQL_TEXT,
  MEM.MON$MEMORY_USED
FROM MON$MEMORY_USAGE MEM
  NATURAL JOIN MON$STATEMENTS STMT
  ORDER BY MEM.MON$MEMORY_USED DESC
```

To enumerate all session-level context variables for the current connection:

```
SELECT
  VAR.MON$VARIABLE_NAME,
  VAR.MON$VARIABLE_VALUE
FROM MON$CONTEXT_VARIABLES VAR
  WHERE VAR.MON$ATTACHMENT_ID = CURRENT_CONNECTION
```

Terminating a Client

The MON\$ structures are, by design, read-only. Thus, user DML operations on them are prohibited. However, a mechanism is built in to allow deleting (only) of records in the MON\$STATEMENTS and MON\$ATTACHMENTS tables. The effect of this mechanism is to make it possible, respectively, to cancel running statements and, for ODS 11.2 databases, to terminate client sessions.

To cancel all current activity for a specified connection:

```
DELETE FROM MON$STATEMENTS
  WHERE MON$ATTACHMENT_ID = 32
```

To disconnect all clients except the “Me” connection:

```
DELETE FROM MON$ATTACHMENTS
  WHERE MON$ATTACHMENT_ID <> CURRENT_CONNECTION
```

Note

1.
 - A statement cancellation attempt becomes a void operation (“no-op”) if the client has no statements currently running.
 - Upon cancellation, the execute/fetch API call returns the *isc_cancelled* error code.
 - Subsequent operations are allowed.
2.
 - Any active transactions in the connection being terminated will have their activities cancelled immediately and they are rolled back.
 - Once terminated, the client session receives the *isc_att_shutdown* error code.
 - Subsequent attempts to use this connection handle will cause network read/write errors.

Security Hardening

Windows Platforms

No SYSDBA Auto-mapping (Windows)

In V.2.1, members of administrative Windows groups were mapped to SYSDBA by default. From V.2.5 forward, automatic SYSDBA mapping is controlled on per-database basis using the new SQL command

ALTER ROLE RDB\$ADMIN SET/DROP AUTO ADMIN MAPPING

Note

For a full overview of the RDB\$ADMIN role, refer to the topic [New RDB\\$ADMIN System Role](#) in the *Administrative Features* chapter.

Chapter 9

Data Definition Language (DDL)

V.2.5 brings a few significant additions and enhancements to DDL.

Quick Links

- [CREATE/ALTER/DROP USER](#)
- [Syntaxes for Altering Views](#)
- [SSP Extension for CREATE VIEW](#)
- [ALTER Mechanism for Computed Columns](#)
- [GRANTED BY Extension for GRANT and REVOKE](#)
- [ALTER ROLE](#)
- [REVOKE ALL](#)
- [Default COLLATION Attribute for a Database](#)
- [ALTER CHARACTER SET](#)
- [The DIFFERENCE FILE Argument for CREATE DATABASE](#)

Although this release emphasises architectural changes in the movement towards Firebird 3, a number of improvements and extensions have been implemented, in many cases as a response to feature requests in the Tracker.

General Alert

In v.2.5 and beyond, it is possible to alter the data type of a column, even if the column is referenced in a stored procedure or trigger, without an exception being thrown. Because compiled PSQL is stored statically as a binary representation (“BLR”) in a BLOB, the original BLR survives even a backup and restore. Being static, the BLR is not updated by the data type change, either.

This means that BLR that has references to the affected columns from the tables or downstream views, together with any variables declared using the TYPE OF syntax with reference to them is more than likely to be broken without warning. At best, the BLR will be flagged as “needing attention” but tests show that the flag is not set under all conditions.

In short, the engine now no longer stops you from changing the type of a field that has any dependencies in compiled PSQL. It will be a matter for your own change control to identify the affected procedures and triggers and recompile them to accommodate the changes.

Visibility of Procedure Definition Changes on Classic

Dmitry Yemanov

Tracker reference [CORE-2052](#).

One such change addressed the problem of the visibility of altered stored procedures to other connections to the Classic server. Now, such changes are made visible to the entire server as soon as the modifying transaction has completed its commit.

CREATE/ALTER/DROP USER

Alex Peshkov

Tracker reference [CORE-696](#).

In v.2.5, Firebird finally has syntax to enable user accounts on the server to be managed by submitting SQL statements when logged in to a regular database.

The CREATE USER and ALTER USER statements also include the parameters GRANT ADMIN ROLE and REVOKE ADMIN ROLE to enable a user with SYSDBA privileges user to grant the RDB\$ADMIN role in the security database to an ordinary user. For the usage description and a full overview of the RDB\$ADMIN role, refer to the topic [New RDB\\$ADMIN System Role](#) in the *Administrative Features* chapter.

Syntax Patterns

The SYSDBA, or a user with SYSDBA privileges in both the current database and the security database, can add a new user:

```
CREATE USER <username> {PASSWORD 'password'}  
  [FIRSTNAME 'firstname']  
  [MIDDLENAME 'middlename']  
  [LASTNAME 'lastname']  
  [GRANT ADMIN ROLE];
```

Note

The PASSWORD clause is required when creating a new user. It should be the initial password for that new user. The user can change it later, using ALTER USER.

The SYSDBA, or a user with SYSDBA privileges in both the current database and the security database, can change one or more of the password and proper name attributes of an existing user. Non-privileged users can use this statement to alter only their own attributes.

```
ALTER USER <username>  
  [PASSWORD 'password']  
  [FIRSTNAME 'firstname']  
  [MIDDLENAME 'middlename']  
  [LASTNAME 'lastname']  
  [{GRANT | REVOKE} ADMIN ROLE];
```

Note

At least one of the optional parameters must be present.

ALTER USER does not allow the <username> to be changed. If a different <username> is required, the old one should be deleted (dropped) and a new one created.

The SYSDBA, or a user with SYSDBA privileges in both the current database and the security database, can delete a user:

```
DROP USER <username>;
```

Restrictions

CREATE USER and DROP USER statements and GRANT | REVOKE ADMIN ROLE are available only for the SYSDBA, or a user that has acquired the RDB\$ADMIN role in both the current database and the security database.

An ordinary user can ALTER his own password and elements of his proper name. An attempt to modify another user will fail.

Examples

SYSDBA or a [user with equivalent privileges](#) in both the current database and the security database, can do:

```
CREATE USER alex PASSWORD 'test';  
  
ALTER USER alex FIRSTNAME 'Alex' LASTNAME 'Peshkov';  
  
ALTER USER alex PASSWORD 'IdQfA';
```

Syntaxes for Altering Views

Adriano dos Santos Fernandes

Previously, in order to alter a view definition, it was necessary to save the view definition off-line somewhere and drop the view, before recreating it with its changes. This made things very cumbersome, especially if there were dependencies. V.2.5 introduces syntaxes for ALTER VIEW and CREATE OR ALTER VIEW.

Tracker references are [CORE-770](#) and [CORE-1640](#).

ALTER VIEW

ALTER VIEW enables a view definition to be altered without the need to recreate (drop and create) the old version of the view and all of its dependencies.

CREATE OR ALTER VIEW

With CREATE OR ALTER VIEW, the view definition will be altered (as with ALTER VIEW) if it exists, or created if it does not exist.

Syntax Pattern

```
{create [ or alter ] | alter } view <view_name>  
  [ ( <field list> ) ]  
as <select statement>
```

Example

```
create table users (  
    id integer,  
    name varchar(20),  
    passwd varchar(20)  
);  
  
create view v_users as  
    select name from users;  
  
alter view v_users (id, name) as  
    select id, name from users;
```

Extensions for CREATE VIEW

The following extensions have been added for CREATE VIEW.

Specify Stored Procedure in FROM Clause

Adriano dos Santos Fernandes

Tracker reference [CORE-886](#).

A selectable stored procedure can now be specified in the FROM clause of a view definition.

Example

```
create view a_view as  
select * from a_procedure(current_date);
```

Create UNION View Without Column List

Dmitry Yemanov

Tracker reference [CORE-1402](#).

The column list can now be omitted from CREATE VIEW when the set is defined by a UNION.

Example

```
recreate view V1 as  
select d.rdb$relation_id from rdb$database d  
union all  
select d.rdb$relation_id from rdb$database d  
  
recreate view V2 as  
select d.rdb$relation_id as q from rdb$database d  
union all  
select d.rdb$relation_id as w from rdb$database d
```

Inferred Column Names

Adriano dos Santos Fernandes

Tracker reference [CORE-2424](#).

CREATE VIEW now has the capability to infer column names for views involving a GROUP BY clause or a derived table.

Example

```
create view V as
  select d.rdb$relation_id from rdb$database d
  group by d.rdb$relation_id

create view V as
  select a from (select 1 a from rdb$database);
```

ALTER Mechanism for Computed Columns

Adriano dos Santos Fernandes

Tracker reference [CORE-1454](#).

A column defined as COMPUTED BY <expression> can now be altered using the ALTER TABLE...ALTER COLUMN syntax. This feature can be used only to change the <expression> element of the column definition to a different expression. It cannot convert a computed column to non-computed or vice versa.

Syntax Pattern

```
alter table <table-name>
  alter <computed-column-name>
  [type <data-type>]
  COMPUTED BY (<expression>);
```

Examples

```
create table test (
  n integer,
  dn computed by (n * 2)
);
commit;
alter table test
  alter dn computed by (n + n);
```

Extensions for SQL Permissions

Alex Peshkov

The following extensions have been implemented in the area of SQL permissions (privileges).

GRANTED BY Clause

A GRANTED BY (or its alternative, AS) clause can now be optionally included in GRANT and REVOKE statements, enabling the grantor to be a user other than the CURRENT_USER (the default).

Syntax Pattern

```
grant <right> to <object>
  [ { granted by | as } [ user ] <username> ]
--
revoke <right> from <object>
  [ { granted by | as } [ user ] <username> ]
```

Tip

```
{ granted by | as }
```

GRANTED BY and AS are equivalent. GRANTED BY is the form recommended by the SQL standard. We support AS for compatibility with some other servers (Informix, for example).

Example

Logged in as SYSDBA:

```
create role r1; -- SYSDBA owns the role
/* next, SYSDBA grants the role to user1
with the power to grant it to others */
grant r1 to user1 with admin option;
/* SYSDBA uses GRANTED BY to exercise
user1's ADMIN OPTION */
grant r1 to public granted by user1;
```

In isql, we look at the effects:

```
SQL>show grant;
/* Grant permissions for this database */
GRANT R1 TO PUBLIC GRANTED BY USER1
GRANT R1 TO USER1 WITH ADMIN OPTION
SQL>
```

ALTER ROLE

Tracker reference [CORE-1660](#).

The new ALTER ROLE statement has a specialised function to control the assignment of SYSDBA permissions to Windows administrators during trusted authentication. It has no other purpose currently.

Note

For the usage description of `ALTER ROLE` and a full overview of the `RDB$ADMIN` role, refer to the topic `New RDB$ADMIN System Role` in the *Administrative Features* chapter.

REVOKE ALL

Tracker reference [CORE-2113](#).

When a user is removed from the security database or another authentication source, such as the operating system ACL, any associated cleanup of SQL privileges in databases has to be performed manually. This extension adds the capability to revoke all privileges in one stroke from a particular user or role.

Syntax Pattern

```
REVOKE ALL ON ALL FROM { <user list> | <role list> }
```

Example

Logged in as SYSDBA:

```
# gsec -del guest
# isql employee
fbs bin # ./isql employee
Database: employee
SQL> REVOKE ALL ON ALL FROM USER guest;
SQL>
```

Default COLLATION Attribute for a Database

Adriano dos Santos Fernandes

Tracker references [CORE-1737](#) and [CORE-1803](#).

An ODS 11.2 or higher database can now have a default `COLLATION` attribute associated with the default character set, enabling all text column, domain and variable definitions to be created with the same collation order unless a `COLLATE` clause for a different collation is specified.

The `COLLATION` clause is optional. If it is omitted, the default collation order for the character set is used.

Tip

Note also that the default collation order for a character set used in a database can now also be changed, thanks to the introduction of syntax for `ALTER CHARACTER SET`.

Syntax Pattern

```
create database <file name>
  [ page_size <page size> ]
```

```
[ length = <length> ]
[ user <user name> ]
[ password <user password> ]
[ set names <charset name> ]
[ default character set <charset name>
  [ [ collation <collation name> ] ] ]
[ difference file <file name> ]
```

Note

The parameter **DIFFERENCE FILE** is not a new one for CREATE DATABASE. It was quietly introduced in association with the *nBackup* utility in V.2.0 and has been lurking undocumented for three years. For more information, see [Evolution of CREATE DATABASE](#) at the end of this chapter.

Example

```
create database 'test.fdb'
  default character set win1252
  collation win_ptbr;
```

ALTER CHARACTER SET Command

Adriano dos Santos Fernandes

Tracker reference [CORE-1803](#).

New syntax introduced in this version, enabling the default collation for a character set to be set for a database.

The default collation is used when table columns are created with a given character set (explicitly, through a CHARACTER SET clause in the column or domain definition, or implicitly, through the default character set attribute of the database) and no COLLATION clause is specified.

Note

String constants also use the default collation of the connection character set.

The character set and collation of existing columns are not affected by ALTER CHARACTER SET changes.

Syntax Pattern

```
ALTER CHARACTER SET <charset_name>
  SET DEFAULT COLLATION <collation_name>
```

Example

```
create database 'people.fdb'
  default character set win1252;

alter character set win1252
  set default collation win_ptbr;
```

```
create table person (  
  id integer,  
  name varchar(50) /* will use the database default  
                   character set and the win1252  
                   default collation */  
);  
  
insert into person  
  values (1, 'adriano');  
insert into person  
  values (2, 'ADRIANO');  
  
/* will retrieve both records  
   because win_ptbr is case insensitive */  
select * from person where name like 'A%';
```

Tip

Another improvement allows the current value of RDB\$DEFAULT_COLLATE_NAME in the system table RDB\$CHARACTER_SETS to survive the backup/restore cycle.

Evolution of CREATE DATABASE

DDL support for the database header attributes introduced to register and change the nBackup states has been evolving since Firebird 2.0. Users of nBackup will be familiar with the ALTER DATABASE statements to begin and end the storage of the “delta” data, in a file apart from the main database, during a full *nBackup*.

Naming the Delta File for nBackup

ALTER DATABASE also has another argument that allows you to set the name that will be used for the file that stores the delta data. To quote from Paul Vinkenoog's excellent manual for *nBackup*:

By default, the delta file lives in the same directory as the database itself. It also has the same name as the database file, but with **.delta** appended. There is usually no reason to change this, but it can be done if need be—though not via nbackup itself.

Make a connection to the database with any client that allows you to enter your own SQL statements and give the command:

```
alter database  
  add difference file 'path-and-filename'
```

The custom delta file specification is persistent in the database; it is stored in the system table RDB\$FILES.

To revert to the default behaviour, issue the following statement:

```
alter database  
  drop difference file
```

Those who are still curious may study the details in the V.2.0 release notes or in the *nBackup* manual.

The *DIFFERENCE FILE* Argument for *CREATE DATABASE*

C. Valderrama

In Firebird 2.0, syntax for prescribing a custom name for the delta file appeared as an extra, optional argument for *CREATE DATABASE*. You can observe its placement in the [syntax pattern given above](#) for the topic *Default COLLATION Attribute for a Database*. As with *ALTER DATABASE*, the keyword for the argument is *DIFFERENCE FILE* and the argument is a valid file specification. It allows you to specify a custom name for the delta file that will be created whenever *ALTER DATABASE BEGIN BACKUP* is called, or when the equivalent *nBackup* shell command is invoked.

Examples of Usage

```
..\bin> isql -user sysdba -pass masterke
```

```
Use CONNECT or CREATE DATABASE to specify a database
SQL> create database 'ticks' difference file 'jaguar';
SQL> shell dir jaguar;
Volume in drive F is Firebird
Volume Serial Number is BCD9-4211
```

```
Directory of ..\bin
```

```
File Not Found
```

This is correct, we only defined the file name. Now it will be used:

```
SQL> alter database begin backup;
SQL> shell dir jaguar;
Volume in drive F is Firebird
Volume Serial Number is BCD9-4211
```

```
Directory of ..\bin
```

```
10-11-2009  00:59                8.192 jaguar
              1 File(s)                8.192 bytes
              0 Dir(s) 16.617.979.904 bytes free
```

```
SQL> alter database end backup;
SQL> shell dir jaguar;
Volume in drive F is Firebird
Volume Serial Number is BCD9-4211
```

```
Directory of ..\bin
```

```
SQL> drop database;
SQL> ^Z
```

Since the argument is a file name, it goes inside single quotes. Double-quotes are not valid: the statement will fail and return a confusing error message.

```
..\bin> isql -user sysdba -pass masterke
```

```
Use CONNECT or CREATE DATABASE to specify a database
```

```
SQL> create database 'ticks' DIFFERENCE FILE 'jaguar';
SQL> alter database add difference file 'leopard';
```

```
Statement failed, SQLCODE = -607
unsuccessful metadata update
-Difference file is already defined
```

The message is correct. Even though the delta was deleted by the ALTER DATABASE END BACKUP call, the name of the difference file is stored persistently and only one may exist.

```
SQL> alter database drop difference file;
SQL> alter database begin backup;
```

This does not break anything, because the engine will rescue the situation and create the delta using the default mechanism.

```
SQL> alter database add difference file 'leopard';
SQL> alter database begin backup;
SQL> alter database drop difference file;
```

```
Statement failed, SQLCODE = -607
unsuccessful metadata update
-Cannot change difference file name while database is in backup mode
```

This is correct validation.

```
SQL> alter database end backup;
SQL> drop database;
SQL> ^Z
```

Chapter 10

Data Manipulation Language (DML)

In this chapter are the additions and improvements that have been added to the SQL data manipulation language subset in Firebird 2.5.

Quick Links

- [RegEx Search Support using SIMILAR TO](#)
- [Hex Literal Support](#)
- [New UUID Conversion Functions](#)
- [Extension to LIST\(\) Function](#)
- [SOME_COL = ? OR ? IS NULL Predication](#)
- [Extensions to DATEADD\(\) and DATEDIFF\(\) Functions](#)
- [BIN_NOT\(\) Function Added](#)
- [Write to Temporary Tables in a Read-Only Database](#)
- [Optimizer Improvements](#)

RegEx Search Support using SIMILAR TO

Adriano dos Santos Fernandes

Tracker reference [CORE-769](#).

A new SIMILAR TO predicate is introduced to support regular expressions. The predicate's function is to verify whether a given SQL-standard regular expression matches a string argument. It is valid in any context that accepts Boolean expressions, such as WHERE clauses, CHECK constraints and PSQL IF() tests.

Syntax Patterns

```
<similar predicate> ::=
  <value> [ NOT ] SIMILAR TO <similar pattern> [ ESCAPE <escape character> ]

<similar pattern> ::= <character value expression> regular expression>

<regular expression> ::=
  <regular term>
  | <regular expression> <vertical bar> <regular term>

<regular term> ::=
  <regular factor>
  | <regular term> <regular factor>

<regular factor> ::=
  <regular primary>
```

```

| <regular primary> <asterisk>
| <regular primary> <plus sign>
| <regular primary> <question mark>
| <regular primary> <repeat factor>

```

```

<repeat factor> ::=
    <left brace> <low value> [ <upper limit> ] <right brace>

```

```

<upper limit> ::= <comma> [ <high value> ]

```

```

<low value> ::= <unsigned integer>

```

```

<high value> ::= <unsigned integer>

```

```

<regular primary> ::=
    <character specifier>
    | <percent>
    | <regular character set>
    | <left paren> <regular expression> <right paren>

```

```

<character specifier> ::=
    <non-escaped character>
    | <escaped character>

```

```

<regular character set> ::=
    <underscore>
    | <left bracket> <character enumeration>... <right bracket>
    | <left bracket> <circumflex> <character enumeration>... <right bracket>
    | <left bracket> <character enumeration include>... <circumflex> <character enumeration excl
        <right bracket>

```

```

<character enumeration include> ::= <character enumeration>

```

```

<character enumeration exclude> ::= <character enumeration>

```

```

<character enumeration> ::=
    <character specifier>
    | <character specifier> <minus sign> <character specifier>
    | <left bracket> <colon> <character class identifier> <colon> <right bracket>

```

```

<character specifier> ::=
    <non-escaped character>
    | <escaped character>

```

```

<character class identifier> ::=
    ALPHA
    | UPPER
    | LOWER
    | DIGIT
    | SPACE
    | WHITESPACE
    | ALNUM

```

Note

1. <non-escaped character> is any character except <left bracket>, <right bracket>, <left paren>, <right paren>, <vertical bar>, <circumflex>, <minus sign>, <plus sign>, <asterisk>, <underscore>, <percent>, <question mark>, <left brace> and <escape character>.
2. <escaped character> is the <escape character> succeeded by one of <left bracket>, <right bracket>, <left paren>, <right paren>, <vertical bar>, <circumflex>, <minus sign>, <plus sign>, <asterisk>, <underscore>, <percent>, <question mark>, <left brace> or <escape character>.

Table 10.1. Character class identifiers

Identifier	Description	Note
ALPHA	All characters that are simple latin letters (a-z, A-Z)	Includes latin letters with accents when using accent-insensitive collation
UPPER	All characters that are simple latin uppercase letters (A-Z)	Includes lowercase letters when using case-insensitive collation
LOWER	All characters that are simple latin lowercase letters (a-z)	Includes uppercase letters when using case-insensitive collation
DIGIT	All characters that are numeric digits (0-9)	
SPACE	All characters that are the space character (ASCII 32)	
WHITESPACE	All characters that are whitespaces (vertical tab (9), newline (10), horizontal tab (11), carriage return (13), formfeed (12), space (32))	
ALNUM	All characters that are simple latin letters (ALPHA) or numeric digits (DIGIT)	

Usage Guide

Return true for a string that matches <regular expression> or <regular term>:

```
<regular expression> <vertical bar> <regular term>
```

```
'ab' SIMILAR TO 'ab|cd|efg' -- true
'efg' SIMILAR TO 'ab|cd|efg' -- true
'a' SIMILAR TO 'ab|cd|efg' -- false
```

Match zero or more occurrences of <regular primary>: <regular primary> <asterisk>

```
'' SIMILAR TO 'a*' -- true
'a' SIMILAR TO 'a*' -- true
'aaa' SIMILAR TO 'a*' -- true
```

Match one or more occurrences of <regular primary>: <regular primary> <plus sign>

```
' ' SIMILAR TO 'a+'      -- false
'a' SIMILAR TO 'a+'      -- true
'aaa' SIMILAR TO 'a+'    -- true
```

Match zero or one occurrence of <regular primary>: <regular primary> <question mark>

```
' ' SIMILAR TO 'a?'      -- true
'a' SIMILAR TO 'a?'      -- true
'aaa' SIMILAR TO 'a?'    -- false
```

Match exact <low value> occurrences of <regular primary>: <regular primary> <left brace> <low value> <right brace>

```
' ' SIMILAR TO 'a{2}'    -- false
'a' SIMILAR TO 'a{2}'    -- false
'aa' SIMILAR TO 'a{2}'   -- true
'aaa' SIMILAR TO 'a{2}'  -- false
```

Match <low value> or more occurrences of <regular primary>: <regular primary> <left brace> <low value> <comma> <right brace>:

```
' ' SIMILAR TO 'a{2,}'   -- false
'a' SIMILAR TO 'a{2,}'   -- false
'aa' SIMILAR TO 'a{2,}'  -- true
'aaa' SIMILAR TO 'a{2,}' -- true
```

Match <low value> to <high value> occurrences of <regular primary> <regular primary> <left brace> <low value> <comma> <high value> <right brace>:

```
' ' SIMILAR TO 'a{2,4}'  -- false
'a' SIMILAR TO 'a{2,4}'  -- false
'aa' SIMILAR TO 'a{2,4}' -- true
'aaa' SIMILAR TO 'a{2,4}' -- true
'aaaa' SIMILAR TO 'a{2,4}' -- true
'aaaaa' SIMILAR TO 'a{2,4}' -- false
```

Match any (non-empty) character: <underscore>

```
' ' SIMILAR TO '_'       -- false
'a' SIMILAR TO '_'       -- true
'l' SIMILAR TO '_'       -- true
'al' SIMILAR TO '_'      -- false
```

Match a string of any length (including empty strings): <percent>

```
' ' SIMILAR TO '%'       -- true
'az' SIMILAR TO 'a%z'    -- true
```

```
'a123z' SIMILAR TO 'a%z' -- true
'azx' SIMILAR TO 'a%z' -- false
```

Group a complete <regular expression> to use as one single <regular primary> as a sub-expression: <left paren> <regular expression> <right paren>

```
'ab' SIMILAR TO '(ab){2}' -- false
'aabb' SIMILAR TO '(ab){2}' -- false
'abab' SIMILAR TO '(ab){2}' -- true
```

Match a character identical to one of <character enumeration>: <left bracket> <character enumeration>... <right bracket>

```
'b' SIMILAR TO '[abc]' -- true
'd' SIMILAR TO '[abc]' -- false
'9' SIMILAR TO '[0-9]' -- true
'9' SIMILAR TO '[0-8]' -- false
```

Match a character not identical to one of <character enumeration>: <left bracket> <circumflex> <character enumeration>... <right bracket>

```
'b' SIMILAR TO '[^abc]' -- false
'd' SIMILAR TO '[^abc]' -- true
```

Match a character identical to one of <character enumeration include> but not identical to one of <character enumeration exclude>: <left bracket> <character enumeration include>... <circumflex> <character enumeration exclude>...

```
'3' SIMILAR TO '[:,DIGIT:]^3]' -- false
'4' SIMILAR TO '[:,DIGIT:]^3]' -- true
```

Match a character identical to one character included in <character class identifier>. Refer to the table of [Character Class Identifiers](#), above. May be used with <circumflex> to invert the logic (see above): <left bracket> <colon> <character class identifier> <colon> <right bracket>

```
'4' SIMILAR TO '[:,DIGIT:]' -- true
'a' SIMILAR TO '[:,DIGIT:]' -- false
'4' SIMILAR TO '[^[:,DIGIT:]]' -- false
'a' SIMILAR TO '[^[:,DIGIT:]]' -- true
```

Examples

```
create table department (
  number numeric(3) not null,
  name varchar(25) not null,
  phone varchar(14)
  check (phone similar to '\([0-9]{3}\) [0-9]{3}\-[0-9]{4}' escape '\')
);
```

```
insert into department
  values ('000', 'Corporate Headquarters', '(408) 555-1234');
insert into department
  values ('100', 'Sales and Marketing', '(415) 555-1234');
insert into department
  values ('140', 'Field Office: Canada', '(416) 677-1000');

insert into department
  values ('600', 'Engineering', '(408) 555-123'); -- check constraint violation

select * from department
  where phone not similar to '\\([0-9]{3}\\) 555\\-%' escape '\\';
```

Hex Literal Support

Bill Oliver

Adriano dos Santos Fernandes

Tracker reference [CORE-1760](#).

Support for hexadecimal numeric and binary string literals has been introduced.

Syntax Patterns

```
<numeric hex literal> ::=
  { 0x | 0X } <hexit> [ <hexit>... ]

<binary string literal> ::=
  { x | X } <quote> [ { <hexit> <hexit> }... ] <quote>

<digit> ::=
  0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9

<hexit> ::=
  <digit> | A | B | C | D | E | F | a | b | c | d | e | f
```

Numeric Hex Literals

- The number of <hexit> in the string cannot exceed 16.
- If the number of <hexit> is greater than eight, the constant data type is a signed BIGINT. If it is eight or less, the data type is a signed INTEGER.

Tip

That means 0xF0000000 is -268435456 and 0x0F0000000 is 4026531840.

Binary String Literals

The resulting string is defined as a CHAR($n/2$) CHARACTER SET OCTETS, where n is the number of <hexit>.

Examples

```
select 0x10, cast('0x0F0000000' as bigint)
  from rdb$database;
select x'deadbeef'
  from rdb$database;
```

New UUID Conversion Functions

Adriano dos Santos Fernandes

Tracker references [CORE-1656](#) and [CORE-1682](#).

Two new built-in functions, `UUID_TO_CHAR` and `CHAR_TO_UUID`, enable conversion between a UUID in the form of a CHAR(16) OCTETS string and the RFC4122-compliant form.

CHAR_TO_UUID()

The function `CHAR_TO_UUID()` converts the CHAR(32) ASCII representation of a UUID (XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX) to the CHAR(16) OCTETS representation, optimized for storage.

Syntax Model

```
CHAR_TO_UUID( <string> )
```

Example

```
select char_to_uuid('93519227-8D50-4E47-81AA-8F6678C096A1')
  from rdb$database;
```

UUID_TO_CHAR()

The function `UUID_TO_CHAR()` converts a CHAR(16) OCTETS UUID (as returned by the `GEN_UUID()` function) to the CHAR(32) ASCII representation (XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX).

Syntax Model

```
UUID_TO_CHAR( <string> )
```

Example

```
select uuid_to_char(gen_uuid())
  from rdb$database;
```

SOME_COL = ? OR ? IS NULL Predication

Adriano dos Santos Fernandes

Tracker reference [CORE-2298](#)

By popular request, particularly from Delphi programmers, support has been implemented for a predication that is able to “OR” test both the equivalence between a column and a parameter and whether the value passed to the parameter is NULL. This construct is often desired as a way to avoid the need to prepare one query to request a filtered result set and another for the same query without the filter.

Users of Delphi and other programming interfaces that apply client-side object names to parameters wanted the ability for the DSQL engine to recognise a usage like the following:

```
WHERE col1 = :param1 OR :param1 IS NULL
```

At the API level, the language interface translates the query to

```
WHERE col1 = ? OR ? IS NULL
```

That presented two problems:

1. While the programmer treated the parameter **:param1** as though it were a single variable with two references, the API could not: it is presented with two *parameters*
2. The second parameter is of an unknown data type and the program has no way to assign to it

What was needed to solve this problem was to introduce a new data type to handle the “? is NULL” condition and teach Firebird to do the right thing when it received such a request.

The implementation works like this. To resolve the first problem, the request must supply *two* parameters (for our Delphi example):

```
WHERE col1 = :param1 OR :param2 IS NULL
```

- If “param1” is not NULL, the language interface is required to assign the correct value for the first parameter, set the **XSQLVAR.sqlind** to NOT NULL and leave XSQLVAR.sqldata NULL.
- If “param2” is NULL, the language interface is required to set the XSQLVAR.sqlind of *both* parameters to NULL and leave the XSQLVAR.sqldata NULL.

In other words, for the parameter (?) in **? IS NULL**

:

- XSQLVAR.sqlind should be set in accordance with NULL/NON-NULL state of the parameter. This is the type of parameter that is described by the new constant **SQL_NULL**.
- The XSQLVAR.sqldata of a **SQL_NULL** type of parameter should always be passed by the client application as a NULL pointer and should never be accessed.

NULL Specified in the Output Set

When NULL is specified as an output constant (**select NULL from ...**), it continues to be described as CHAR(1), rather than by SQL_NULL. That may change in a future version.

Extension to LIST() Function

Adriano dos Santos Fernandes

Tracker reference [CORE-1453](#)

A string expression is now allowed as the delimiter argument of the LIST() function.

Example

```
SELECT
  DISCUSSION_ID,
  LIST(COMMENT, ASCII_CHAR(13))
FROM COMMENTS
GROUP BY DISCUSSION_ID;
```

Extension to DATEADD and DATEDIFF() Functions

Adriano dos Santos Fernandes

The WEEK unit was introduced for functions DATEADD and DATEDIFF.

MILLISECOND, SECOND, MINUTE and HOUR units are no longer invalid units to use with DATE arguments.

BIN_NOT() Function Added

Adriano dos Santos Fernandes

Completing the set of built-in binary functions added in v.2.1, new function BIN_NOT() returns the result of a bitwise NOT operation performed on its argument.

Syntax Pattern

```
BIN_NOT( <number> )
```

Example

```
select bin_not(flags) from x;
```

Write to Temporary Tables in a Read-Only Database

Vladyslav Khorsun

(V.2.5.1) Write operations to global temporary tables in a read-only database are enabled.

Tracker reference [CORE-3399](#).

Optimizer Improvements

Changes in optimizer logic that address recognised problems include:

CROSS JOIN Logic (D. Yemanov)

When a CROSS JOIN involved an empty table, the optimizer had no special logic to detect that the query was futile and return the empty set immediately. That shortcut logic has now been implemented.

Tracker reference [CORE-2200](#).

Note

The same change was implemented in V.2.1.2.

Derived Tables (A. dos Santos Fernandes)

The limit on the number of contexts available when using derived tables has been raised.

Tracker reference [CORE-2029](#).

Timing of DEFAULT Evaluation (A. dos Santos Fernandes)

Under rare conditions, the early evaluation of a DEFAULT value or expression defined for a column might give rise to a confused situation regarding the evaluation of an input value supplied for that column. The possibility was addressed by deferring the evaluation of DEFAULT and not actually performing the evaluation at all unless it was actually necessary.

Tracker reference [CORE-1842](#).

Index Use for NOT IN Searches (A. dos Santos Fernandes)

Better performance has been achieved for the NOT IN predicate by enabling the use of an index.

Tracker reference [CORE-1137](#).

Undo Log Memory Consumption (D. Yemanov)

Excessive memory consumption by the Undo log after a lengthy series of updates in a single transaction has been avoided.

Tracker reference [CORE-1477](#).

Other Improvements

Other changes to smooth out the little annoyances include:

FREE_IT Error Detection (A. dos Santos Fernandes)

Previously, a UDF declared with FREE_IT would crash if the pointer returned had not been allocated by the `ib_util_malloc()` function. Now, such a condition is detected, an exception is thrown and the pointer remains in its allocated state.

Tracker reference [CORE-1937](#).

“Expression evaluation not supported” message improved (C. Valderrama)

A number of secondary GDS codes were introduced to provide more details about an operation that fails with an “Expression evaluation not supported” exception, for example:

```
'Argument for @1 in dialect 1 must be string or numeric'  
'Strings cannot be added to or subtracted from DATE or TIME types'  
'Invalid data type for subtraction involving DATE, TIME or TIMESTAMP types'  
etc.
```

These detailed messages follow the GDS code for the **isc_expression_eval_err (expression evaluation not supported)** error in the status vector.

Tracker reference [CORE-1799](#).

Chapter 11

Procedural SQL (PSQL)

Several significant changes appear in Firebird's procedural language (PSQL), the language set for triggers, stored procedures and dynamic executable blocks, especially with regard to new extensions to the capabilities of EXECUTE STATEMENT. This release also heralds the arrival of the “autonomous transaction”.

Quick Links

- [Autonomous Transactions](#)
- [Borrow Database Column Type for a PSQL Variable](#)
- [New Extensions to EXECUTE STATEMENT](#)
 - [Context Issues](#)
 - [Authentication](#)
 - [Transaction Behaviour](#)
 - [Inherited Access Privileges](#)
 - [External Queries from PSQL](#)
 - [EXECUTE STATEMENT with Dynamic Parameters](#)
 - [Examples Using EXECUTE STATEMENT](#)
- [Other PSQL Improvements](#)
 - [Subqueries as PSQL Expressions](#)
 - [SQLSTATE as Context Variable](#)

Autonomous Transactions

Adriano dos Santos Fernandes

Tracker reference [CORE-1409](#).

This new implementation allows a piece of code to run in an autonomous transaction within a PSQL module. It can be handy for a situation where you need to raise an exception but do not want the database changes to be rolled back.

The new transaction is initiated with the same isolation level as the one from which it is launched. Any exception raised in a block within the autonomous transaction will cause changes to be rolled back. If the block runs through until its end, the transaction is committed.

Warning

Because the autonomous transaction is independent from the one from which is launched, you need to use this feature with caution to avoid deadlocks.

Syntax Pattern

```
IN AUTONOMOUS TRANSACTION
```

```
DO
  <simple statement | compound statement>
```

Example of Use

```
create table log (
  logdate timestamp,
  msg varchar(60)
);

create exception e_conn 'Connection rejected';

set term !;

create trigger t_conn on connect
as
begin
  if (current_user = 'BAD_USER') then
  begin
    in autonomous transaction
    do
    begin
      insert into log (logdate, msg) values (current_timestamp, 'Connection rejected');
    end

    exception e_conn;
  end
end!

set term ;!
```

Borrow Database Column Type for a PSQL Variable

Adriano dos Santos Fernandes

Tracker reference [CORE-1356](#).

This feature extends the implementation in v.2 whereby domains became available as “data types” for declaring variables in PSQL. Now it is possible to borrow the data type of a column definition from a table or view for this purpose.

Syntax Pattern

```
data_type ::=
  <builtin_data_type>
  | <domain_name>
  | TYPE OF <domain_name>
  | TYPE OF COLUMN <table or view>.<column>
```

Note

TYPE OF COLUMN gets only the type of the column. Any constraints or default values defined for the column are ignored.

Examples

```
CREATE TABLE PERSON (
  ID INTEGER,
  NAME VARCHAR(40)
);

CREATE PROCEDURE SP_INS_PERSON (
  ID TYPE OF COLUMN PERSON.ID,
  NAME TYPE OF COLUMN PERSON.NAME
)
AS
DECLARE VARIABLE NEW_ID TYPE OF COLUMN PERSON.ID;
BEGIN
  INSERT INTO PERSON (ID, NAME)
  VALUES (:ID, :NAME)
  RETURNING ID INTO :NEW_ID;
END
```

Hidden Trap!

In v.2.5 and beyond, it is possible to alter the data type of a column, even if the column is referenced in a stored procedure or trigger, without an exception being thrown. Because compiled PSQL is stored statically as a binary representation (“BLR”) in a BLOB, the original BLR survives even a backup and restore. Being static, the BLR is not updated by the data type change, either.

This means that, for variables declared using the TYPE OF syntax, as well as the affected columns from the tables, together with any view columns derived from them, the compiled BLR is broken by the change of data type. At best, the BLR will be flagged as “needing attention” but tests show that the flag is not set under all conditions.

In short, the engine now no longer stops you from changing the type of a field that has any dependencies in compiled PSQL. It will be a matter for your own change control to identify the affected procedures and triggers and recompile them to accommodate the changes.

New Extensions to EXECUTE STATEMENT

Unusually for our release notes, we begin this chapter with the full, newly extended syntax for the EXECUTE STATEMENT statement in PSQL and move on afterwards to explain the various new features and their usage.

```
[FOR] EXECUTE STATEMENT <query_text> [( <input_parameters> )]
  [ON EXTERNAL [DATA SOURCE] <connection_string>]
  [WITH {AUTONOMOUS | COMMON} TRANSACTION]
  [AS USER <user_name>]
  [PASSWORD <password>]
  [ROLE <role_name>]
  [WITH CALLER PRIVILEGES]
  [INTO <variables>]
```


Note

The order of the optional clauses is not fixed so, for example, a statement based on the following model would be just as valid:

```
[ON EXTERNAL [DATA SOURCE] <connection_string>]
[WITH {AUTONOMOUS | COMMON} TRANSACTION]
[AS USER <user_name>]
[PASSWORD <password>]
[ROLE <role_name>]
[WITH CALLER PRIVILEGES]
```

Clauses cannot be duplicated.

Context Issues

If there is no ON EXTERNAL DATA SOURCE clause present, EXECUTE STATEMENT is normally executed within the CURRENT_CONNECTION context. This will be the case if the AS USER clause is omitted, or it is present with its <user_name> argument equal to CURRENT_USER.

However, if <user_name> is not equal to CURRENT_USER, then the statement is executed in a separate connection, established without Y-Valve and remote layers, inside the same engine instance.

Note

In the absence of an AS USER <user_name> clause, CURRENT_USER is the default.

Authentication

Where server authentication is needed for a connection that is different to CURRENT_CONNECTION, e.g., for executing an EXECUTE STATEMENT command on an external datasource, the AS USER and PASSWORD clauses are required. However, under some conditions, the PASSWORD may be omitted and the effects will be as follows:

1. On Windows, for the CURRENT_CONNECTION (i.e., no external data source), trusted authentication will be performed *if it is active* and the AS USER parameter is missing, null or equal to CURRENT_USER.
2. If the external data source parameter is present and its <connection_string> refers to the same database as the CURRENT_CONNECTION, the effective user account will be that of the CURRENT_USER.
3. If the external data source parameter is present and its <connection_string> refers to a different database than the one CURRENT_CONNECTION is attached to, the effective user account will be the operating system account under which the Firebird process is currently running.

In any other case where the PASSWORD clause is missing, only *isc_dpb_user_name* will be presented in the DPB (attachment parameters) and native authentication will be attempted.

Transaction Behaviour

The new syntax has an optional clause for setting the appropriate transaction behaviour: `WITH AUTONOMOUS TRANSACTION` and `WITH COMMON TRANSACTION`. `WITH COMMON TRANSACTION` is the default and does not need to be specified. Transaction lifetimes are bound to the lifetime of `CURRENT_TRANSACTION` and are committed or rolled back in accordance with the `CURRENT_TRANSACTION`.

The behaviour for `WITH COMMON TRANSACTION` is as follows:

- a. Causes any transaction in an external data source to be started with the same parameters as `CURRENT_TRANSACTION`; otherwise
- b. Executes the statement inside the `CURRENT_TRANSACTION`; or
- c. May use another transaction that is started internally in `CURRENT_CONNECTION`.

The `WITH AUTONOMOUS TRANSACTION` setting starts a new transaction with the same parameters as `CURRENT_TRANSACTION`. That transaction will be committed if the statement is executed without exceptions or rolled back if the statement encounters an error.

Inherited Access Privileges

Vladyslav Khorsun

Tracker reference [CORE-1928](#).

By design, the original implementation of `EXECUTE STATEMENT` isolated the executable code from the access privileges of the calling stored procedure or trigger, falling back to the privileges available to the `CURRENT_USER`. In general, the strategy is wise, since it reduces the vulnerability inherent in providing for the execution of arbitrary statements. However, in hardened environments, or where privacy is not an issue, it could present a limitation.

The introduction of the optional clause `WITH CALLER PRIVILEGES` now makes it possible to have the executable statement inherit the access privileges of the calling stored procedure or trigger. The statement is prepared using any additional privileges that apply to the calling stored procedure or trigger. The effect is the same as if the statement were executed by the stored procedure or trigger directly.

Important

The `WITH CALLER PRIVILEGES` option is not compatible with the `ON EXTERNAL DATA SOURCE` option.

External Queries from PSQL

Vladyslav Khorsun

Tracker reference [CORE-1853](#).

`EXECUTE STATEMENT` now supports queries against external databases by inclusion of the `ON EXTERNAL DATA SOURCE` clause with its `<connection_string>` argument.

The <connection_string> Argument

The format of <connection_string> is the usual one that is passed through the API function *isc_attach_database()*, viz.

```
[<host_name><protocol_delimiter>]database_path
```

Character Set

The connection to the external data source uses the same character set as is being used by the CURRENT_CONNECTION context.

Access Privileges

If the external data source is on another server then the clauses AS USER <user_name> and PASSWORD <password> will be needed.

The clause WITH CALLER PRIVILEGES is a no-op if the external data source is on another server.

MORE INFORMATION REQUIRED. ROLES?

Note

Use of a two-phase transaction for the external connection is not available in V.2.5.

EXECUTE STATEMENT with Dynamic Parameters

Vladyslav Khorsun
Alex Peshkov

Tracker reference [CORE-1221](#).

The new extensions provide the ability to prepare a statement with dynamic input parameters (placeholders) in a manner similar to a parameterised DSQL statement. The actual text of the query itself can also be passed as a parameter.

Syntax Conventions

The mechanism employs some conventions to facilitate the run-time parsing and to allow the option of “naming” parameters in a style comparable with the way some popular client wrapper layers, such as Delphi, handle DSQL parameters. The API's own convention, of passing unnamed parameters in a predefined order, is also supported. **However, named and unnamed parameters cannot be mixed.**

The New Binding Operator

At this point in the implementation of the dynamic parameter feature, to avoid clashes with equivalence tests, it was necessary to introduce a new assignment operator for binding run-time values to named parameters. The new operator mimics the Pascal assignment operator:“:=”.

Syntax for Defining Parameters

```
<input_parameters> ::=
    <named_parameter> | <input_parameters>, <named_parameter>

<named_parameter> ::=
    <parameter name> := <expression>
```

Example for **named input parameters**

For example, the following block of PSQL defines both <query_text> and named <input_parameters> (<named_parameter>):

```
EXECUTE BLOCK AS
  DECLARE S VARCHAR(255);
  DECLARE N INT = 100000;
  BEGIN
    /* Normal PSQL string assignment of <query_text> */
    S = 'INSERT INTO TTT VALUES (:a, :b, :a)';

    WHILE (N > 0) DO
      BEGIN
        /* Each loop execution applies both the string value
           and the values to be bound to the input parameters */

        EXECUTE STATEMENT (:S) (a := CURRENT_TRANSACTION, b := CURRENT_CONNECTION)
        WITH COMMON TRANSACTION;
        N = N - 1;
      END
    END
  END
```

Example for **unnamed input parameters**

A similar block using a set of unnamed input parameters instead and passing constant arguments directly:

```
EXECUTE BLOCK AS
  DECLARE S VARCHAR(255);
  DECLARE N INT = 100000;
  BEGIN
    S = 'INSERT INTO TTT VALUES (?, ?, ?)';

    WHILE (N > 0) DO
      BEGIN
        EXECUTE STATEMENT (:S) (CURRENT_TRANSACTION, CURRENT_CONNECTION, CURRENT_TRANSACTION);
        N = N - 1;
      END
    END
  END
```

END
END

Note

Observe that, if you use both <query_text> and <input_parameters> then the <query_text> must be enclosed in parentheses, viz.

```
EXECUTE STATEMENT (:sql) (p1 := 'abc', p2 := :second_param) ...
```

Exception Handling

The handling of exceptions depends on whether the **ON EXTERNAL DATA SOURCE** is present.

ON EXTERNAL DATA SOURCE clause is present

If ON EXTERNAL DATA SOURCE clause is present, Firebird cannot interpret error codes supplied by the unknown data source so it interprets the error information itself and wraps it as a string into its own error wrapper (**isc_eds_connection** or **isc_eds_statement**).

The text of the interpreted remote error contains both error codes and corresponding messages.

1. Format of **isc_eds_connection** error

Template string

```
Execute statement error at @1 :\n@Data source : @3
```

Status-vector tags

```
isc_eds_connection,  
isc_arg_string, <failed API function name>,  
isc_arg_string, <text of interpreted external error>,  
isc_arg_string, <data source name>
```

2. Format of **isc_eds_statement** error

Template string

```
Execute statement error at @1 :\n@Statement : @3\nData source : @4
```

Status-vector tags

```
isc_eds_statement,  
isc_arg_string, <failed API function name>,  
isc_arg_string, <text of interpreted external error>,  
isc_arg_string, <query>,  
isc_arg_string, <data source name>
```

At PSQL level the symbols for these errors can be handled by treating them like any other *gdscodes*. For example

```
WHEN GDSCODE eds_statement
```

Note

Currently, the originating error codes are not accessible in a WHEN statement. The situation could be improved in future.

ON EXTERNAL DATA SOURCE clause is not present

If ON EXTERNAL DATA SOURCE clause is not present, the original status-vector with the error is passed as-is to the caller PSQL code.

For example, if a dynamic statement were to raise the *isc_lock_conflict* exception, the exception would be passed to the caller and could be handled using the usual handler:

```
WHEN GDSCODE lock_conflict
```

Examples Using EXECUTE STATEMENT

The following examples offer a sampler of ways that the EXECUTE STATEMENT extensions might be applied in your applications.

Test Connections and Transactions

A couple of tests you can try to compare variations in settings:

Test a) :Execute this block few times in the same transaction - it will create three new connections to the current database and reuse it in every call. Transactions are also reused.

```
EXECUTE BLOCK
  RETURNS (CONN INT, TRAN INT, DB VARCHAR(255))
AS
  DECLARE I INT = 0;
  DECLARE N INT = 3;
  DECLARE S VARCHAR(255);
BEGIN
  SELECT A.MON$ATTACHMENT_NAME FROM MON$ATTACHMENTS A
  WHERE A.MON$ATTACHMENT_ID = CURRENT_CONNECTION
  INTO :S;

  WHILE (i < N) DO
  BEGIN
    DB = TRIM(CASE i - 3 * (I / 3)
      WHEN 0 THEN '\\.\' WHEN 1 THEN 'localhost:' ELSE '' END) || :S;

    FOR EXECUTE STATEMENT
      'SELECT CURRENT_CONNECTION, CURRENT_TRANSACTION
      FROM RDB$DATABASE'
      ON EXTERNAL :DB
      AS USER CURRENT_USER PASSWORD 'masterkey' -- just for example
      WITH COMMON TRANSACTION
      INTO :CONN, :TRAN
```

```
DO SUSPEND;

  i = i + 1;
END
END
```

Test b) : Execute this block few times in the same transaction - it will create three new connections to the current database on every call.

```
EXECUTE BLOCK
  RETURNS (CONN INT, TRAN INT, DB VARCHAR(255))
AS
  DECLARE I INT = 0;
  DECLARE N INT = 3;
  DECLARE S VARCHAR(255);
BEGIN
  SELECT A.MON$ATTACHMENT_NAME
    FROM MON$ATTACHMENTS A
  WHERE A.MON$ATTACHMENT_ID = CURRENT_CONNECTION
    INTO :S;

  WHILE (i < N) DO
  BEGIN
    DB = TRIM(CASE i - 3 * (I / 3)
      WHEN 0 THEN '\\.\'
      WHEN 1 THEN 'localhost:'
      ELSE '' END) || :S;

    FOR EXECUTE STATEMENT
      'SELECT CURRENT_CONNECTION, CURRENT_TRANSACTION FROM RDB$DATABASE'
      ON EXTERNAL :DB
      WITH AUTONOMOUS TRANSACTION -- note autonomous transaction
      INTO :CONN, :TRAN
    DO SUSPEND;

    i = i + 1;
  END
END
```

Input Evaluation Demo

Demonstrating that input expressions evaluated only once:

```
EXECUTE BLOCK
  RETURNS (A INT, B INT, C INT)
AS
BEGIN
  EXECUTE STATEMENT (
    'SELECT CAST(:X AS INT),
      CAST(:X AS INT),
      CAST(:X AS INT)
    FROM RDB$DATABASE' )
    (x := GEN_ID(G, 1))
  INTO :A, :B, :C;

  SUSPEND;
```

END

Insert Speed Test

Recycling our earlier examples for input parameter usage for comparison with the non-parameterised form of EXECUTE STATEMENT:

```

RECREATE TABLE TTT (
  TRAN INT,
  CONN INT,
  ID INT);

-- Direct inserts:

EXECUTE BLOCK AS
  DECLARE N INT = 100000;
BEGIN
  WHILE (N > 0) DO
  BEGIN
    INSERT INTO TTT VALUES (CURRENT_TRANSACTION, CURRENT_CONNECTION, CURRENT_TRANSACTION);
    N = N - 1;
  END
END

-- Inserts via prepared dynamic statement
-- using named input parameters:

EXECUTE BLOCK AS
  DECLARE S VARCHAR(255);
  DECLARE N INT = 100000;
BEGIN
  S = 'INSERT INTO TTT VALUES (:a, :b, :a)';

  WHILE (N > 0) DO
  BEGIN
    EXECUTE STATEMENT (:S)
      (a := CURRENT_TRANSACTION, b := CURRENT_CONNECTION)
      WITH COMMON TRANSACTION;
    N = N - 1;
  END
END

-- Inserts via prepared dynamic statement
-- using unnamed input parameters:

EXECUTE BLOCK AS
  DECLARE S VARCHAR(255);
  DECLARE N INT = 100000;
BEGIN
  S = 'INSERT INTO TTT VALUES (?, ?, ?)';

  WHILE (N > 0) DO
  BEGIN
    EXECUTE STATEMENT (:S) (CURRENT_TRANSACTION, CURRENT_CONNECTION, CURRENT_TRANSACTION);
    N = N - 1;
  END
END

```


Other PSQL Improvements

Improvements made to existing PSQL syntax include the following:

Subqueries as PSQL Expressions

A. dos Santos Fernandes

Tracker reference [CORE-2580](#)

Previously, a subquery used as a PSQL expression would return an exception, even though it was logically valid in SQL terms. For example, the following constructions would all return errors:

```
var = (select ... from ...);  
if ((select ... from ...) = 1) then  
if (1 = any (select ... from ...)) then  
if (1 in (select ... from ...)) then
```

Now, such potentially valid expressions are allowed, removing the need to jump through hoops to fetch the output of a scalar subquery into an intermediate variable using `SELECT...INTO`.

SQLSTATE as Context Variable

D. Yemanov

Tracker reference [CORE-2890](#)

(v.2.5.1) `SQLSTATE` is made available as a PSQL context variable, to be used with `WHEN` in an exception block, like `GDSCODE` and `SQLCODE`.

Chapter 12

International Language Support (INTL)

Adriano dos Santos Fernandes

Some improvements appear in this release to tighten and enhance Firebird's handling capabilities for international language environment.

Default COLLATION Attribute for a Database

Databases of ODS 11.2 and higher can now optionally be created with a default collation associated with the default character set. For details, please see [Default COLLATION Attribute for a Database](#) in the DDL chapter.

ALTER CHARACTER SET Command

DDL syntax has been introduced to enable the default collation for a character set to be set at database level. For details, please see [ALTER CHARACTER SET Command](#) in the DDL chapter.

Connection Strings & Character Sets

Capability has been implemented in the API database connection (DPB) area to interoperate with the character set and/or code page of server and client, to avoid the previous problems that could occur when file names contained non-ASCII characters.

Refer to the topic [Connection Strings & Character Sets](#) in the chapter *Changes to the Firebird API and ODS*. Even if you are not normally interested in the API, this topic will be a worthwhile read if you have been bothered with such issues.

Other Improvements

Introducer Syntax Usage

The usage of introducer syntax, i.e., prefixing an underscore to a character set name, to force the succeeding text literal to transliterated to that character set, has caused some problems in situations where a single SQL statement

entails usage of more than one character set. The actual problems differ from version to version, showing up as transliteration errors, malformed string errors or just as some kind of unexpected behaviour.

Problems could occur in two different usage scenarios.-

1. One query is employing the introducer syntax when another query performs a select from MON\$STATEMENTS
2. Introducer syntax was used in a PSQL module

To enable a workaround for such problems, it is now possible to transform the literal string into the hex representation of the ASCII characters being submitted by the introducer. For example:

```
select _dos850 '123áé456' from rdb$database
```

may be transformed to

```
select _dos850 X'313233A082343536' from rdb$database
```

Malformed UNICODE_FSS Characters Disallowed

Tracker reference [CORE-1600](#).

Malformed characters are no longer allowed in data for UNICODE_FSS columns.

Repair Switches for Malformed Strings

New restore switches were added to the *gbak* utility code for the purpose of repairing malformed UNICODE_FSS data and metadata by restoring a backup of the affected database. Details are in the *gbak* section of the Utilities chapter.

Numeric Sort Attributes

Tracker reference: [CORE-1945](#))

For UNICODE collations only, a custom attribute NUMERIC-SORT has been enabled for specifying the order by which to sort numerals.

Format & Usage

NUMERIC-SORT={0 | 1}

The default, 0, sorts numerals in alphabetical order. For example:

1
10
100
2
20

1 sorts numerals in numerical order. For example:

1
2
10
20
100

Example

```
create collation unicode_num for utf8
from unicode 'NUMERIC-SORT=1';
```

Character Sets and Collations

UNICODE_CI_AI

Tracker reference [CORE-824](#).

UNICODE_CI_AI: case-insensitive, accent-insensitive collation added for UTF8.

WIN_1258

Tracker reference [CORE-2185](#).

Added alias WIN_1258 for WIN1258 character set, for consistency with other WIN* character sets.

SJIS and EUCJ Character Sets

Tracker reference [CORE-2103](#).

Strings in SJIS and EUCJ character sets are now verified for well-formedness.

Character set GB18030

Tracker reference [CORE-2636](#).

GB18030 is a Chinese national standard describing the required language and character support necessary for software in China. It has been activated from ICU.

Chapter 13

Command-line Utilities

Incompatibility with Older Clients

To enable the 32-bit tools to work correctly with new structures that enable statistics routines to work properly with a 64-bit server, it was necessary to introduce some new internal API functions (**struct perf64** and **perf64_xxx**) and change *isql* and *qli* to use them. This means that the *isql* and *qli* programs in V.2.5 are not compatible with older Firebird clients.

For more information, see the topic [Statistics Now Work Properly with 64-bit Values](#) in the Engine chapter.

fbtracemgr

Vlad Khorsun

Tracker reference [CORE-2524](#).

This is a new CLI utility for interfacing with the [new trace facilities](#). Usage is as follows.-

```
fbtracemgr <action> [<parameters>]
```

Action Switches

```
-STA[RT]      Start trace session
-STO[P]       Stop trace session
-SU[SPEND]    Suspend trace session
-R[ESUME]     Resume trace session
-L[IST]       List existing trace sessions
```

Parameters

Action parameters

```
-N[AME]       <string>  Session name
-I[ID]        <number>  Session ID
-C[ONFIG]     <string>  Trace configuration file name
```

Connection parameters

```
-SE[RVICE]    <string>  Service name
```

```
-U[SER]      <string>  User name
-P[ASSWORD] <string>  Password
-FE[TCH]    <string>  Fetch password from file
-T[RUSTED]  <string>  Force trusted authentication
```

Examples using fbtracemgr

```
fbtracemgr -SE remote_host:service_mgr -USER SYSDBA -PASS masterkey -LIST
fbtracemgr -SE service_mgr -START -NAME my_trace -CONFIG my_cfg.txt
fbtracemgr -SE service_mgr -SUSPEND -ID 2
fbtracemgr -SE service_mgr -RESUME -ID 2
fbtracemgr -SE service_mgr -STOP -ID 4
```

Notes

1. All switches and parameter identifiers are case-insensitive
2. To stop an interactive trace session on any platform, press Ctrl-C.

V.2.5.1 Improvements

The diagnostics of internal trace errors were improved. (Tracker reference [CORE-3413](#)).

Output is now flushed periodically. (Tracker reference [CORE-3324](#)).

Retrieve Password from a File or Prompt

Alex Peshkov

Any command-line utility that takes a **-password** parameter is vulnerable to password sniffing, especially when the utility is run from a script. Since v.2.1, the [PASSWORD] argument has displayed in the process list on POSIX platforms as an asterisk (*), which was an improvement on showing it in clear.

As a second stage towards hiding the password from unauthorised eyes, this release enables it to be retrieved from a file or (on POSIX) from STDIN.

New -fetch_password Switch

Firebird 2.5 introduces the new switch *-fet[ch_password]* as an optional replacement for *-pa[ssword]* for all command-line utilities that take a password for authentication purposes. The switch may be progressively abbreviated from the right, conforming to the established rules.

PLEASE NOTE

1. The exception to the rules is the *qli* utility, for which only *-F* is valid.
2. The new switch *cannot* be applied to substitute for the *-pw* switch of the *gsec* utility.

Usage of `-fetch_password`

The switch requires one parameter, an unquoted string that is the file path for the file containing the password. If the call is not made by a system user with Superuser/Administrator privileges, the location must be accessible by the system user making the call.

For example,

```
isql -user sysdba -fet passfile server:employee
```

extracts the first line of from a file named “passfile” in the current working directory and loads it into the [PASSWORD] argument of the call.

The filename can be specified as *stdin*:

```
isql -user sysdba -fet stdin server:employee
```

If *stdin* is the terminal, a prompt is presented—

```
Enter password:
```

—requiring the operator to type in the password.

Tip

On POSIX, the operator will also be prompted if s/he specifies

```
-fetch /dev/tty
```

This technique could be useful if, for example, you needed to restore from *stdin* (all one line):

```
bunzip2 -c emp.fbk.bz2 | gbak -c stdin /db/new.fdb  
-user sysdba -fetch /dev/tty
```

gsec

The following improvements have been added for *gsec*:

Mapping Switch for Windows Administrators

Alex Peshkov

Since v.2.1, Windows domain administrators have had full access to the user management functions. In v.2.5 they do not get these privileges automatically unless the SYSDBA has configured the security database to make it happen automatically.

In the *Administrative Features* chapter is a [detailed overview](#) of the new system role RDB\$ADMIN. There, you will find descriptions of the new ALTER ROLE syntax that can be used by the SYSDBA to enable or disable the automatic mapping of Windows administrators to the RDB\$ADMIN role in databases, including the security database which they access when creating, altering and dropping users.

This automatic mapping can also be done in a *gsec* command-line call, using the new **-mapping** switch.

Mapping an OS Administrator to the RDB\$ADMIN Role

The new **-mapping** switch is used to enable or disable the association of an operating system user with the RDB\$ADMIN role in the security database. It takes one argument: either *set* to enable the association or *drop* to disable it. The syntax is:

```
-mapping {set | drop}
```

Granting the RDB\$ADMIN Role to a Firebird User

The introduction of the RDB\$ADMIN system role has made it possible to escalate the privileges of an ordinary user. However, it was (and still is) not possible any for any user, even SYSDBA, to attach directly to the security database and grant the required permissions for the user to manage other users. A parameter—GRANT ADMIN ROLE—was included in the new CREATE USER and ALTER USER statement syntaxes to enable SYSDBA, or another user that has already acquired the RDB\$ADMIN role in the security database, to have the RDB\$ADMIN role applied to an ordinary user “at arm’s length”, as it were.

The same can be achieved in *gsec* using the new switch **-admin**. It takes one argument: either YES (to grant the RDB\$ADMIN role to the specified user in *security2.fdb*) or NO (to revoke it). The syntax is:

```
-admin {YES | NO}
```

Command-line Help for gsec

Claudio Valderrama

Tracker reference: [CORE-756](#)

Parameter help has been implemented for *gsec*, accessible by using the **-help** or **-?** switches.

fbsvcmgr

Additions made to *gsec* and the Service Parameter Block (SPB) relating to the system role RDB\$ADMIN are covered by the appropriate support in the *fbsvcmgr* utility.

- for *gsec* **-mapping**, two new tag items: **isc_action_svc_set_mapping** and **isc_action_svc_drop_mapping**
- for *gsec* **-admin**: the new parameter **isc_spb_sec_admin**, of *spb_long* with either 0 as its value (meaning REVOKE ADMIN ROLE) or non-zero (meaning GRANT ADMIN ROLE).

For a full overview of the RDB\$ADMIN role, refer to the topic [New RDB\\$ADMIN System Role](#) in the *Administrative Features* chapter.

gbak

Repair Switches for Malformed Strings

Adriano dos Santos Fernandes

Tracker reference [CORE-1831](#).

The *gbak* utility has two new restore switches intended to repair malformed UNICODE_FSS character data and metadata, respectively, when restoring the backup of an affected database.

Switch Syntax

```
-FIX_FSS_D(ATA) <charset> -- fix malformed UNICODE_FSS data  
-FIX_FSS_M(METADATA) <charset> -- fix malformed UNICODE_FSS metadata
```

Hints with Malformed String Exceptions for Restores

([CORE-2754](#), A. dos Santos Fernandes)

When a restore fails with malformed string errors, *gbak* will supply a hint in verbose output, referring the user to the **-FIX_FSS_METADATA** and **-FIX_FSS_DATA** switches.

Preserve Character Set Default Collation

Adriano dos Santos Fernandes

An improvement allows the current value of RDB\$DEFAULT_COLLATE_NAME in the system table RDB\$CHARACTER_SETS to survive the backup/restore cycle.

Tracker reference [CORE-789](#).

Improve Insertion Performance for Restore

Adriano dos Santos Fernandes

(v.2.5.1) Improved performance in the records insertion phase of *restore*.

Tracker reference [CORE-3433](#).

nBackup

New Switches Added

Four switches were added to NBackup:

-FE <filename> Fetch password from file
-Z Print version information
-? Help
-D ON|OFF Force direct I/O on or off

- -FE :: Supports the feature allowing the authentication password to be fetched from a file. For details, see [New -fetch_password Switch](#) in this chapter.
- -Z :: Prints the version details of the *nbackup* executable.
- -? :: Prints a terse list of usage instructions for the switches and options.
- -D :: ON option enables direct input and output operations; OFF disables them. The default settings depend on the operating system and Firebird version, as follows:

Windows service-capable platforms

ON in all versions.

POSIX

OFF in 2.0-2.0.5, 2.1-2.1.2, 2.1.4, 2.5 and higher

ON in 2.1.3 and 2.0.6, provided O_DIRECT is available; otherwise OFF. POSIX_FADV_NOREUSE is also set, if available.

Note

Using -D ON has no effect if neither O_DIRECT nor POSIX_FADV_NOREUSE is available, although no error or warning will be raised.

Note

Support for the new -D switch is also included among changes to the Services API in this version. For details, see [isc_spb_nbk_direct on/off](#) in the chapter *Changes to the Firebird API and ODS*.

I/O Resource Load Reduced on POSIX

An improvement has been done for POSIX versions to address a problem whereby the full backup tool of the *nBackup* incremental backup utility would hog I/O resources when backing up large databases, bringing production work to a standstill. Now, nBackup tries to read from the operating system cache before attempting to read from disk, thus reducing the I/O load substantially.

Note

The “cost” may be a 10 to 15 percent increase in the time taken to complete the full backup under high-load conditions.

Tracker reference [CORE-2316](#).

isql

Some changes to the *isql* interactive query tool have been implemented.

SQLSTATE instead of SQLCODE

Claudio Valderrama

isql now returns the SQLSTATE completion code in diagnostics, instead of the now deprecated SQLCODE. For more information, see the topic [Support for SQLSTATE Completion Codes](#) in the chapter *Changes to the Firebird API and ODS*.

Improvement for Exponential Number Output

Claudio Valderrama

Tracker reference [CORE-1171](#).

isql has always output different formatting of numbers on Windows and POSIX for two-digit exponents. The default behaviour of Microsoft and Intel compilers is to zero-pad the exponent to three digits regardless. For example,

```
select cast ('-2.488355210669293e+39' as double precision)
  from rdb$database;
```

- On POSIX, the result is **-2.488355210669293e+39**
- On Windows, the result was **-2.488355210669293e+039**

The *isql* output has been modified so that the Windows output now conforms with that on other platforms.

SHOW COLLATIONS in isql Help

Help for SHOW COLLATIONS has been added to the command-line help for *isql*. (Tracker reference [CORE-2432](#), A. dos Santos Fernandes).

SET ROWCOUNT Statement Added

Mark O'Donohue

Added mainly to assist testers, this command enables a limit to be set on the number of rows a query returns to the interactive *isql* shell.

Example of Use

The following *isql* statement will stop returning rows after the 10,000th row has been output:

```
SQL>set rowcount 10000;
```

gpre (Precompiler)

Some Updates

Stephen Boyd
Adriano dos Santos Fernandes

Tracker reference [CORE-1527](#).

GPRES now supports the IS NOT DISTINCT predicate and CASE/NULLIF/COALESCE/SUBSTRING functions, as well as the whole set of CURRENT_* context variables.

Deprecated Features with Future Impact on Utilities

In anticipation of the dropping of the intrinsic function `isc_ddl` from the Firebird 3 codebase, certain features currently available in the `gdef` and `gpre` tools are deprecated—meaning that, whilst they may work in V.2.5, they will fail in Firebird 3. More details can be found in the [Compatibility chapter](#).

gstat

Claudio Valderrama

Tracker reference [CORE-1411](#).

The `gstat` statistics reporting utility now has `-?` and `-help` switches for requesting help about available switches and arguments.

Installation Notes

Installation And Migration Guide

The latest version of Installation and Migration Guide for Firebird versions 2.0.x and 2.1.x is still applicable to the v.2.5 series. If a copy of this document is not present in your /doc/ directory, you can download it from the Firebird website.

Some improvements have been done to solve issues that could arise with the binary installation packages.

Linux (POSIX)

The following improvements apply to POSIX installations.

Installation Scripts Cleanup

Alex Peshkov

- ([CORE-2195](#)): the Linux Classic installation scripts were reviewed to improve the assignment of ownership and access rights to documentation and other components.
- ([CORE-2392](#)): Cleanups of installation scripts for all active POSIX ports for Superserver and Superclassic were done to address a problem with the Guardian on these platforms.
- ([CORE-2626](#)): The startup scripts in `/etc/init.d` did not previously take into account the presence of the directories `/var/run/firebird` and `/tmp/firebird` on the host system. Various types of startup failures were traceable back to this deficiency.

The startup scripts distributed with the V.2.5 and later builds address and solve the problem.

- (**v.2.5.1** [CORE-3467](#)): A silent install switch (`-silent`) was added to `make install` to enable a simpler installation and the ability to automate a build and tests run. With this switch, Firebird is installed without any prompts, generating a random password for SYSDBA. The random password is saved in a file at the same location as the security database, `security2.fdb`.

Dedicated Firebird Switches for 'configure'

Alex Peshkov

A lot of the standard **configure** switches for fine-tuning the installation directories on POSIX platforms do not work for Firebird.

It was close to impossible to make the standard GNU switches work without changing the defaults for them, a rigmarole that is far from obvious or easy. Instead, a set of new switches for the **configure** has been added to enable fine-level configuration of the locations of Firebird's files.

The **ib_util** loader was also improved to make the engine work correctly with the configured layout.

Available Switches

```
--with-fbbin executables DIR (PREFIX/bin)
--with-fbsbin system admin executables DIR (PREFIX/bin)
--with-fbconf config files DIR (PREFIX)
--with-fbllib object code libraries DIR (PREFIX/lib)
--with-fbinclude C/C++ header files DIR (PREFIX/include)
--with-fbdoc documentation root DIR (PREFIX/doc)
--with-fbudf UDF DIR (PREFIX/UDF)
--with-fbsample examples DIR (PREFIX/examples)
--with-fbsample-db examples database DIR (PREFIX/examples/empbuild)
--with-fbhelp QLI help DIR (PREFIX/help)
--with-fbintl international DIR (PREFIX/intl)
--with-fbmisc misc DIR (PREFIX/misc)
--with-fbsecure-db security database DIR (PREFIX)
--with-fbmsg message files DIR (PREFIX)
--with-fblog log files DIR (PREFIX)
--with-fbglock guardian lock DIR (PREFIX)
--with-fbplugins plugins DIR (PREFIX)
```

Detection of Path to Firebird's Binaries

Adriano dos Santos Fernandes

Tracker reference: [CORE-2398](#))

A feature has been implemented for POSIX builds, to have Firebird correctly detect the path where its binaries are installed.

Important

This is an experimental state currently and is disabled in the distributed binaries. To activate it, build Firebird from source and pass **--enable-binreloc** to `autogen.sh` when building.

Windows

Deployment Structure for Embedded

Adriano dos Santos Fernandes

Vlad Khorsun

Tracker entry: [CORE-1814](#)

In this release it is possible to be a little more flexible about the location of your embedded application components than previously. Changes made here addressed a handful of interrelated issues that had been problematic for application developers, viz.

- An external function library that depended on another DLL, such as `ib_util.dll` or the client code in `fbembed.dll` or some totally external library, could not be just placed in the `..\UDF` folder beneath the emulated Firebird root (`<root>`) because the depended-on files should not be in that location. It was necessary to be somewhat creative with the `%PATH%` variable.
- The previous rules for determining the emulated `<root>` required `fbembed.dll` (renamed or not) to be in the same folder as the application code. One effect of that was to make it necessary either to keep separate copies of `fbembed.dll` for each application that needed to use it, or to place all of the client executables in one location.

Determination of the `<root>` for the embedded v.2.5 engine has been changed so that it no longer has to be the application's path. It is by the location of `fbembed.dll`, wherever it might be located. This means that, with the v.2.5 engine, you can set up a standard, self-contained structure for the emulated Firebird `<root>` that is available to any of your embedded applications, to third-party UDF libraries and to local copies of the command-line tools, once they have loaded `fbembed.dll`.

If you separate the emulated Firebird tree from your application[s], be sure to set the *RootDirectory* parameter in the structure's `firebird.conf` file to point to the absolute path location that is its `<root>`.

Note

The new rules do not break your existing embedded structures. You can still structure your embedded applications just as you did previously. The difference now is that it is not a requirement any more to deploy your executable with its own dedicated copy of an emulated Firebird `<root>` structure.

Managing MSVC8 Assemblies

Vlad Khorsun

Tracker entry: [CORE-2243](#)

Note

Because the changes took effect from V.2.1.2, this discussion also appears as a special topic in the *V.2 Installation and Migration document*.

Firebird 2.5 is built by the Microsoft MSVC8 compiler in Visual Studio 2005. Because all the Firebird binaries are built to use dynamic linking, they all require run-time libraries.

To avoid the dll-hell issue Microsoft introduced new rules for the distribution of components that may be shared by multiple applications. From Windows XP forward, shared libraries—such as the Visual C++ and Visual C runtimes `msvcp80.dll`, `msvcr80.dll` and `msvcvm80.dll`—must be distributed as shared or as private assemblies.

- The Microsoft MSI Installer installs shared assemblies into the common special folder SxS for use by multiple applications.
- Private assemblies are distributed with applications and should be put into the application folder. Use of the `\system32` folder for assemblies is now prohibited on the XP, Server2003 and Vista platform families.

Installing Runtimes as a Shared Assembly

To install the runtimes as a shared assembly, the deployment system must have MSI 3.0 installed and the user must have administrative privileges. Often, this is not possible with an application being deployed with Firebird Embedded: it must be installed ready-to-run. In that case, do not plan to install the runtimes as a shared assembly.

Installing Runtimes as a Private Assembly

To install the MSVC8 run-time libraries as a private assembly its contents—the three DLLs mentioned above and the assembly's manifest file, `Microsoft_VC80_CRT.manifest`—must be put into every folder where a dependent binary (.exe or .dll) resides, because of built-in checks for the folders that are the expected location of the runtimes that are equivalent to the compile-time libraries that were used.

A typical installation of Firebird Embedded would thus require three complete copies of the MSVC8 run-time assembly: one in the application folder and one each into the `\intl` and `\udf` folders. To avoid the issue of bloating the installation, some changes were done for V.2.1.2 in the way some of the Firebird binaries are built. (See also [Tracker entry CORE-2243](#)).

These are the changes that enable Firebird Embedded to work even if the application structure does not incorporate the MSVC8 runtime assembly:

- a. The libraries `ib_util.dll`, `fbudf.dll`, `ib_udf.dll`, `fbintl.dll` are built without any embedded manifest. The effect is to avoid having the loader search for a MSVC8 assembly in the same folder as corresponding DLL. For this to work, the host process must have already loaded the MSVC8 run-time via manifest before any attempt is made to load these secondary DLL's.
- b. `fbembed.dll` now has code to create and activate the activation context from its own manifest before loading any secondary DLL that might be required.

Notes

- a. It is highly recommended to use the Microsoft redistribution package to install the MSVC8 run-time! The executable installer `vcredist_x86.exe` or `vcredist_x64.exe` (as appropriate to your kit selection) should be present in the zip kits for the full installation and the Embedded version. If not, it can be downloaded from the [Microsoft download site](#).
- b. Third party UDFs must satisfy *one of the following requirements* if a MSVC8 run-time assembly is installed as private assembly. When compiling the UDF library, the MSVC8 runtime *EITHER*:
 - is NOT used
 - is used but the build is done without the embedded manifest
 - is used and the build is done with the embedded manifest—the default option in the MSVC IDE. In this case the MSVC8 assembly must be in the same folder as the UDF library

Chapter 15

Compatibility Issues

Dmitry Yemanov

For migrating v.2.0.x or v.2.1.x databases to Firebird 2.5, a number of incompatibilities that are likely to affect existing databases or existing applications should be noted. It is not recommended that you begin a migration until you have resolved these.

Incompatibility with Older Clients

To enable the 32-bit tools to work correctly with new structures that enable statistics routines to work properly with a 64-bit server, it was necessary to introduce some new internal API functions (**struct perf64** and **perf64_xxx**) and change *isql* and *qli* to use them. This means that the *isql* and *qli* programs in V.2.5 are not compatible with older Firebird clients.

For more information, see the topic [Statistics Now Work Properly with 64-bit Values](#) in the Engine chapter.

Effects of Unicode Metadata

If you have not previously updated text objects within the metadata of your databases to be in character set UTF8, restoring a database until V.2.5 will fail with “malformed string” errors. To resolve this it is necessary to pay attention to the files in the `/misc/upgrade/metadata` directory of your installation and to use the new `-fix_fss_data` and `-fix_fss_metadata` switches in the *gbak* command line.

Configuration Parameters Removed

The deprecated configuration parameters `OldParameterOrdering` and `CreateInternalWindow` are no longer supported and have been removed from `firebird.conf`.

Two parameters that allowed tuning of the Lock Manager in previous versions are not required with the new Lock Manager implementation and have been removed. They are `LockSemCount` and `LockSignal`.

The parameter `MaxFileSystemCache` has been renamed to `FileSystemCacheThreshold`.

SQL Language Changes

It will be necessary to pay attention to some changes in the SQL language implementation.

Reserved Words

While some new reserved keywords are introduced, generally the overall list of reserved words has been reduced dramatically by making Firebird's parser grammar allow most of the previously reserved non-standard keywords as non-reserved. The list of those still reserved, and of SQL standard keywords newly reserved, is available in the chapter [Reserved Words and Changes](#).

Execution Results

Some changes will now cause exceptions during run-time execution of queries, including those that are run during the execution of the *gbak* utility code (backups and restores).

Malformed String Errors

Well-formedness checks are now performed on UNICODE_FSS strings and text blobs. If new or existing UNICODE_FSS is malformed, it will now cause exceptions at execution time.

Logic Change in SET Clause

Previously, when the SET clause of the UPDATE statement assigned new values to columns, the new value replaced the old value immediately. If the same column was assigned or assigned to more than once, the current value would be that of the assignment most recently done. In other words, previously, assignment order mattered.

To bring Firebird in line with the standard, from this version forward, only the original value of a column is accessible to any assignment in the SET clause.

For a period, it is possible to revert to the legacy behavior by setting the temporary parameter [OldSetClauseSemantics](#) in `firebird.conf`. This parameter will be deprecated and removed in future releases.

Utilities

Be on the watch for the effects of the following changes to the Firebird command-line utilities.

fb_lock_print

Because v.2.5 maintains separate lock structures for each database on the server, *fb_lock_print* now requires a database path name in order to print the lock table. Include the new switch `-d <path name>` in the command line to specify *the filesystem path* to the database you wish to analyse.

Deprecated Features with Future Impact

In anticipation of the dropping of the intrinsic function **isc_ddl** from the Firebird 3 codebase, certain features currently available in the *gdef* and *gpre* tools are deprecated—meaning that, whilst they may work in V.2.5, they will fail in Firebird 3.

- *gdef* will no longer be supported at all. Instead, *isql* should be used, with regular DDL commands.
- For *gpre* pre-processing, replace all DDL operations with

```
EXEC SQL
EXECUTE IMMEDIATE "..."
```

- In all custom applications, calls to **isc_ddl** must be replaced with SQL DDL statement requests.

API Changes

Notice the following changes to the application programming interface (API) that is implemented in the client libraries.

Rejection of Inconsistent TPB Options

The API functions **isc_start_transaction()** and **isc_start_multiple()** will now reject combinations of transaction parameter buffer (TPB) items that do not “belong together”.

For example, a non-zero *wait timeout* is inconsistent with the *no wait* option; and *no record version* is inconsistent with any *transaction isolation mode* other than ReadCommitted. Now, instead of making some arbitrary (and possibly incorrect) assumption about the inherent ambiguities, the engine will reject such combinations as invalid.

For more information, see the topic [Transaction Diagnostics](#) in the chapter *Changes in the Firebird Engine*.

Addition of SQL_NULL Constant

New SQL_NULL constant was introduced to enable the predication **OR ? IS NULL** to be recognised and processed with the expected outcome and without engendering the “Data type unknown” exception. This affects how the XSQLVAR structures are populated for such queries. For information, refer to the topic [SOME_COL = ? OR ? IS NULL Predication](#) in the DML chapter.

Security Hardening

The following changes should be noted.

No SYSDBA Auto-mapping (Windows)

In V.2.1, members of administrative Windows groups were mapped to SYSDBA by default. From V.2.5 forward, SYSDBA mapping is controlled on per-database basis using the new SQL command

```
ALTER ROLE RDB$ADMIN SET/DROP AUTO ADMIN MAPPING
```

For more details, refer to [the chapter on Security](#).

Default Authentication Method (Windows)

In V.2.1, where support for Windows trusted authentication was introduced, the default authentication method applied was *mixed*, i.e., the DPB or SPB would accept either native Firebird logins or trusted user logins. Thus, the *Authentication* parameter in `firebird.conf` showed *mixed* as the default.

From V.2.5 forward, the default is *native*. To have *mixed* or *trusted*, it is now necessary to configure this parameter specifically.

Tracker reference [CORE-2376](#))

Access Path for Services

In some previous Firebird versions, if a full server path was supplied as the database name argument for remote access to a service, it would appear to be valid if the SYSDBA password was the same on both the client and the server. Under these conditions, no exception would be thrown and access would succeed.

For example, the following syntax for a remote call to *gbak* would appear valid if the remote client was running a Firebird server that had the same SYSDBA password as the remote server:

```
gbak -b -se dbhost:service_mgr dbhost:dbalias  
/var2/backups/mydb.fbk -user SYSDBA -password masterke
```

This anomaly was a bug that, unfortunately, some developers have long regarded as an undocumented feature. In v.2.5, this syntax will except under any condition. The correct syntax should be:

```
gbak -b -se dbhost:service_mgr dbalias  
/var2/backups/mydb.fbk -user SYSDBA -password masterke
```

```
gbak -b -se dbhost:service_mgr d:\databases\mydb.fdb  
x:\backups\mydb.fbk -user SYSDBA -password masterke
```

Known Platform Issues

Engine and other changes may have some platform-specific effects other than those noted generally. Known issues are listed below.

MacOSX

To make the new engine work properly under multi-threaded conditions we had to use [Grand Central Dispatch](#), which was first released in MacOSX 10.6 (Snow Leopard). Thus, MacOSX users should be aware that Firebird 2.5 will run only on MacOSX 10.6 or higher versions.

Note

If you want to use an earlier version of OSX you will need to use earlier versions of Firebird.

Chapter 16

Platform Ports

In this chapter are topics about ports of Firebird to non-mainstream platforms, along with notes about changes and improvements to ports that have been done previously.

HPPA

Linux/HPPA

D. Ivanov
A. Peshkov

Port of Firebird 2.5.1 to Linux on HPPA.

Tracker reference [CORE-3184](#)

Linux/Alpha

D. Ivanov
A. Peshkov

Port of Firebird 2.5.1 to Linux on Alpha.

Tracker reference [CORE-3184](#)

IBM eServer z-Series

Linux/s390 (32-bit)

D. Ivanov
A. Peshkov

Tracker reference [CORE-2625](#)

Port for Linux/s390 (32-bit) platform, built on s390x architecture.

The patch for this port removes the `-DS390X` define from `prefix.linux_s390x` and replaces it with checks for `__s390__` and `__s390x__` defines which are presented by `gcc`. This way, both ports can use the one prefix file.

Note

s390 has no alignment restrictions.

Linux/s390x (64-bit)

D. Horak
A. Peshkov

Tracker reference [CORE-2559](#)

Port for Linux/s390x (64-bit) platform.

Linux/sh4 (Renesas SH)

N. Iwamatsu
D. Ivanov
A. Peshkov

Tracker reference [CORE-2655](#)

Port for Linux/sh4 (Renesas SH) platform

Linux/sh4 (Renesas SH) architecture comes in little- and big-endian variants. The port supports both. SH has alignment constraints.

HP-UX

Lock Table Improvement for HP-UX

A. Peshkov

Tracker reference [CORE-2644](#)

All of the POSIX platform ports except HP-UX enable the lock table to grow when the space initially allocated for it becomes exhausted. For the HP-UX port, the hardware limitations of the PA-RISC platform prevented the same implementation of dynamic resizing because it does not support remapping the same file region to a different virtual address in the same process. **ISC_remap_file** therefore does not work and the lock table could not grow.

The SAS team had developed a solution in their edition of the Vulcan code to make the lock table expandable on this platform. It is this solution that has been imported to the Firebird 2.5 port for HP-UX.

Port for Very Old Windows 32-bit Platforms

At the request of N. Samofatov

Tracker reference [CORE-2609](#)

It had been signalled that versions of Firebird beyond V. 2.1.x would not support Microsoft Windows 98, ME, or NT4. However, Red Soft, a commercial company which wants to support Win98/ME and an idiomatic version of NT4 used by and maintained for government departments in Russia, requested permission for the inclusion of some re-implemented compiler conditions, isolated in separate units, that would enable the building of “ports” of Firebird 2.5 to those old or idiomatic Windows platforms.

These ports are not part of the Firebird 2.5 main stream distribution!

- The Windows binaries distributed by the Firebird Project do not contain this support.
- No such ports are subjected to any QA testing by the project.
- Any QA, future maintenance and development or bug-fixing affecting such ports are the responsibility of those who use them.

Bugs Fixed

Firebird 2.5.1 Release

The following improvements and bug fixes were reported as fixed prior to the v.2.5.1 release:

Core Engine/API

([CORE-3537](#)) **Improvement** :: The “undoing” on transaction rollback of changes made to global temporary tables created with the ON COMMIT DELETE ROWS option was unnecessary and was removed.

implemented by V. Khorsun

~ ~ ~

([CORE-3536](#)) **Improvement** :: Garbage collection in global temporary tables was being delayed unnecessarily by active transactions in other attachments.

implemented by V. Khorsun

~ ~ ~

([CORE-3457](#)) **Improvement** :: The temporary space manager has been made more optimal with regard to small chunk allocations.

implemented by D. Yemanov

~ ~ ~

([CORE-3323](#)) **Improvement** :: Lock Manager was provided with the capability to cancel waiting.

Considering the following example:

```
tx1: update table t ... where id = 1
tx2: update table t ... where id = 1
```

If transaction tx2 is in WAIT mode, it will wait for the end of tx1 forever and this wait cannot be broken using either a **DELETE FROM MON\$xxx** or a **fb_cancel_operation** request.

The improvement provides Lock Manager with the ability to break such interminable waits.

implemented by V. Khorsun

~ ~ ~

([CORE-3295](#)) **Improvement ::** Estimate the actual record compression ratio in the optimizer, thus allowing a more precise guess about table cardinalities (number of stored records).

implemented by D. Yemanov

~ ~ ~

([CORE-3560](#)) The V.2.5 Classic Server was using more memory than V.2.1.5 when caching metadata.

fixed by A. Peshkov

~ ~ ~

([CORE-3549](#)) Database corruption could occur at the end of a session, throwing the error “page xxx is of wrong type expected 4 found 7”

fixed by V. Khorsun

~ ~ ~

([CORE-3547](#)) A floating-point negative zero was not the same as a positive zero in indexes.

fixed by D. Yemanov

~ ~ ~

([CORE-3535](#)) The write target of a dirty page could be undefined if an error occurred when the *nbackup* state was changed.

fixed by V. Khorsun

~ ~ ~

([CORE-3533](#)) On SuperServer, a memory leak would occur if a connection was terminated without explicitly releasing the handles for all the statements prepared and cached by the client application.

fixed by V. Khorsun

~ ~ ~

([CORE-3532](#)) The server would hang on starting a new session when a trace was running.

fixed by A. Peshkov

~ ~ ~

([CORE-3525](#)) Autonomous transactions were wrongly inheriting the run-time flags of the "parent" transaction.

fixed by V. Khorsun

~ ~ ~

([CORE-3515](#)) Index corruption could occur under a complex set of conditions that was causing index updates to get out of synch and lose entries. Validation would pick it up and write the "missing entries" message into the firebird.log

fixed by V. Khorsun

~ ~ ~

[\(CORE-3512\)](#)

fixed by A. Peshkov The server could hang when a trace was running.

~ ~ ~

[\(CORE-3509\)](#)

ALTER PROCEDURE was permitting the addition of a parameter with the same name as an existing one.

fixed by V. Khorsun

~ ~ ~

[\(CORE-3502\)](#)

DROP VIEW was ignoring the existing non-column dependencies.

fixed by V. Khorsun

~ ~ ~

[\(CORE-3494\)](#)

Attach would fail anyway, after a shutdown had been rejected by the handler installed in fb_shutdown_callback().

fixed by A. Peshkov

~ ~ ~

[\(CORE-3491\)](#)

Altering a TYPE OF COLUMN parameter in a PSQL module would affect the original column.

fixed by A. dos Santos Fernandes

~ ~ ~

[\(CORE-3461\)](#)

DDL operations would fail after a backup and restore.

fixed by A. Peshkov

~ ~ ~

[\(CORE-3443\)](#)

Race conditions could occur during UDF library lookup.

fixed by A. Peshkov

~ ~ ~

[\(CORE-3418\)](#)

Database trigger created as INACTIVE was coming up active.

fixed by V. Khorsun

~ ~ ~

[\(CORE-3398\)](#)

GRANT ADMIN ROLE was not being accepted.

fixed by A. Peshkov

~ ~ ~

([CORE-3397](#)) Unresolved symbols were found in intl and trace libraries.

fixed by A. Peshkov

~ ~ ~

([CORE-3394](#)) Failed attempt to violate a unique constraint could leave an unnecessary “lock conflict” error in the status vector.

fixed by V. Khorsun

~ ~ ~

([CORE-3341](#)) Events posted from inside an autonomous transaction could get lost and never be delivered to the listening client.

fixed by D. Yemanov

~ ~ ~

([CORE-3340](#)) In an autonomous transaction with an empty exception handler, duplicate values could be inserted into a primary of unique key column, leading to an unrestorable backup.

fixed by D. Yemanov

~ ~ ~

([CORE-3327](#)) The thread pool in a network server could create more threads than were necessary.

fixed by V. Khorsun

~ ~ ~

([CORE-3326](#)) A fast mutex could be left in locked state by a dead process.

fixed by V. Khorsun

~ ~ ~

([CORE-3325](#)) Under high load it was possible that a new process would fail to map shared memory.

fixed by V. Khorsun

~ ~ ~

([CORE-3315](#)) The audit plugin would record a second *FAILED* EXECUTE_STATEMENT_FINISH after a “normal” one.

fixed by V. Khorsun

~ ~ ~

([CORE-3314](#)) Dependencies were not being removed after a procedure and the table it depended on were dropped in the same transaction.

fixed by D. Yemanov

~ ~ ~

([CORE-3312](#)) The join plan was sub-optimal when the slave table depended on the master one by means of the OR predicate.

fixed by D. Yemanov

~ ~ ~

([CORE-3306](#)) An invariant sub-query was being treated as variant, causing multiple invocations of a nested stored procedure.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3283](#)) A bad plan was produced with a query using a LEFT OUTER JOIN in a sub-query.

fixed by D. Yemanov

~ ~ ~

([CORE-3282](#)) EXECUTE STATEMENT was using the wrong character set when parsing the SQL text argument.

fixed by V. Khorsun

~ ~ ~

([CORE-3266](#)) A race condition could occur between the asynchronous service detach request and a running user trace service.

fixed by V. Khorsun

~ ~ ~

([CORE-3256](#)) Error “request depth exceeded” could appear while preparing a select query against a view with explicit plan.

fixed by D. Yemanov

~ ~ ~

([CORE-3237](#)) Stored procedures were being compiled too slowly.

fixed by D. Yemanov

~ ~ ~

([CORE-3207](#)) AccessViolationException inside FB_965910463_Class.isc_start_multiple on beginning transaction.

fixed by D. Yemanov

~ ~ ~

([CORE-3205](#)) isc_dsqli_exec_immed2() was not returning error codes isc_stream_eof and isc_sing_select_err when they occurred.

fixed by D. Yemanov

~ ~ ~

([CORE-3188](#)) Error “Page 0 is of wrong type (expected 6, found 1)” was occurring unexpectedly.

fixed by V. Khorsun

~ ~ ~

([CORE-3176](#)) A view with a column derived from a sub-query would join the sub-query output to the table without using an index.

fixed by D. Yemanov

~ ~ ~

([CORE-3168](#)) The `exclude_filter` was not working for the `<services>` section of the trace facility.

fixed by V. Khorsun

~ ~ ~

([CORE-3157](#)) Adding a parameter description to a stored procedure could lead to heavy memory consumption and protracted execution time.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3151](#)) Memory allocated in `sqlda_sup` could be left unreleased in some cases.

fixed by A. Peshkov

~ ~ ~

([CORE-3148](#)) Dangerous code was discovered in `SQZ_apply_differences`.

fixed by D. Kovalenko, A. Peshkov

~ ~ ~

([CORE-3140](#)) Comments for parameters were not being preserved after altering procedures.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3137](#)) Partial rollback was possible for a selectable procedure that was modifying data.

fixed by D. Yemanov

~ ~ ~

([CORE-3131](#)) `WIN1257_LV` (Latvian) collation was wrong for the four letters A, E, I and U.

fixed by D. Yemanov

~ ~ ~

([CORE-3125](#)) An access violation was occurring in the routine Worker::shutdown().

fixed by V. Khorsun

~ ~ ~

([CORE-3058](#)) New generators were being created with wrong values when more than 32,767 generators had been created previously.

fixed by D. Yemanov

~ ~ ~

([CORE-3029](#)) Bugcheck “Too many savepoints (287)” at rollback after an exception from EXECUTE BLOCK that contained an exception handler.

fixed by D. Yemanov

~ ~ ~

([CORE-3024](#)) Error “no current record for fetch operation” would occur after ALTER VIEW.

fixed by A. Peshkov

~ ~ ~

([CORE-2835](#)) Natural was being used to select, when the primary key index should have been used.

fixed by D. Yemanov

~ ~ ~

([CORE-2827](#)) Prepare was very slow for complex interrelated metadata when the operation being prepared indirectly involved many triggers.

fixed by D. Yemanov

~ ~ ~

([CORE-2709](#)) Indexed reads in a compound index with NULLs were excessive.

fixed by D. Yemanov

~ ~ ~

([CORE-1752](#)) Results of a join with different collations exhibited variations according to the execution plan used.

fixed by D. Yemanov

~ ~ ~

([CORE-1274](#)) Results would be wrong when PLAN MERGE was chosen and the data types of the equality predicate arguments were different.

fixed by D. Yemanov

~ ~ ~

Server Crashes

([CORE-3554](#)) Passing an empty SQL query remotely could crash the server during the **prepare** or throw an incorrect parsing error.

fixed by D. Yemanov

~ ~ ~

([CORE-3557](#)) The server could crash while preparing a query against a table that was in the process of being dropped.

fixed by V. Khorsun

~ ~ ~

([CORE-3524](#)) The server would crash while recompiling a stored procedure that was in use.

fixed by V. Khorsun

~ ~ ~

([CORE-3503](#)) ALTER VIEW would crash the server if the new version had an artificial (aggregate or union) stream at the position of a regular context in the older version.

fixed by V. Khorsun

~ ~ ~

([CORE-3477](#)) The server would crash whenever non-existent SQL parameters were passed to it.

fixed by D. Yemanov

~ ~ ~

([CORE-3440](#)) The server would crash if isc_que_events() had queued 0 events.

fixed by V. Khorsun

~ ~ ~

([CORE-3247](#)) The server was crashing with BLOBs in character set UTF8.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3419](#)) Server could hang or crash when an autonomous transaction was rolled back.

fixed by V. Khorsun

~ ~ ~

([CORE-3400](#)) Server crashes were occurring frequently on FreeBSD8.2R.

fixed by A. Peshkov

~ ~ ~

([CORE-3374](#)) The server could crash or corrupt data if SELECT WITH LOCK was issued against records not in the latest format.

fixed by D. Yemanov

~ ~ ~

([CORE-3320](#)) A certain MERGE syntax was able to crash the server. Details are not published.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3255](#)) The server could crash using views with GROUP BY.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3219](#)) Trace manager would crash the server with DSQL_unprepare.

fixed by V. Khorsun

~ ~ ~

([CORE-3217](#)) The server would crash inside the Lock Manager when multiple connections were attaching or detaching simultaneously.

fixed by D. Yemanov

~ ~ ~

([CORE-3202](#)) execute_immediate API call family could crash the remote server.

fixed by D. Yemanov

~ ~ ~

([CORE-3180](#)) ALTER VIEW with non-matched columns in the declaration and selection would crash the server.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3138](#)) An internal error or a crash would occur when accessing any MON\$ table after altering its structure.

fixed by D. Yemanov

~ ~ ~

([CORE-3064](#)) Using both the identifier of a procedure and its alias in an explicit plan would crash the server.

fixed by D. Yemanov

~ ~ ~

Data Manipulation Language

([CORE-3523](#)) SIMILAR TO was giving false matches on descending ranges.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3493](#)) Adding a value to a timestamp earlier than '16.11.1858 00:00:01' would throw the error "value exceeds the range for valid timestamp".

fixed by D. Yemanov

~ ~ ~

([CORE-3489](#)) Blob transliteration sometimes failed to happen inside a union.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3479](#)) The ASCII_VAL() function would raise an error for empty strings instead of returning 0.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3355](#)) Comparison of DATE and TIMESTAMP would be wrong if an index was used.

fixed by D. Yemanov

~ ~ ~

([CORE-3353](#)) Predicate (blob_field LIKE ?) was describing the parameter as VARCHAR(30) rather than as BLOB.

fixed by D. Yemanov

~ ~ ~

([CORE-3335](#)) Internal wrapping was occurring for the multi-byte blob SUBSTRING function and its boundary arguments, causing wrong results.

fixed by D. Yemanov

~ ~ ~

([CORE-3311](#)) The error "data type unknown" would be thrown while preparing an UPDATE or DELETE statement with a parameterized ROWS clause.

fixed by D. Yemanov

~ ~ ~

([CORE-3302](#)) DISTINCT aggregates were returning wrong (duplicated) data.

fixed by D. Yemanov

~ ~ ~

([CORE-3277](#)) The RIGHT() function was giving a wrong result for varchars of character set UTF8.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3245](#)) A SUBSTRING() operation on a long text BLOB would return a BLOB containing only 32767 characters if the optional third argument was not present.

fixed by D. Yemanov

~ ~ ~

([CORE-3244](#)) The result of POSITION() was wrong for an empty string (") if the third argument was present.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3233](#)) LIKE, STARTING and CONTAINING would fail if the second operand was 32KB or more.

fixed by A. dos Santos Fernandes, D. Yemanov

~ ~ ~

([CORE-3228](#)) The RIGHT() function would fail with multi-byte text blobs more than 1024 bytes in length.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3227](#)) The ASCII_VAL() function would fail if the argument contained multi-byte characters anywhere.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3222](#)) A view with "WITH CHECK OPTION" did not like a TRIM() function call in the WHERE clause.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3211](#)) String truncation would occur when selecting from a view that contained a NOT IN condition.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3210](#)) Error "no current record for fetch operation" would occur unexpectedly in a SELECT query.

fixed by D. Yemanov

~ ~ ~

([CORE-3208](#)) Recursive queries were exhibiting significant memory leaks.

fixed by D. Yemanov

~ ~ ~

([CORE-3203](#)) UPDATE OR INSERT with RETURNING would cause an “Invalid Cursor” error.

fixed by D. Yemanov

~ ~ ~

([CORE-3173](#)) Selecting from a stored procedure that contained two Common Table Expressions, the second with a GROUP BY clause, as well as an inner join, would wrongly return nothing in the result.

fixed by D. Yemanov

~ ~ ~

([CORE-3164](#)) Parameterized requests involving BLOB fields would fail when the client was connected using character set UTF8.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3141](#)) The last column in a view was being returned as NULL, even when it should contain a value.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3091](#)) The built-in function POWER(X, Y) did not work when the X argument was negative and the Y value was a scaled numeric of scale 0.

fixed by A. dos Santos Fernandes

~ ~ ~

Command-line Utilities

gbak

([CORE-3236](#)) *gbak* would throw the “unavailable database” error when both the service manager switch and localhost:db were specified.

fixed by A. Peshkov

~ ~ ~

([CORE-3249](#)) If an already existing file was used as the output file for *gbak -b* and the size of the new backup file was smaller than the existing one, the backup file was not being truncated.

fixed by A. Peshkov

~ ~ ~

([CORE-3232](#)) A non-transportable backup created without the “-se service_mgr” switch would produce a backup file that was about 50 percent larger than the source database.

fixed by A. Peshkov

~ ~ ~

nBackup

([CORE-3521](#)) nBackup delta file contents were not being flushed to disk.

fixed by V. Khorsun

~ ~ ~

([CORE-3482](#)) nBackup on Linux would segfault on Ctrl-C and leave the database locked and thus the delta file continuing to grow.

fixed by A. Peshkov

~ ~ ~

([CORE-3297](#)) nBackup would exit with no information if firebird.conf was missing.

fixed by A. Peshkov

~ ~ ~

([CORE-3199](#)) On POSIX, nBackup would fail if the requesting user was not root or Owner, due to the O_NOATIME flag being open.

fixed by A. Peshkov

~ ~ ~

fbtracemgr

([CORE-3487](#)) The fbtracemgr utility would segfault sometimes when ended with Ctrl-C.

fixed by A. Peshkov

~ ~ ~

fb_lock_prt

([CORE-3454](#)) “fb_lock_print -c” would cause the server to hang.

fixed by A. Peshkov

~ ~ ~

gpre

([CORE-3486](#)) GPRE language modules could not be compiled with gcc 4.4.

fixed by A. Peshkov

~ ~ ~

([CORE-3022](#)) *gpre* was getting C++ compiler warnings with GCC 4.4.1.

fixed by D. Dodson

~ ~ ~

Database Monitoring/Administration

([CORE-2305](#)) **Improvement** :: Make MON\$STATEMENT_ID value constant among monitoring snapshots.

implemented by D. Yemanov

~ ~ ~

([CORE-3508](#)) MON\$DATABASE_NAME and MON\$ATTACHMENT_NAME fields were substituting question marks for non-ASCII characters regardless of the connection character set.

fixed by D. Yemanov

~ ~ ~

([CORE-3218](#)) A statement cancellation request might be silently ignored by the currently running SQL code.

fixed by D. Yemanov

~ ~ ~

Services Manager

([CORE-3261](#)) Assertion when running restore service.

fixed by A. Peshkov

~ ~ ~

POSIX-Only Bugs

([CORE-3589](#)) On MacOSX and FreeBSD, semaphores were not detached from a shared file when the shared file was closed, causing an internal resource leak. The bug was effectively hidden, due to missing checks

of errors returned by the **ISC_event_init()** function. It showed up on MacOSX 10.7 as failure of Superserver and Superclassic to start.

fixed by A. Peshkov

~ ~ ~

([CORE-3544](#)) **make install** on Linux was failing.

fixed by A. Peshkov

~ ~ ~

([CORE-3447](#)) Some collations were not being installed on Linux with icu versions greater than 4.2.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3377](#)) During the Firebird build on POSIX, attempts were being made to place records about a missing fbintl.conf in firebird.log in the destination directory instead of the build directory.

fixed by A. Peshkov

~ ~ ~

([CORE-3259](#)) On POSIX, deadlock and SEGV would occur when processing Ctrl-C (terminate) in user code.

fixed by A. Peshkov

~ ~ ~

([CORE-3257](#)) 'make install' would fail on Linux.

fixed by A. Peshkov

~ ~ ~

([CORE-3229](#)) Something was causing error records to appear in the firebird.log “Operating system directive open failed, Too many links”.

fixed by A. Peshkov

~ ~ ~

([CORE-3212](#)) Compiling for FreeBSD was failing with an error.

fixed by A. Peshkov

~ ~ ~

([CORE-3194](#)) Number of connections to Superclassic on Linux was limited to 508.

fixed by A. Peshkov

~ ~ ~

([CORE-3185](#)) Firebird compilation by a non-root user on a POSIX box that was already running a Firebird server would encounter access conflicts in temporary lock space.

fixed by A. Peshkov

~ ~ ~

([CORE-3166](#)) The script *changeMultiConnectMode.sh*, for switching from Classic mode to Superclassic on POSIX, was wrongly included in the Superserver installation.

fixed by A. Peshkov

~ ~ ~

([CORE-3150](#)) On Linux, segmentation fault would occur in when *gbak* was interrupted with Ctrl-C.

fixed by A. Peshkov

~ ~ ~

([CORE-3143](#)) On Linux, segmentation fault could occur when a user interrupted *gstat*.

fixed by A. Peshkov

~ ~ ~

([CORE-2921](#)) 'make install' would not work in FreeBSD.

fixed by A. Peshkov

~ ~ ~

Windows-Only Bugs

([CORE-3329](#)) The Windows Administrator was acquiring the RDB\$ADMIN role unexpectedly.

fixed by A. Peshkov

~ ~ ~

([CORE-3059](#)) RemoteFileOpenAbility set True on a Windows server would fail.

fixed by D. Yemanov

~ ~ ~

Remote Interface/API

([CORE-3248](#)) **Improvement** :: Set unused bytes of varchar values in the message buffer to zero.

implemented by A. Peshkov

~ ~ ~

([CORE-2752](#)) **Improvement** :: Set the SO_KEEPALIVE option on the client TCP socket.

implemented by D. Yemanov

~ ~ ~

([CORE-3511](#)) Unquoted role names with non-ASCII characters were being upper-cased wrongly when passed in the database parameter buffer (DPB).

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3389](#)) isc_dsql_exec_immed2 with a zero transaction handle could lead to a BUGCHECK(147).

fixed by V. Khorsun

~ ~ ~

([CORE-3387](#)) The client library could hang indefinitely waiting for a reply packet on a forcibly disconnected server socket.

fixed by D. Yemanov

~ ~ ~

([CORE-3351](#)) The Windows client could put 10054 error messages into firebird.log at connection time.

fixed by V. Khorsun

~ ~ ~

([CORE-3328](#)) The client was writing “Unsuccessful detach from database” error messages into firebird.log when a database was shut down.

fixed by V. Khorsun

~ ~ ~

([CORE-3220](#)) API function isc_info_svc_get_users was returning error messages in the result cluster.

fixed by A. Peshkov

~ ~ ~

([CORE-3170](#)) The engine could enter an infinite loop if EVENTS were posted but no subscribers existed.

fixed by V. Khorsun

~ ~ ~

([CORE-3119](#)) The remote protocol code related to events processing was causing an endless loop and 100 per cent CPU usage.

fixed by V. Khorsun

~ ~ ~

([CORE-3095](#)) Client would receive an event count of one regardless of how many times in the same transaction the event was posted.

fixed by V. Khorsun

~ ~ ~

Firebird 2.5.0 Release

The following bugs were reported as fixed prior to the final 2.5.0 release:

Core Engine/API

([CORE-3115](#)) The internal record compression routines exhibited some bugs.

fixed by A. Peshkov, D. Kovalenko

~ ~ ~

([CORE-3103](#)) A SELECT statement incurred more non-indexed reads in the third v.2.5 release candidate than the same statement in v.2.1.3.

fixed by D. Yemanov

~ ~ ~

([CORE-3101](#)) ALTER DOMAIN was not possible in a database that had been migrated from an earlier version.

fixed by A. dos Santos Fernandes, D. Yemanov

~ ~ ~

([CORE-3100](#)) The WAIT mode and lock timeout parameters for the external transaction of EXECUTE STATEMENT were not matched to the corresponding parameters of the local transaction.

fixed by V. Khorsun

~ ~ ~

([CORE-3096](#)) A regression introduced with the fix for CORE-2893 caused double processing of nodes when preparing a statement, prompting an abort to occur in the debug build.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3094](#)) Parameters would not work with NOT IN from a selectable stored procedure.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3090](#)) Incorrect LEFT JOIN result using table and derived constant subquery.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-3089](#)) Attempt to run EXECUTE STATEMENT on an external data source failed when the data source was an InterBase 4.1 (ODS 8) database.

fixed by V. Khorsun

~ ~ ~

([CORE-3079](#)) Batch inserts were noticeably slowed down if they were all executed in a single transaction and involved triggers that posted events.

fixed by V. Khorsun

~ ~ ~

Server Crashes

([CORE-3109](#)) The server would crash when `isc_dql_exec_immed3_m()` was called for **CREATE DATABASE ...** in a NULL transaction.

fixed by A. dos Santos Fernandes

~ ~ ~

Command-line Utilities

gsec

([CORE-3116](#)) The `gsec` utility was sending its output list of users to `stderr` instead of `stdout`.

fixed by A. Peshkov

~ ~ ~

Old Bugs Fixed

The following bugs were reported as fixed during Firebird v.2.5 development:

Core Engine/API

([CORE-3034](#)) If a request was cancelled for some reason, such as database shutdown or by a user request, while inserting key into an expression index, it could lead to bugcheck 300 (can't find shared latch).

fixed by V. Khorsun

~ ~ ~

([CORE-3016](#)) On disconnect a “Fatal lock manager error: invalid lock id (0), errno: 0” could show up in `firebird.log`.

fixed by V. Khorsun

~ ~ ~

([CORE-3015](#)) Various “Cannot initialize the shared memory region” errors were reported.

fixed by V. Khorsun

~ ~ ~

([CORE-3003](#)) New checks for the existence of a SUSPEND statement in a procedure being called via SELECT statement were working properly: if the procedure was called via a SELECT statement and SUSPEND was not present, an error would be thrown, e.g., “Procedure ... is not selectable (it does not contain a SUSPEND statement)”.

However, if the same error condition occurred during a RESTORE, it could cause the restore to fail.

fixed by D. Yemanov

~ ~ ~

([CORE-2995](#)) A single error in the status vector was being reported twice.

fixed by V. Khorsun

~ ~ ~

([CORE-2993](#)) On a highly loaded system, the fatal lock manager error “Invalid lock id (NNN)” could occur while working with monitoring tables.

fixed by V. Khorsun, D. Yemanov

~ ~ ~

([CORE-2900](#)) Using a request containing an aggregated DISTINCT could cause regular but random memory access violations.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2977](#)) The server was working incorrectly with indexed fields of type DATE in databases with ODS lower than 10.

fixed by V. Khorsun

~ ~ ~

([CORE-2965](#)) The ROW_COUNT value returned after a subquery involving a SINGULAR predicate was inconsistent with the rules, viz., ROW_COUNT is meant to count rows affected by an update or delete and should not return non-zero value from subqueries.

fixed by D. Yemanov

~ ~ ~

([CORE-2956](#)) Problems with requests processing procedure parameters.

fixed by V. Khorsun

~ ~ ~

([CORE-2943](#)) An error would occur during parsing of a recursive query with two recursive parts.

fixed by V. Khorsun

~ ~ ~

([CORE-2936](#)) If two consecutive leaf index pages were removed (garbage collected) from an index simultaneously by two different connections, the linked list of sibling pages could be broken, with the sibling pointer at another index page left pointing to the freed index page. When the freed page was allocated again, index corruption would be reported as “Wrong page type (expected 7 found N)”.

fixed by V. Khorsun

~ ~ ~

([CORE-2916](#)) Error handling was broken in the case of a conversion error occurring during index creation.

fixed by D. Yemanov

~ ~ ~

([CORE-2893](#)) An expression within a subquery could be treated as invariant and produce incorrect results.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2879](#)) Sweep could raise the error **page 0 is of wrong type (expected 6, found 1)**.

fixed by V. Khorsun

~ ~ ~

([CORE-2876](#)) Unintelligent error handling when using **ALTER DATABASE ADD DIFFERENCE FILE** could cause the engine to get confused about the backup mode of a database.

fixed by A. Peshkov

~ ~ ~

([CORE-2875](#)) Comparing a CHAR column longer than 4096 bytes with a string constant would cause a string right truncation error.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2871](#)) If a derived table or a view contained both a left/right join and an ORDER BY clause and the outer query also contained an ORDER BY clause, the outer ORDER BY clause would have no effect.

fixed by D. Yemanov

~ ~ ~

([CORE-2858](#)) Memory trashing was possible when raising some exceptions to signal failed security checks.

fixed by C. Valderrama

~ ~ ~

([CORE-2856](#)) A non-NULL key in a unique index could not be found when the key was removed

fixed by V. Khorsun

~ ~ ~

([CORE-2833](#)) Changing data that affected an expression index would fail if that data contained references to null date fields.

fixed by D. Yemanov

~ ~ ~

([CORE-2826](#)) Join condition would fail for UTF-8 databases.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1089](#)) Selecting from a view that used DISTINCT and LEFT JOIN returned records in the wrong order if the ORDER BY clause did not include columns from the right-side (non-mandatory) table.

fixed by D. Yemanov

~ ~ ~

([CORE-195](#)) Regression of an old bug, previously fixed in v.1.5.1, whereby a bugcheck 291 (cannot find back record version) would occur when updating the same record that had already fired an action in a BEFORE UPDATE trigger. The regression that was reintroduced in v.2.0 was less destructive, insofar as it affected only the record that was physically first in the table.

fixed by A. Peshkov

~ ~ ~

([CORE-2822](#)) The error “no current row for fetch operation” was being thrown when a subquery included a non-trivial derived table.

fixed by D. Yemanov

~ ~ ~

([CORE-2820](#)) Queries with PLAN ORDER were exhibiting small memory leaks as a side effect of an earlier, major fix.

fixed by V. Khorsun

~ ~ ~

([CORE-2815](#)) Tidy-up of a logic condition that could cause the page inventory page to be marked for change after it was already changed.

fixed by V. Khorsun

~ ~ ~

([CORE-2785](#)) An improperly handled BLOB transliteration problem showed up under specific conditions, viz., where an input argument string in single-byte Cyrillic text, for storing as COMMENT with an object definition, overstepped the maximum 64 KB size of a non-initial BLOB segment when transliterated. Its surplus bytes were causing an immediate transliteration error instead of being deferred for inclusion in the subsequent chunk, as they should have been.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2783](#)) An access violation would occur when a recursive query was used as a subquery in the SELECT list and its result was specified as an ORDER BY criterion.

fixed by V. Khorsun

~ ~ ~

([CORE-2730](#)) A bus error could occur when working with DB_KEY on RISC machines: when casting from a literal (which has no alignment requirements) to db_key, alignment is necessary at the QWORD boundary and should have been enforced.

fixed by A. Peshkov

~ ~ ~

([CORE-2722](#)) Storage of a malformed blob was being allowed when copying from a blob with NONE/OCTETS character set.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2660](#)) COUNT(*) would incorrectly return 0 when a match was not found for an outer join condition. Now it correctly returns NULL.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2659](#)) The query plan for an outer join involving complex views could be sub-optimal, by not making use of an available index.

fixed by D. Yemanov

~ ~ ~

([CORE-2640](#)) Under some conditions, the lock manager could fail to detect a regular deadlock and cause the server to hang.

fixed by V. Khorsun, D. Yemanov

~ ~ ~

([CORE-2635](#)) A unique index could be corrupted at level 1 if it contained a lot of NULL keys.

fixed by V. Khorsun

~ ~ ~

([CORE-2632](#)) Unexpected “Invalid BLOB ID” errors could occur when working with the monitoring tables.

fixed by V. Khorsun

~ ~ ~

([CORE-2616](#)) Error “page <N> is of wrong type (expected 7, found 5)” could occur under load, giving the impression that something had corrupted the database. On restart, there would be no evidence of corruption.

fixed by V. Khorsun

~ ~ ~

([CORE-2608](#)) On latter versions of Windows (64-bit XP or later, 32-bit Vista or later), when Firebird is working with large databases, the operating system dedicates all RAM to the filesystem cache and stops responding. Windows may also crash.

It is a [documented issue](#). For Firebird, it has been addressed by implementing a new parameter, *FileSystemCacheSize* in *firebird.conf* to control the amount of RAM used for filesystem caching.

fixed by N. Samofatov

~ ~ ~

([CORE-2602](#)) Attachments using character set NONE could fail to read from the monitoring tables.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2591](#)) High mutex wait ratio and degraded performance would start to show up after a period of normal performance.

fixed by D. Yemanov

~ ~ ~

([CORE-2581](#) and [CORE-2582](#)) Infinity and NAN results from any expressions, including internal and external function calls, were not trapped by the engine as exceptions in some circumstances.

fixed by C. Valderrama

~ ~ ~

([CORE-2578](#)) Selecting RDB\$DBKEY from a view with more than one table joined would return a conversion error.

fixed by A. dos Santos Fernandes, A Peshkov

~ ~ ~

([CORE-2514](#)) An error concerning CreateFile was being reported when space on the 'temp' drive was insufficient.

fixed by D. Yemanov

~ ~ ~

([CORE-2422](#) and [CORE-2321](#)) The server was not switching between multiple locations configured in TempDirectories. Instead, the sorting would fail when the first configured temporary directory had insufficient free space. Typically, the user would see a sorted or other query that required a large amount of temp space fail with an error like 'operating system directive write failed. Invalid argument.'

fixed by D. Yemanov

~ ~ ~

([CORE-2315](#)) Firebird FLOAT support did not conform to the original InterBase specification. According to the IB documentation, a FLOAT value should have a range from 1.175E-38 to 3.402E38. Cross-platform tests proved that the largest value that would not overflow was < 3.4E38.

fixed by W. Oliver

~ ~ ~

([CORE-1991](#)) When a UDF was declared with BLOB parameters, the engine was storing null in RDB\$FUNCTION_PARAMETERS.RDB\$FIELD_LENGTH. This was causing a run-time 'message length error'.

fixed by D. Sibiryakov

~ ~ ~

([CORE-1781](#)) The engine would throw a consistency check error when a subquery was ordered by an aggregate function from another context.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2564](#)) Bus errors would be thrown when working with the monitoring tables on RISC machines.

fixed by A. Peshkov

~ ~ ~

([CORE-2550](#)) Bus errors would be thrown when working with DB_KEY on big-Endian machines.

fixed by A. Peshkov

~ ~ ~

([CORE-2538](#)) PSQL would not allow use of a singleton query result as an input parameter to a stored procedure if the procedure was called using EXECUTE PROCEDURE.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2532](#)) Volume sizes in a multi-volume database were being assigned incorrectly.

fixed by V. Khorsun

~ ~ ~

([CORE-2526](#)) The server could be shut down without regard to connections to services.

fixed by A. Peshkov, D. Yemanov

~ ~ ~

([CORE-2505](#)) Built-in trigonometric functions could produce NaN and Infinity.

fixed by C. Valderrama

~ ~ ~

([CORE-2501](#)) Binary shift functions would give wrong results with negative shift values.

fixed by C. Valderrama

~ ~ ~

([CORE-2499](#)) The implementation limit of DISTINCT items was not being enforced, causing incorrect BLR to be generated.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2482](#)) Monitoring tables data collection was unstable when the database was accepting attachments or detachments.

fixed by A. Peshkov

~ ~ ~

([CORE-2475](#)) External table data was not visible to Classic sessions other than the first to access the table.

fixed by V. Khorsun

~ ~ ~

([CORE-2449](#)) An unexpected “lock conflict” error could be thrown in lieu of the expected exception.

fixed by D. Yemanov

~ ~ ~

([CORE-2444](#)) The engine could hang when multiple attachments registered their interest in events simultaneously and free space in the events table became exhausted.

fixed by V. Khorsun

~ ~ ~

([CORE-2426](#)) ALTER TABLE was not respecting collations.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2416](#)) Preparing a query with aggregation over a derived table would cause an access violation.

fixed by V. Khorsun

~ ~ ~

([CORE-2411](#)) The optimizer would choose a slower PLAN for certain types of query than it would in versions 2.0.4 and 2.1.1. (The bug affected versions 2.0.5 and 2.1.2 as well.)

fixed by D. Yemanov

~ ~ ~

([CORE-2397](#)) Dropping more than one index on a table within the same transaction could cause corruption.

fixed by V. Khorsun

~ ~ ~

([CORE-2359](#)) Logical multibyte maximum string length was not respected when assigning numbers.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2331](#)) ALTER DOMAIN would result in an invalid RDB\$FIELD_SUB_TYPE value being stored.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2272](#)) The server would start returning garbage when killing an events connection attempt.

fixed by A. Peshkov

~ ~ ~

([CORE-1971](#)) Set the fixed and documented evaluation order (always left to right) for the WHERE clause and other predicates.

fixed by D. Yemanov

~ ~ ~

([CORE-1346](#)) LPAD() and RPAD() functions hit implementation restrictions when applied to more than one column in a statement.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2356](#)) On Windows the listener process of Classic Server was unable to create the necessary resources after restart if any worker process was present.

fixed by V. Khorsun

~ ~ ~

([CORE-2355](#)) Incorrect handling of LOWER/UPPER when result string shrinks in terms of byte length.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2351](#)) It was not possible to create a database whose <file specification> was an alias, even though the alias existed.

fixed by A. Peshkov

~ ~ ~

([CORE-2349](#)) The “Invalid SQLDA” error was being falsely thrown.

fixed by V. Khorsun

~ ~ ~

([CORE-2348](#)) More database corruption problems showed up resulting from transaction numbers overflowing 32-bit signed integer.

fixed by V. Khorsun

~ ~ ~

([CORE-2340](#)) Bugcheck 258 (page slot not empty) could occur under high concurrent load.

fixed by V. Khorsun

~ ~ ~

([CORE-2320](#)) A complex recursive query did not always return all rows.

fixed by V. Khorsun

~ ~ ~

([CORE-2313](#)) INF_* functions could invalidate the whole output buffer with isc_info_truncated at the beginning, due to a boundary condition.

fixed by C. Valderrama

~ ~ ~

([CORE-2311](#)) A WITH RECURSIVE query could cause memory leakage.

fixed by V. Khorsun

~ ~ ~

([CORE-2300](#)) The second evaluation of SUBSTRING() would throw an unexpected *arithmetic exception, numeric overflow, or string truncation* error.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2289](#)) The wrong name was being reported for the referenced primary key when a foreign key violation occurred during foreign key creation.

fixed by V. Khorsun

~ ~ ~

([CORE-2242](#)) The engine was incorrectly populating integer containers in the blob parameter buffer (BPB) with integers in machine-local format, causing errors on Big Endian platforms.

fixed by A. Peshkov

~ ~ ~

([CORE-2241](#)) If an ALTER TABLE ALTER COLUMN.. operation was performed on a table in the course of a bulk insert operation, minor index corruption could occur causing subsequent queries to return the wrong number of records. The bug was traced to legacy code in BTR\compress_root().

fixed by V. Khorsun

~ ~ ~

([CORE-2230](#)) Input parameters for EXECUTE BLOCK were not being domain-checked.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2186](#)) In the Windows embedded server, *fbintl.dll* was not being unloaded after the **isc_dsql_execute_immediate()** during the processing for CREATE DATABASE.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2182](#)) It was not possible to drop an existing UDF whose name was duplicated by the name of a new built-in function.

fixed by D. Yemanov

~ ~ ~

([CORE-2154](#)) A “request synchronization error” would occur when calling **isc_dsql_sql_info()** with the **isc_info_sql_records** parameter after the last record had been fetched with EXECUTE PROCEDURE.

fixed by V. Khorsun

~ ~ ~

([CORE-2132](#)) Indexed retrieval could not be chosen if a stored procedure call was used in the comparison predicate.

fixed by D. Yemanov

~ ~ ~

([CORE-2117](#)) Incorrect ROW_COUNT values were being returned with indexed retrieval and subqueries.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2115](#)) For some long queries the query plan could go missing.

fixed by D. Yemanov

~ ~ ~

([CORE-2101](#)) A Bugcheck 249 (*pointer page vanished*) error would be thrown when an attempt was made to fetch beyond the end-of-stream mark of an open PSQL cursor.

fixed by D. Yemanov

~ ~ ~

([CORE-2098](#)) It was not possible to create a view that selected from a global temporary table.

fixed by V. Khorsun

~ ~ ~

([CORE-2078](#)) If selective non-indexed predicates were involved in a join, the join plan was not optimized as well as it could be.

fixed by D. Yemanov

~ ~ ~

([CORE-2075](#)) Parts of the RDB\$DB_KEY of views could be inverted when using outer joins.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2073](#)) Expression indices bug: incorrect result for the inverted boolean.

fixed by D. Yemanov

~ ~ ~

([CORE-2069](#)) Incorrect view expansion when RDB\$DB_KEY was used inside a view body.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2068](#)) Comparison would return a wrong result with the **IN** <subquery_expression> operand if the <subquery_expression> argument involved the RDB\$DB_KEY.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2067](#)) GROUP BY and RDB\$DB_KEY problems

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2066](#)) Conversion of SQL_TEXT / SQL_VARCHAR to SQL_TIMESTAMP / SQL_TIME / SQL_DATE

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2053](#)) Computed expressions could suffer from poor optimization if used inside the RETURNING clause of the INSERT statement.

fixed by D. Yemanov

~ ~ ~

([CORE-2045](#)) A v.2.1 regression was picked up, whereby references to non-existent system fields with blr_field were not being resolved to NULL, whereas a parallel change involving blr fld was exhibiting the proper corrective behaviour.

fixed by dos Santos Fernandes

~ ~ ~

([CORE-2044](#)) Incorrect result for UPDATE OR INSERT ... RETURNING OLD and non-nullable columns

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2041](#)) UPDATE OR INSERT with GEN_ID() was causing the generator to step by 3.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2039](#)) Domain-level CHECK constraints were processing NULL values wrongly.

fixed by D. Yemanov

~ ~ ~

([CORE-2031](#)) NULL in the first record in a condition on RDB\$DB_KEY

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2027](#)) The buffer size for ORDER BY expressions involving system fields was being calculated wrongly.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2026](#)) Problem with a read-only marked database

fixed by V. Khorsun

~ ~ ~

([CORE-2008](#)) NOT NULL flag for procedure parameters in the system schema

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2002](#)) A conversion error from a UDF result would leave a memory leak if the result was marked with FREE_IT.

fixed by C. Valderrama

~ ~ ~

([CORE-2001](#)) When trying to show a conversion error, the message *arithmetic exception or string truncation* was sometimes appearing instead.

fixed by C. Valderrama

~ ~ ~

([CORE-2000](#)) The Lock Manager could report false deadlocks under high load.

fixed by V. Khorsun

~ ~ ~

([CORE-1986](#)) Altering the name of a domain was causing dependencies on the domain to be dropped.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1984](#)) The Lock Manager could falsely report a deadlock if one of the alleged participants was waiting with a permitted timeout.

fixed by V. Khorsun

~ ~ ~

([CORE-1980](#)) Sweeper could consume 100% of CPU indefinitely.

fixed by V. Khorsun

~ ~ ~

([CORE-1970](#)) Lock conversion denied (215) errors could occur.

fixed by V. Khorsun

~ ~ ~

([CORE-1962](#)) The function **EXTRACT (MILLISECONDS FROM aTimeStampOrTime)** was returning incorrect results.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1958](#)) A Bugcheck 179 (*decompression overran buffer*) consistency check error would be thrown when attempts were made to update the same record multiple times in a transaction.

fixed by D. Yemanov

~ ~ ~

([CORE-1957](#)) Long access control lists (ACLs) were being truncated, causing privileges to disappear.

fixed by A. Peshkov

~ ~ ~

([CORE-1943](#)) A statement that aggregated on a RAND() expression would return infinite rows.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1938](#)) Bugcheck 243 (missing pointer page) was being thrown on preparing or executing statements that referred to a table being dropped or recreated by another connection.

fixed by D. Yemanov

~ ~ ~

([CORE-1936](#)) The built-in function **LOG(base, number)** was not checking parameters and would deliver NAN values for out-of-range input instead of excepting.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1914](#)) If a problem occurred during table creation, the database could be left in an inconsistent state.

fixed by A. Peshkov

~ ~ ~

([CORE-1812](#)) For some date/time expressions in dialect 1, indexes were not being used when they should have been.

fixed by D. Yemanov

~ ~ ~

([CORE-1650](#)) An improbable case was demonstrated whereby something like **SELECT GEN_ID(..) FROM RDB\$DATABASE** with a **GROUP BY** operation would cause rows to be generated infinitely.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1607](#)) A correlated subquery that depended on a UNION stream would be poorly optimized.

fixed by D. Yemanov

~ ~ ~

([CORE-1606](#)) Ability to insert child record if parent record is locked but foreign key target unchanged

fixed by A. Potapchenko, V. Khorsun

~ ~ ~

([CORE-1575](#)) Multiple updates to a table in a single transaction would throw up a serious memory bug.

fixed by D. Yemanov

~ ~ ~

([CORE-1544](#)) When a user application created “temporary” stored procedures in run-time for some run-time purpose, the internal generator for the RDB\$PROCEDURES.RDB\$PROCEDURE_ID column could easily overflow the 32K limit (a signed SMALLINT) in its internal generator and cause a “numeric overflow” exception on trying to create the new stored procedure.

The fix wraps around the generated value at the 32K boundary, allowing reuse of existing gaps in the ID numbering. A similar fix was applied to RDB\$GENERATORS and RDB\$EXCEPTIONS as well.

fixed by D. Yemanov

~ ~ ~

([CORE-1343](#)) Simple case and a subquery

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1246](#)) Outer joins with derived tables returned incorrect results.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1245](#)) Outer joins with views returned incorrect results.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-903](#)) Firebird's behaviour with regard to multiple assignments referring to the same column in the SET clause of an UPDATE statement did not comply with the standard.

fixed by D. Yemanov

~ ~ ~

([CORE-501](#)) COALESCE exhibited an optimization problem.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-216](#)) Privileges could be lost if a database contained too many.

fixed by A. Peshkov

~ ~ ~

([CORE-1421](#)) SuperServer could not shut down immediately if the shutdown request followed a failed login attempt.

fixed by A. Peshkov

~ ~ ~

([CORE-1907](#)) Dropping and adding a domain constraint in the same transaction would leave incorrect dependencies.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1905](#)) Hash sign (#) in filenames in aliases.conf was being handled incorrectly.

fixed by C. Valderrama

~ ~ ~

([CORE-1887](#)) Newly created databases had wrong access rights.

fixed by A. Peshkov

~ ~ ~

([CORE-1869](#)) Roles granting/revoking logic differed between v.2.0 and v.2.1.

fixed by A. Peshkov

~ ~ ~

([CORE-1841](#)) If some VIEW used derived tables and long table names/aliases, it was possible to overflow RDB\$VIEW_RELATIONS.RDB\$CONTEXT_NAME.

fixed by V. Khorsun

~ ~ ~

([CORE-1840](#)) Each DDL execution would cause a small memory leak.

fixed by D. Yemanov

~ ~ ~

([CORE-1838](#)) SET STATISTICS INDEX on an index for a GTT could wrongly change the index id by the maximum available number for the database page size.

fixed by V. Khorsun

~ ~ ~

([CORE-1830](#)) Possible index corruption with multiple updates of the same record in the same transaction with savepoints being used.

fixed by V. Khorsun

~ ~ ~

([CORE-1817](#)) The RelaxedAliasChecking parameter was having no effect with regard to RDB\$DB_KEY.

fixed by V. Khorsun

~ ~ ~

([CORE-1811](#)) Parser reacted incorrectly to the unquoted usage of the keyword "VALUE".

fixed by D. Yemanov

~ ~ ~

([CORE-1798](#)) RDB\$DB_KEY was being evaluated as NULL in INSERT ... RETURNING.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1797](#)) OLD/NEW.RDB\$DB_KEY returned incorrect result in triggers.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1784](#)) Error with EXECUTE PROCEDURE inside EXECUTE STATEMENT.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1777](#)) Conflicting table reservation specifications were being allowed in the TPB.

fixed by C. Valderrama

~ ~ ~

([CORE-1775](#)) Security checking was performing poorly during prepare.

fixed by V. Khorsun

~ ~ ~

([CORE-1770](#)) Bugcheck 291 in DDL.

fixed by A. Peshkov

~ ~ ~

([CORE-1735](#)) Behaviour problem with SET DEFAULT action argument in referential integrity triggers.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1731](#)) Sometimes engine would hang for several minutes, using 1000% CPU load but with no I/O activity.

fixed by V. Khorsun

~ ~ ~

([CORE-1730](#)) Problems arose if one of the directories specified in the TempDirectories config setting was not available.

fixed by D. Yemanov

~ ~ ~

([CORE-1724](#)) Common table expressions could not be used in computed columns nor in quantified predicates (IN / ANY / ALL).

fixed by V. Khorsun

~ ~ ~

([CORE-1694](#)) Bug in CREATE/ALTER database trigger, where comments were in Russian.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1693](#)) Error in EXECUTE STATEMENT inside CONNECT / TRANSACTION START triggers.

fixed by A. dos Santos Fernandes, D. Yemanov

~ ~ ~

([CORE-1689](#)) “There are <n> dependencies” error message shows the wrong count of dependent objects

fixed by C. Valderrama

~ ~ ~

([CORE-1357](#)) DummyPacketInterval mechanism was broken.

fixed by D. Yemanov

~ ~ ~

([CORE-1307](#)) Switch -s of fb_inet_server was not being processed correctly

fixed by A. Peshkov

~ ~ ~

([CORE-479](#)) Grants would overwrite previous entries in RDB\$SECURITY_CLASSES.

fixed by A. Peshkov

~ ~ ~

Server/Client Crashes

([CORE-3011](#)) The server could hang or crash while monitoring connections that were repeatedly attaching and detaching.

fixed by D. Yemanov

~ ~ ~

([CORE-2908](#)) The engine could crash or raise unexpected errors when working with an ODS 8.x database.

fixed by V. Khorsun

~ ~ ~

([CORE-2888](#)) A source of memory corruption could cause incorrect query evaluation and could even crash the server.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2576](#)) The server could crash when parsing wrong or truncated BLR.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2455](#)) The server would fail when doing DROP DATABASE right after an error occurred in a statistical function.

fixed by A. Peshkov

~ ~ ~

([CORE-2441](#)) The server could crash on executing an UPDATE OR INSERT statement.

fixed by A. Peshkov

~ ~ ~

([CORE-2306](#)) Superserver could terminate abnormally when some worker thread failed to start.

fixed by A. Peshkov

~ ~ ~

[\(CORE-2368\)](#) An access violation would follow the call to `isc_cancel_events()` if the event was not found.

fixed by V. Khorsun

~ ~ ~

[\(CORE-2222\)](#) Storing a text blob with a transliterating blob filter could cause an access violation in the engine.

fixed by V. Khorsun

~ ~ ~

[\(CORE-2137\)](#) A database restore could crash the server when the configuration parameter *DummyPacketInterval* was set explicitly.

fixed by D. Yemanov

~ ~ ~

[\(CORE-2121\)](#) The Client library could crash in the course of an operation involving BLOBs.
c

fixed by A. Peshkov

~ ~ ~

[\(CORE-1983\)](#) An out-of-memory condition in the operating system would cause an access violation.

fixed by A. Peshkov

~ ~ ~

[\(CORE-1965\)](#) The Lock Manager would crash with an invalid lock ID under concurrent DDL load.

fixed by D. Yemanov

~ ~ ~

[\(CORE-1894\)](#) Circular dependencies between computed fields would crash the server.

fixed by A. dos Santos Fernandes

~ ~ ~

[\(CORE-1963\)](#) The server could crash on commit when granting/revoking privileges from multiple connections simultaneously.

fixed by D. Yemanov

~ ~ ~

[\(CORE-1506\)](#) Server crashes with `isc_dsqli_execute_immediate()` and a zero-length string.

fixed by A. Peshkov

~ ~ ~

([CORE-210](#)) The Classic server would crash if the same stored procedure was being altered in two separate processes.

fixed by D. Yemanov

~ ~ ~

([CORE-1930](#)) Possible server crash if procedure was altered to have no outputs and dependent procedures were not recompiled

fixed by V. Khorsun

~ ~ ~

([CORE-1919](#)) Memory corruptions in EXECUTE STATEMENT could crash the server.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1884](#)) Random crashes using stored procedures with expressions as default values for input parameters.

fixed by V. Khorsun

~ ~ ~

([CORE-1839](#)) Server could crash when sorting by a field that was calculated using a recursive CTE.

fixed by A. Peshkov

~ ~ ~

([CORE-1793](#)) Server crashes at prepare of query with unused parameterized CTE.

fixed by V. Khorsun

~ ~ ~

([CORE-1512](#)) Server crashes due to the wrong parsing of the DEFAULT clause.

fixed by D. Yemanov

~ ~ ~

Database Monitoring/Administration

([CORE-2209](#)) Monitoring requests in high load conditions could become very slow and even block other activity during that time.

fixed by D. Yemanov

~ ~ ~

[\(CORE-2171\)](#) The column MON\$CALLER_ID of table MON\$CALL_STACK was reporting invalid IDs.
fixed by D. Yemanov

~ ~ ~

[\(CORE-2017\)](#) I/O statistics for stored procedures were not being kept account of in the monitoring tables.
fixed by D. Yemanov

~ ~ ~

[\(CORE-1944\)](#) On Big Endian platforms, the monitoring tables contained wrong data.
fixed by A. Peshkov

~ ~ ~

[\(CORE-1890\)](#) Database monitoring process could hang under high load.
fixed by D. Yemanov

~ ~ ~

[\(CORE-1881\)](#) Database monitoring could crash the server or badly affect its page locking logic.
fixed by D. Yemanov

~ ~ ~

[\(CORE-1728\)](#) Database monitoring would not work after a fresh Linux install.
fixed by A. Peshkov

~ ~ ~

Data Manipulation Language

[\(CORE-1910\)](#) Invalid fields were accepted in the insert clause for MERGE.
fixed by A. dos Santos Fernandes

~ ~ ~

[\(CORE-1859\)](#) Arithmetic overflow or division by zero could occur in MAX() function.
fixed by A. dos Santos Fernandes

~ ~ ~

[\(CORE-1828\)](#) Error with ABS() function in dialect 1.
fixed by A. dos Santos Fernandes

~ ~ ~

Command-line Utilities

isql

([CORE-2831](#)) Database and user name should not be in the output when a script is extracted.

fixed by C. Valderrama

~ ~ ~

([CORE-2741](#)) Metadata extract would misinterpret the DDL of a CHECK constraint if the CHECK keyword was in any character mix other than all lower case or all upper case.

fixed by C. Valderrama

~ ~ ~

([CORE-915](#)) The metadata extract tools of *isql* were doubling line breaks in PSQL body code in cases where the PSQL body had been written in a third-party Windows text editor.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2408](#)) The *isql* metadata extraction process was placing default values for procedure parameters before the NOT NULL and COLLATE flags.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2407](#)) The *isql* metadata extraction process was omitting the PAGE_SIZE clause from the CREATE DATABASE statement.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2370](#)) An SQL plan of more than 2048 characters was not printed at all in *isql*.

fixed by C. Valderrama

~ ~ ~

([CORE-2270](#)) When run in a *zlogin* console, *isql* would consume all memory and crash.

fixed by J. Swierczynski, A. Peshkov

~ ~ ~

([CORE-1891](#)) SHOW VIEW would show nonsense information for view fields with expressions.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1875](#)) Errors in scripts with CURRENT_DATE.

fixed by V. Khorsun

~ ~ ~

([CORE-1862](#)) Extracted script was unusable with interdependent selectable procedures in FB 2.1

fixed by C. Valderrama

~ ~ ~

([CORE-1782](#)) Isql would crash when fetching data for a column having an alias longer than 30 characters.

fixed by D. Yemanov

~ ~ ~

([CORE-1749](#)) DDL statement with AUTODDL ON was not showing statistics.

fixed by D. Yemanov

~ ~ ~

([CORE-1507](#)) ISQL linecount facility in scripts goes out of sync after an INPUT command.

fixed by C. Valderrama

~ ~ ~

([CORE-1363](#)) ISQL would crash when the string converted from a double was longer than 23 bytes.

fixed by C. Valderrama

~ ~ ~

gsec

([CORE-2928](#)) Buffer overflow in *gsec*

fixed by A. Peshkov

For reasons unknown, the *gsec* code copies the value of the password hash to an internal user data structure during a display operation. Since V.2.0, when the newer hash algorithm made the hash longer than previously, the buffer used for storing it could be too short.

This does not create a vulnerability because the hash value does not travel anywhere. It is harmless, anyway: the buffer overflow cannot be exploited because the first, middle and last names are filled immediately after the password.

It is fixed now, thus avoiding having newer versions of *glibc* detecting this overflow.

~ ~ ~

([CORE-2528](#)) The gsec utility did not return error codes to the operating system.

fixed by C. Valderrama

~ ~ ~

([CORE-1680](#)) Gsec DISPLAY was showing only a few of the first users when the security databases contained more than 50 users.

fixed by A. Peshkov

~ ~ ~

gbak

([CORE-2914](#)) Restoring a database having an expression index referencing a non-existent UDF would cause the server to crash.

fixed by D. Yemanov

~ ~ ~

([CORE-2793](#)) The binary representation of a backup file was demonstrated to be not consistent from one backup to another, when tested over multiple backup/restore cycles of the same data. It was reported as applicable to all *gbak* versions since v.1.5.4. To date, it has been fixed here, in v.2.1.4 and in the v.3.0 alpha.

fixed by D. Yemanov

~ ~ ~

([CORE-2634](#)) A regression had arisen whereby performance dropped when restoring a database with a large amount of metadata.

fixed by A. Peshkov

~ ~ ~

([CORE-2291](#)) The error *Bugcheck 284 (cannot restore singleton select data)* would be thrown on bad trigger code involving [FOR] SELECT, when the engine should have been detecting the error and throwing the proper exception.

fixed by V. Khorsun

~ ~ ~

([CORE-2285](#)) A database with a large number of grants could become corrupted after a backup/restore.

fixed by A. Peshkov

~ ~ ~

([CORE-2245](#)) A database with long exception messages defined would exhibit errors when being restored from a backup.

fixed by C. Valderrama

~ ~ ~

[\(CORE-2223\)](#) gbak was encountering several bugs when operating on the access control lists (ACLs) that store SQL privileges.

fixed by A. Peshkov

~ ~ ~

[\(CORE-2214\)](#) Security classes were being restored incorrectly.

fixed by A. dos Santos Fernandes

~ ~ ~

[\(CORE-1911\)](#) Backup and restore were not thread-safe when using the Services API.

fixed by C. Valderrama

~ ~ ~

[\(CORE-1843\)](#) Gbak with Service Manager would not allow paths with spaces.

fixed by A. Peshkov

~ ~ ~

[\(CORE-1703\)](#) Delays/lockups when the gbak output was redirected to another process.

fixed by D. Yemanov

~ ~ ~

nBackup

[\(CORE-2750\)](#) Physical backup could not restart operations after the explicit difference file had been dropped.

fixed by C. Valderrama

~ ~ ~

[\(CORE-2648\)](#) NBackup's delta file was not respecting the "Forced Writes" database setting.

fixed by V. Khorsun

~ ~ ~

[\(CORE-2266\)](#) nBackup's database locking was not working correctly, causing database file growth to continue when database writes should have been in suspension.

fixed by V. Khorsun

~ ~ ~

([CORE-1696](#)) Deadlock would occur in the Lock Manager when the NBackup utility was in use.

fixed by R. Simakov

~ ~ ~

([CORE-1876](#)) Incremental backups with NBACKUP were broken in v.2.1.

fixed by N. Samofatov

~ ~ ~

gfix

([CORE-2846](#)) If **gfix -shut <mode> -attach <timeout>** failed after the specified timeout due to connections being still active, connection to the database would become impossible.

fixed by D. Yemanov

~ ~ ~

([CORE-97](#)) A lock on the database file was being left by **gfix -shut -force**, preventing a subsequent restore.

fixed by D. Yemanov

~ ~ ~

([CORE-2268](#)) Non-valid transaction numbers would cause **gfix** to throw BUGCHECK errors.

fixed by V. Khorsun

~ ~ ~

([CORE-2271](#)) The *gfix* utility had a legacy bug that exhibited itself during the database validation/repair routines on large databases. The privilege level of the user running these routines was being checked too late in the operation, thus allowing a non-privileged user (i.e., not SYSDBA or Owner) to start a validation operation. Once the privilege check occurred, the database validation could halt in mid-operation and thus be left unfinished, resulting in logical corruption that might not have been there otherwise.

fixed by A. Peshkov

~ ~ ~

([CORE-1961](#)) A Bugcheck 210 (*page in use during flush*) consistency check error would be thrown during database validation.

fixed by D. Yemanov, R. Simakov

~ ~ ~

gstat

([CORE-2519](#)) Output from **gstat** was incorrect for tables with more than 2 billion records.

fixed by V. Khorsun

~ ~ ~

([CORE-1412](#)) Some long-standing bugs in gstat's processing of parameters needed fixing.

fixed by C. Valderrama

~ ~ ~

fb_lock_print

([CORE-2598](#)) **fb_lock_print -c[onsistency]** switch was not working on Windows.

fixed by D. Yemanov

~ ~ ~

([CORE-2354](#)) “fb_lock_print -ia” output was not being flushed to the file between iterations.

fixed by A. Peshkov

~ ~ ~

qli Query Utility for GDML

([CORE-2247](#)) In the QLI utility, message and descriptor buffers were not properly aligned.

fixed by A. Peshkov

~ ~ ~

Services Manager

([CORE-1982](#)) Simultaneous backups or restores invoked through the Services API could interfere with one another.

fixed by A. dos Santos Fernandes

~ ~ ~

Remote Interface/API

([CORE-2563](#)) It was possible to shut down the Superserver's main port (3050 by default) by sending a malformed packet of some special format, that would lead to a Denial of Service condition for new incoming connections. This exploit could be used by an unauthenticated client.

Reported 15-Jul-2009 by Core Security Technologies.

fixed by D. Yemanov

~ ~ ~

([CORE-2437](#)) A buffer overflow could occur on the client when receiving events.

fixed by A. Peshkov

~ ~ ~

([CORE-2307](#)) API information requests were returning incomplete values in the results.

fixed by D. Yemanov

~ ~ ~

([CORE-2234](#)) Sometimes, terminated worker processes in Classic on Windows were still considered to be alive after termination, due to improper checks on the Firebird server's part. The same bug could cause the Firebird server to misbehave with prolonged deadlocks when under load.

fixed by V. Khorsun

~ ~ ~

([CORE-2151](#)) When a temporary directory path had spaces within it, it was (wrongly) being truncated at the rightmost space.

fixed by V. Khorsun

~ ~ ~

([CORE-2033](#)) The symbol `_Unwind_GetIP` in the client library was being left unresolved due to a missing static library linkage.

fixed by A. Peshkov

~ ~ ~

([CORE-2018](#)) Only a single client could access a read-only database.

fixed by V. Khorsun

~ ~ ~

([CORE-1972](#)) A non-SYSDBA user was able to change the Forced Writes mode of any database, along with several other database characteristics that should be restricted to the SYSDBA. This long-standing, legacy loophole in the handling of DPB parameters could lead to database corruptions or give ordinary users access to SYSDBA-only operations.

The changes could affect several existing applications, database tools and connectivity layers (drivers, components).

Same fix was backported to v.2.1.2 and v.2.0.5.

fixed by A. Peshkov

~ ~ ~

([CORE-1868](#)) Client library was crashing inside `isc_dsqli_free_statement()`.

fixed by A. Peshkov

~ ~ ~

([CORE-1763](#)) The client library was not setting the options SO_KEEPALIVE nor TCP_NODELAY for the socket at connection.

fixed by V. Khorsun

~ ~ ~

([CORE-1755](#) and [CORE-1756](#)) A couple of crash scenarios could occur in `isc_start_transaction()`.

fixed by D. Kovalenko, A. Peshkov

~ ~ ~

([CORE-1726](#)) Failure in `isc_service_start()`.

fixed by A. Peshkov

~ ~ ~

([CORE-1079](#)) Every attach of `fbclient/fbembed` library to the host process would leak 64KB of memory.

fixed by A. Peshkov

~ ~ ~

Security

([CORE-2657](#)) A poorly documented SPB tag provided the undesirable ability to pass an arbitrary trusted user name to a service and enable that user to acquire any permissions, including those of SYSDBA.

fixed by A. Peshkov

~ ~ ~

([CORE-2087](#)) When the configuration parameter `RemoteBindAddress` specified a hostname instead of an IP address, or specified a non-existent IP address, it would be silently ignored and the server would bind to all interfaces, without any notification in `firebird.log` or the system log. This was considered a potential security risk if the system had ports open to the Internet. Now, an invalid or unavailable IP address will be resolved to `localhost (127.0.0.1)`.

fixed by A. Peshkov

~ ~ ~

([CORE-2055](#)) Buffer overflow in Firebird client library.

fixed by A. Peshkov

~ ~ ~

([CORE-1845](#)) Some standard calls would show the server installation directory to regular users.

fixed by A. Peshkov

~ ~ ~

([CORE-1810](#)) Some issues appeared concerning usernames with '.' characters.

fixed by A. Peshkov

~ ~ ~

International Language Support

~ ~ ~

([CORE-2642](#)) An obscure ICU initialization problem on Windows could cause misbehaviour in a multi-threaded environment.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2607](#)) Introducer syntax(_charset) problems were occurring when it was used in association with monitoring queries or PSQL modules.

For information about the problem and a workaround provided for it, see the topic [Introducer Syntax Usage](#) under *Other Improvements* in the International Language Support chapter.

fixed by A. dos Santos Fernandes

~ ~ ~

~ ~ ~

([CORE-2361](#)) String truncation was occurring when reading a character set 8859_1 Spanish column using `isc_dsql_fetch()` with a UTF8 connection.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2227](#)) Problems were occurring in some environments when creating triggers that referred to column names with accented characters.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-2123](#)) Problem with getting UNICODE_FSS data in the CP943C connection charset

fixed by A. dos Santos Fernandes, D. Kovalenko

~ ~ ~

([CORE-2122](#)) Translation of large text blobs between UNICODE_FSS / UTF8 and other charsets

fixed by A. dos Santos Fernandes, D. Kovalenko

~ ~ ~

([CORE-2095](#)) Bug in CVJIS_eucj_to_unicode().

fixed by A. dos Santos Fernandes, D. Kovalenko

~ ~ ~

([CORE-2019](#)) A UTF-8 conversion error (string truncation) was being thrown unexpectedly.

fixed by dos Santos Fernandes

~ ~ ~

([CORE-1989](#)) A column with UNICODE_CI collation for UTF8 could not be used in a foreign key constraint.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1927](#)) The procedure **sp_register_character_set** could generate a negative value for RDB\$CHARACTER_SETS.RDB\$CHARACTER_SET_ID.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1690](#)) A condition in tables with UTF8 text was causing the error *Arithmetic exception, numeric overflow, or string truncation*.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1596](#)) Bug in CsConvert::convert

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1432](#)) The collation attribute of columns was not being propagated between record formats.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-316](#)) A database with multi-byte characters in its name could not be opened.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1802](#)) Some issues were reported concerning maximum key size using the PXW_CSY collation.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1774](#)) A problem appeared with collate ES_ES_CI_AI.

fixed by A. dos Santos Fernandes

~ ~ ~

([CORE-1254](#)) Problem with DISTINCT and insensitive collations.

fixed by A. dos Santos Fernandes

~ ~ ~

POSIX-specific

([CORE-3067](#)) On 64-bit POSIX, except HP-UX, objects were not being unmapped when shared memory was closed.

The 64-bit pointer had been masked with a 32-bit mask, which turned out to be a really bad idea: with `Object == 0x7FAB12345678` and `page size == 4K`, the upper bits were lost and the starting address would be at `0x12345000` instead of the intended `0x7FAB12345000`.

fixed by A. Peshkov

~ ~ ~

([CORE-3019](#)) On a recent version of Gentoo Linux, an attempt to source the function library `/etc/init.d/functions.sh` would fail with the message “* ERROR: firebird does not have a start function” and neither SuperServer nor SuperClassic would start.

fixed by A. Peshkov

~ ~ ~

([CORE-3001](#)) Install was failing to create the `firebird` user and group.

fixed by A. Peshkov

~ ~ ~

([CORE-2919](#)) The Linux installation script was ignoring non-standard ports.

fixed by A. Peshkov

~ ~ ~

([CORE-2845](#)) The build process on Solaris would stop with the message: “need to use SFIO”. On the whole, Solaris 10 no longer requires SFIO - the 64-bit builds do not suffer the issue with the maximum 255 file descriptors and there are simple work-arounds for 32-bit builds. See [this article](#) for details.

However, the issue still affects Solaris 10 3/05 through Solaris 10 11/06. [Some patches](#) are needed to correct the problem.

For the time being, the easiest solution is to leave the *define* in the code, but comment it with the relevant information for Solaris 10. The *define* can be removed if the patches have been applied or the user is running an unaffected version of Solaris 10.

Updated information concerning this issue has been inserted into *common.h*. Anyone building 32-bit Firebird on Solaris 10 will have to make an informed decision according to the version and build/patch level of Solaris before uncommenting the SFIO define.

fixed by P. Beach

~ ~ ~

([CORE-2814](#)) On SPARC, the server was crashing in the routine `map_sort_data` with a bus error.

fixed by V. Khorsun

~ ~ ~

([CORE-2601](#)) A lot of the standard **configure** switches for fine-tuning the installation directories on POSIX platforms do not work for Firebird.

It was close to impossible to make the standard GNU switches work without changing the defaults for them, a rigmarole that is far from obvious or easy. Instead, a set of new switches for the **configure** has been added to enable fine-level configuration of the locations of Firebird's files.

The switches are listed in the Installation chapter, in the topic [Dedicated Firebird Switches for 'configure'](#).

fixed by A. Peshkoff

~ ~ ~

([CORE-2572](#)) Locks of type `LCK_page_space` were being processed incorrectly on big-endian machines.

fixed by A. Peshkov

~ ~ ~

([CORE-2221](#)) On POSIX platforms, any attachment to any database would fail after the access rights for `security2.fdb` were modified from 0660 to 0666.

fixed by P. Beach, A. Peshkov

~ ~ ~

([CORE-2093](#)) SuperServer startup would fail on Solaris 64-bit.

fixed by A. Peshkov

~ ~ ~

([CORE-1909](#)) Garbage in `firebird.log` on linux/amd64

fixed by A. Peshkov

~ ~ ~

([CORE-1885](#)) `CREATE COLLATION` caused lost connection under Posix.

fixed by A. dos Santos Fernandes, A. Peshkov

~ ~ ~

([CORE-1854](#)) Value of CURRENT_USER might not be in upper case when using Unix OS authentication.

fixed by A. Peshkov

~ ~ ~

([CORE-1826](#)) changeRunUser.sh and restoreRootRunUser.sh scripts were not changing the run user in the init.d scripts.

fixed by A. Peshkov

~ ~ ~

([CORE-1818](#)) Temporary files used for temporary page spaces were not deleted after use on Posix platforms.

fixed by A. Peshkov

~ ~ ~

([CORE-1807](#)) Server was being assigned to a non-canonical port after abnormal termination.

fixed by A. Peshkov

~ ~ ~

([CORE-1766](#)) Owner and group of isc_monitor1 file on Linux classic server were incorrect.

fixed by A. Peshkov

~ ~ ~

([CORE-1671](#)) atexit() calls in client libraries cause segfaults if the libraries were used in dlopen'ed modules.

fixed by A. Peshkov

~ ~ ~

Windows-specific

([CORE-2769](#)) On heavily loaded Windows systems, local connect (XNET) could fail due to the client timing out while waiting for the server to set the `xnet_response_event`. To help with this problem, the `ConnectionTimeout` parameter in `firebird.conf` has been activated for XNET connections.

Improvement made by D. Yemanov

~ ~ ~

([CORE-2108](#)) When using the Windows local protocol (XNET), the next available map number was calculated incorrectly, thus allowing the server to try to reuse a map number that already existed. If the “new” map's timestamp was equal to the timestamp of the pre-existing map, it would cause the `get_free_slot()` function to fail.

fixed by V. Khorsun

~ ~ ~

([CORE-2107](#)) Establishing a TCP/IP connection to the Windows Classic Server could fail under high load.

fixed by V. Khorsun

~ ~ ~

([CORE-1923](#)) Successful execution of **instsvc.exe remove** was returning 1 as its completion code, instead of 0.

fixed by D. Yemanov

~ ~ ~

([CORE-1820](#)) Setup program was failing to detect a running server.

fixed by P. Reeves, D. Yemanov

~ ~ ~

([CORE-1105](#), [CORE-1390](#), [CORE-1566](#) & [CORE-1639](#)) Aliases would not work properly for XNET connections.

fixed by D. Yemanov

~ ~ ~

MacOSX-specific

([CORE-2065](#)) The MacOSX installation package was in violation of platform rules by not including the client library in the dynamic loader search paths.

fixed by P. Beach

~ ~ ~

Miscellaneous Bugs

([CORE-2282](#)) Truncating UDFs were broken for negative numbers below -1.

fixed by C. Valderrama

~ ~ ~

([CORE-2281](#)) Rounding UDFs were broken for negative numbers.

fixed by C. Valderrama

~ ~ ~

Firebird 2.5 Project Teams

Table 18.1. Firebird Development Teams

Developer	Country	Major Tasks
Dmitry Yemanov	Russian Federation	Full-time database engineer/implementor, core team leader
Alex Peshkov	Russian Federation	Full-time security features coordinator; buildmaster; porting authority
Claudio Valderrama	Chile	Code scrutineer; bug-finder and fixer; ISQL enhancements; UDF fixer, designer and implementor
Vladyslav Khorsun	Ukraine	Full-time DB engineer, SQL feature designer/implementor
Arno Brinkman	The Netherlands	Indexing and Optimizer enhancements; new DSQL features
Adriano dos Santos Fernandes	Brazil	New international character-set handling; text and text BLOB enhancements; new DSQL features; code scrutineering
Nickolay Samofatov	Russian Federation	Engine contributions
Roman Simakov	Russian Federation	Engine contributions
Bill Oliver	U.S.A.	Vulcan fork development, engine contributions
Paul Beach	France	Release Manager; HP-UX builds; MacOS Builds; Solaris Builds
Pavel Cisar	Czech Republic	QA tools designer/coordinator
Philippe Makowski	France	QA tester
Paul Reeves	France	Win32 installers and builds
Roman Rokytssky	Germany	Jaybird implementor and co-coordinator
Evgeny Putilin	Russian Federation	Java stored procedures implementation
Jiri Cincura	Czech Republic	Developer and coordinator of .NET providers
Vladimir Tsvigun	Ukraine	Developer and coordinator of ODBC/JDBC driver for Firebird

Firebird 2.5 Project Teams

Developer	Country	Major Tasks
Stephen Boyd	Canada	GPRE contributions
Paul Vinkenoog	The Netherlands	Coordinator, Firebird documentation project; documentation writer and tools developer/implementor
Norman Dunbar	U.K.	Documentation writer
Pavel Menshchikov	Russian Federation	Documentation translator
Tomneko Hayashi	Japan	Documentation translator
Umberto (Mimmo) Masotti	Italy	Documentation translator
Helen Borrie	Australia	Release notes editor; Chief of Thought Police
also		
Université du Littoral Côte d'Opale Masters students	France	QA tests development

Appendix A: SQLSTATE

SQLSTATE Codes & Messages

In this Appendix are all of the SQLSTATE codes currently supported:

1. The 5-character SQLSTATE code returned by the status array consists of SQL CLASS (2 characters) and SQL SUBCLASS (3 characters)
2. Where existent and known, 1:1 mappings to the deprecated SQLCODE are included.
3. In many cases, SQLCODE:SQLSTATE mappings are not 1:1, which is intentional on the part of the SQL Standards committee. It has been their aim, for many years, that the use of the SQLCODE be deprecated entirely.

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
<i>SQLCLASS 00 (Success)</i>		
00000	Success	
<i>SQLCLASS 01 (Warning)</i>		
01000	General Warning	
01001	Cursor operation conflict	
01002	Disconnect error	
01003	NULL value eliminated in set function	
01004	String data, right-truncated	
01005	Insufficient item descriptor areas	
01006	Privilege not revoked	
01007	Privilege not granted	
01008	Implicit zero-bit padding	
01100	Statement reset to unprepared	
01101	Ongoing transaction has been committed	
01102	Ongoing transaction has been rolled back	
<i>SQLCLASS 02 (No Data)</i>		
02000	No data found or no rows affected	
<i>SQLCLASS 07 (Dynamic SQL error)</i>		

SQLSTATE

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
07000	Dynamic SQL error	
07001	Wrong number of input parameters	
07002	Wrong number of output parameters	
07003	Cursor specification cannot be executed	
07004	USING clause required for dynamic parameters	
07005	Prepared statement not a cursor-specification	
07006	Restricted data type attribute violation	
07007	USING clause required for result fields	
07008	Invalid descriptor count	
07009	Invalid descriptor index	
<i>SQLCLASS 08 (Connection Exception)</i>		
08001	Client unable to establish connection	
08002	Connection name in use	
08003	Connection does not exist	
08004	Server rejected the connection	
08006	Connection failure	
08007	Transaction resolution unknown	
<i>SQLCLASS 0A (Feature Not Supported)</i>		
0A000	Feature Not Supported	
<i>SQLCLASS 0B (Invalid Transaction Initiation)</i>		
0B000	Invalid transaction initiation	
<i>SQLCLASS 0L (Invalid Grantor)</i>		
0L000	Invalid grantor	
<i>SQLCLASS 0P (Invalid Role Specification)</i>		
0P000	Invalid role specification	
<i>SQLCLASS 0U (Attempt to Assign to Non-Updatable Column)</i>		
0U000	Attempt to assign to non-updatable column	
<i>SQLCLASS 0V (Attempt to Assign to Ordering Column)</i>		
0V000	Attempt to assign to Ordering column	
<i>SQLCLASS 20 (Case Not Found For Case Statement)</i>		
20000	Case not found for case statement	

SQLSTATE

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
<i>SQLCLASS 21 (Cardinality Violation)</i>		
21000	Cardinality violation	
21S01	Insert value list does not match column list	
21S02	Degree of derived table does not match column list	
<i>SQLCLASS 22 (Data Exception)</i>		
22000	Data exception	
22001	String data, right truncation	
22002	Null value, no indicator parameter	
22003	Numeric value out of range	
22004	Null value not allowed	
2205	Error in assignment	
2206	Null value in field reference	
2207	Invalid datetime format	
22008	Datetime field overflow	
22009	Invalid time zone displacement value	
2200A	Null value in reference target	
2200B	Escape character conflict	
2200C	Invalid use of escape character	
2200D	Invalid escape octet	
2200E	Null value in array target	
2200F	Zero-length character string	
2200G	Most specific type mismatch	
22010	Invalid indicator parameter value	
22011	Substring error	
22012	Division by zero	
22014	Invalid update value	
22015	Interval field overflow	
22018	Invalid character value for cast	
22019	Invalid escape character	
2201B	Invalid regular expression	
2201C	Null row not permitted in table	

SQLSTATE

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
22020	Invalid limit value	
22021	Character not in repertoire	
22022	Indicator overflow	
22023	Invalid parameter value	
22024	Character string not properly terminated	
22025	Invalid escape sequence	
22026	String data, length mismatch	
22027	Trim error	
22028	Row already exists	
2202D	Null instance used in mutator function	
2202E	Array element error	
2202F	Array data, right truncation	
<i>SQLCLASS 23 (Integrity Constraint Violation)</i>		
23000	Integrity constraint violation	
<i>SQLCLASS 24 (Invalid Cursor State)</i>		
24000	Invalid cursor state	
24504	The cursor identified in the UPDATE, DELETE, SET, or GET statement is not positioned on a row	
<i>SQLCLASS 25 (Invalid Transaction State)</i>		
25000	Invalid transaction state	
25	xxxx	
25S01	Transaction state	
25S02	Transaction is still active	
25S03	Transaction is rolled back	
<i>SQLCLASS 26 (Invalid SQL Statement Name)</i>		
26000	Invalid SQL statement name	
<i>SQLCLASS 27 (Triggered Data Change Violation)</i>		
27000	Triggered data change violation	
<i>SQLCLASS 28 (Invalid Authorization Specification)</i>		
28000	Invalid authorization specification	
<i>SQLCLASS 2B (Dependent Privilege Descriptors Still Exist)</i>		

SQLSTATE

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
2B000	Dependent privilege descriptors still exist	
<i>SQLCLASS 2C (Invalid Character Set Name)</i>		
2C000	Invalid character set name	
<i>SQLCLASS 2D (Invalid Transaction Termination)</i>		
2D000	Invalid transaction termination	
<i>SQLCLASS 2E (Invalid Connection Name)</i>		
2E000	Invalid connection name	
<i>SQLCLASS 2F (SQL Routine Exception)</i>		
2F000	SQL routine exception	
2F002	Modifying SQL-data not permitted	
2F003	Prohibited SQL-statement attempted	
2F004	Reading SQL-data not permitted	
2F005	Function executed no return statement	
<i>SQLCLASS 33 (Invalid SQL Descriptor Name)</i>		
33000	Invalid SQL descriptor name	
<i>SQLCLASS 34 (Invalid Cursor Name)</i>		
34000	Invalid cursor name	
<i>SQLCLASS 35 (Invalid Condition Number)</i>		
35000	Invalid condition number	
<i>SQLCLASS 36 (Cursor Sensitivity Exception)</i>		
36001	Request rejected	
36002	Request failed	
<i>SQLCLASS 37 (Invalid Identifier)</i>		
37000	Invalid identifier	
37001	Identifier too long	
<i>SQLCLASS 38 (External Routine Exception)</i>		
38000	External routine exception	
<i>SQLCLASS 39 (External Routine Invocation Exception)</i>		
39000	External routine invocation exception	
<i>SQLCLASS 3B (Invalid Save Point)</i>		
3B000	Invalid save point	

SQLSTATE

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
<i>SQLCLASS 3C (Ambiguous Cursor Name)</i>		
3C000	Ambiguous cursor name	
<i>SQLCLASS 3D (Invalid Catalog Name)</i>		
3D000	Invalid catalog name	
3D001	Catalog name not found	
<i>SQLCLASS 3F (Invalid Schema Name)</i>		
3F000	Invalid schema name	
<i>SQLCLASS 40 (Transaction Rollback)</i>		
40000	Ongoing transaction has been rolled back	
40001	Serialization failure	
40002	Transaction integrity constraint violation	
40003	Statement completion unknown	
<i>SQLCLASS 42 (Syntax Error or Access Violation)</i>		
42000	Syntax error or access violation	
42702	Ambiguous column reference	
42725	Ambiguous function reference	
42818	The operands of an operator or function are not compatible	
42S01	Base table or view already exists	
42S02	Base table or view not found	
42S11	Index already exists	
42S12	Index not found	
42S21	Column already exists	
42S22	Column not found	
<i>SQLCLASS 44 (With Check Option Violation)</i>		
44000	WITH CHECK OPTION Violation	
<i>SQLCLASS 45 (Unhandled User-defined Exception)</i>		
45000	Unhandled user-defined exception	
<i>SQLCLASS 54 (Program Limit Exceeded)</i>		
54000	Program limit exceeded	
54001	Statement too complex	
54011	Too many columns	

SQLSTATE

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
54023	Too many arguments	
<i>SQLCLASS HY (CLI-specific Condition)</i>		
<i>HY000</i>	CLI-specific condition	
<i>HY001</i>	Memory allocation error	
<i>HY003</i>	Invalid data type in application descriptor	
<i>HY004</i>	Invalid data type	
<i>HY007</i>	Associated statement is not prepared	
<i>HY008</i>	Operation canceled	
<i>HY009</i>	Invalid use of null pointer	
<i>HY010</i>	Function sequence error	
<i>HY011</i>	Attribute cannot be set now	
<i>HY012</i>	Invalid transaction operation code	
<i>HY013</i>	Memory management error	
<i>HY014</i>	Limit on the number of handles exceeded	
<i>HY015</i>	No cursor name available	
<i>HY016</i>	Cannot modify an implementation row descriptor	
<i>HY017</i>	Invalid use of an automatically allocated descriptor handle	
<i>HY018</i>	Server declined the cancellation request	
<i>HY019</i>	Non-string data cannot be sent in pieces	
<i>HY020</i>	Attempt to concatenate a null value	
<i>HY021</i>	Inconsistent descriptor information	
<i>HY024</i>	Invalid attribute value	
<i>HY055</i>	Non-string data cannot be used with string routine	
<i>HY090</i>	Invalid string length or buffer length	
<i>HY091</i>	Invalid descriptor field identifier	
<i>HY092</i>	Invalid attribute identifier	
<i>HY095</i>	Invalid FunctionId specified	
<i>HY096</i>	Invalid information type	
<i>HY097</i>	Column type out of range	
<i>HY098</i>	Scope out of range	
<i>HY099</i>	Nullable type out of range	

SQLSTATE

SQLSTATE Code	Mapped Message	Maps to SQLCODE..
<i>HY100</i>	Uniqueness option type out of range	
<i>HY101</i>	Accuracy option type out of range	
<i>HY103</i>	Invalid retrieval code	
<i>HY104</i>	Invalid LengthPrecision value	
<i>HY105</i>	Invalid parameter type	
<i>HY106</i>	Invalid fetch orientation	
<i>HY107</i>	Row value out of range	
<i>HY109</i>	Invalid cursor position	
<i>HY110</i>	Invalid driver completion	
<i>HY111</i>	Invalid bookmark value	
<i>HYC00</i>	Optional feature not implemented	
<i>HYT00</i>	Timeout expired	
<i>HYT01</i>	Connection timeout expired	
<i>SQLCLASS XX (Internal Error)</i>		
<i>XX000</i>	Internal error	
<i>XX001</i>	Data corrupted	
<i>XX002</i>	Index corrupted	

Appendix B: Licence Notice

The contents of this Documentation are subject to the Public Documentation License Version 1.0 (the “License”); you may only use this Documentation if you comply with the terms of this Licence. Copies of the Licence are available at <http://www.firebirdsql.org/pdfmanual/pdl.pdf> (PDF) and <http://www.firebirdsql.org/manual/pdl.html> (HTML).

The Original Documentation is entitled *Firebird 2.5 Release Notes*.

The Initial Writer of the Original Documentation is: Helen Borrie. Persons named in attributions are Contributors.

Copyright (C) 2004-2009. All Rights Reserved. Initial Writer contact: helebor at users dot sourceforge dot net.