Firebird® Version 1.5.2

Point Release Notes v.152_09

11 December 2004

R E G R E S S I O N S F I X E D

ISSUE	SF Bug # Fixed by
Blobs containing more than 65535 segments were not backed up by GBAK properly.	v.1.0 Regression
Solution Fixed in 1.5.2	N. Samofatov
When a SEGV error (or other asyncronous exception) is thrown from a badly written UDF, the server should log its name. This feature was broken in FB 1.5.	v.1.5.0 regression
Solution Restored in v.1.5.2	V. Horsun
In v.1.5.0, exit(3) was called on critical errors on Windows, precluding the use of the JIT debugger to analyse problems in UDF routines. A fix in v.1.5.1 caused the debugger to be called always, introducing a problem with automatic restarts of the server.	Regression
Solution V. 1.5.2 now calls the debugger only if you set the BugcheckAbort configuration file option to 1. The main goal of the v1.5.2 fix is to avoid Dr.Watson showing its close-or-debug message box.	V. Horsun, N. Samofatov
Support for multi-dimensional array fields was broken	v1.5.1 regression
Solution Restored in v.1.5.2	C. Valderrama
Plans for selectable stored procedures containing multiple FOR loops were being reported in reverse order, compared with v.1.5.0.	v1.5.1 regression

Solution Restored in v.1.5.2	V. Horsun, N. Samofatov

An optimizer regression, present since v.1.5.0, meant certain cases were not optimized properly where outer join syntax was used to perform inner joins.
Example 1
 	
`SELECT *`
`FROM RDB$RELATIONS r`
`  LEFT JOIN RDB$RELATION_FIELDS rf`
`  ON (1 = 1)`
`WHERE`
`  r.RDB$RELATION_NAME = rf.RDB$RELATION_NAME`
FB1.5 returned this plan:
 	
`  PLAN JOIN (R NATURAL,RF NATURAL)`
RDB$RELATION_FIELDS would fetch all records, whereas it ought only to fetch those which match with r.RDB$RELATION_NAME = rf.RDB$RELATION_NAME
 	
Example 2
 	
`SELECT`
`  R.RDB$RELATION_NAME,`
`  RF.RDB$RELATION_NAME`
`FROM`
`  RDB$RELATIONS R`
`  LEFT JOIN RDB$RELATION_FIELDS RF ON`
`  (RF.RDB$RELATION_NAME = R.RDB$RELATION_NAME`
`  AND RF.RDB$RELATION_NAME >= 'zzz')`
The equalities would be distributed wrongly, causing them to be sent to the outer join rather than being applied to the inner stream.
	v1.5.0 regression

Solution In 1.5.2, Example 1 returns this plan: `  PLAN JOIN (R NATURAL,RF INDEX (RDB$INDEX_4))` The issue in Example 2 is also resolved by this fix.	A. Brinkman

---

Documentation for Windows Embedded Server, README.user.embedded in the /doc subdir of the Windows installation, section 2.2 "Database Access" was "corrected" wrongly in v.1.5.1.	Regression

H. Borrie

Solution
This section was unclear in v.1.5.0. In v.1.5.1 it was plain wrong. It is now correct.

B U G S F I X E D

ISSUE	SF Bug # / Fixed by
V. 1.5 Superserver had a potentially severe bug that caused a deadlock to occur if gbak tried to restore a database into a directory to which the DatabaseAccess settings in firebird.conf did not allow access.	v.1.5.0 bug
Solution The fix for this bug has been back-ported to v.1.5.2 from Firebird2 HEAD.	A. Peshkov
A client on Windows XP SP2 was slow to connect to a Linux server.	1065511
Solution Solved.	N. Samofatov
An EXISTS or SINGULAR predicate buried inside an aggregate function caused wrong detection/mapping for aggregate queries.	1063254
Solution Solved.	A. Brinkman
ISQL had a few parsing problems: 1. Semi-colons and '--' sequences were being interpreted as statement terminators and comments, respectively, in multi-line literals 2. Tab characters in literal strings were being translated to spaces	Not logged
Solution All of these problems have now been fixed.	D. Sibiryakov, C. Valderrama, N. Samofatov
The server could crash when execute_immediate was used to release or rollback a transaction to a non-existing savepoint.	Not logged
Solution Now fixed.	N. Samofatov

An old legacy bug that has continued to bug us is that, when a client had some events registered and its network connection had been terminated abnormally (hardware failure, reset button or task manager), then the server would start using 100% of the CPU time until the "parent" port (client connection which called isc_que_events() API routine) reported on its failure. This bug affected all FB versions (more or less, depending on the DummyPacketInterval configuration option) and only TCP/IP connections.	1045970
Solution Further work has been done to rectify the problem in v.1.5.2. It now appears to be solved.	D.Yemanov

---

In v.1.5.1, a bug in the optimiser caused unnecessary fetches in joined relations when "OR" was used on the base relation, as in the following example:
 	
`SELECT * FROM RDB$RELATIONS r`
`LEFT JOIN RDB$RELATION_FIELDS rf`
`ON (r.RDB$RELATION_NAME = rf.RDB$RELATION_NAME)`
`WHERE`
`  r.RDB$RELATION_ID = 0 OR 1 = 0`
	Not logged

Solution Fixed	A. Brinkman

UDF with NULL input parameter. Explained here.	544132
Solution Fixed	C. Valderrama

Left join defeats UDF by mangling a null descriptor. Explained here.	728839
Solution Fixed	C. Valderrama

Error trying to delete from a naturally updatable view containing computed expressions. If you had a view like this:
 	
`create view v_test (f, s)`
`as`
`  select f1, f2 + f3`
`from t`
 	
then the server returned "attempted update of read-only column" error when you tried to perform a DELETE operation.
	Not logged

Solution This is fixed in FB 1.5.2 and above.	D. Yemanov

Numeric data types represented by floating-point variables were being processed incorrectly in an EXECUTE STATEMENT with dialect 1 databases---numerics were being scaled incorrectly: — Not logged

```
create table a (b numeric(18,3));
commit;
insert into a values(12345.678);
commit;
set term ^;
create procedure c
  returns(d numeric(18,3))
as
begin
  for execute statement 'select b from a'
  into :d
  do suspend;
end
^
commit^
set term ;^
select * from c;
```

returned 12.346 instead of 12345.678.

Solution — A. Peshkoff

Fixed

If DISTINCT was used in an aggregate function and the record set being processed (aggregated) is empty, then we had a small memory leak. This memory was not returned until disconnect. — Not logged

This routine would eat 120MB on FB 1.5.1 and previous:

```
CREATE PROCEDURE MEM_LEAK
AS
  DECLARE I INT = 1000;
  DECLARE T INT;
  DECLARE C INT;
BEGIN
  WHILE (I > 0) DO
  BEGIN
    SELECT RDB$INDEX_TYPE,
    COUNT(DISTINCT RDB$RELATION_NAME)
      FROM RDB$INDICES
     WHERE 0=1
    GROUP BY 1
    INTO :T, :C;

    I = I - 1;
  END
END
```

Solution — V. Horsun

Fixed in v.1.5.2

The server leaked resources when an exception was thrown from a selectable stored procedure. The procedural request wasn't freed properly and caused errors like "too many concurrent executions of the same request" after 750-1000 iterations. — Not logged

```
CREATE PROCEDURE P (INP INTEGER)
RETURNS (OUTP INTEGER)
AS
BEGIN
  OUTP = INP / 0;
  SUSPEND;
END

UPDATE T SET ID = 1
  WHERE (SELECT OUTP FROM P(1)) = 1
```

The leaking request blocks were returned on disconnect.

Solution — D. Yemanov

Fixed

There was a problem with deadlock detection when pessimistic locking (WITH LOCK syntax) was used. — Not logged

```
create table test (id integer);
insert into test values(1);
insert into test values(2);
Commit;
```

Transaction 1 (READ COMMITTED, WAIT):

```
select * from test
  where id = 1 with lock;
```

Transaction 2 (READ COMMITTED, WAIT):

```
select * from test
  where id = 2 with lock;
select * from test
  where id = 1 with lock;
```

Transaction 1:

```
select * from test
  where id = 2 with lock;
```

This set of conditions would result in a permanent deadlock.

Solution — N. Samofatov

Fixed. V.1.5.2 detects and reports such a deadlock as an error.

The server would crash when NULL was passed to EXECUTE STATEMENT ... INTO. For example, — Not logged

```
...
VAR = NULL;
EXECUTE STATEMENT :VAR;
...
```

caused the server to die.

Solution — A. Peshkoff

Fixed

The server log was polluted with SIGPIPE errors when running SuperServer on UNIX.The legacy InterBase code was logging sigpipe errors for SS running on *nix. Unfortunately sigpipe errors come in their thousands (when they come at all) with the result that the log filled up rather quickly. In extreme cases this led to filling up the entire partition. — Not logged

Solution — P. Reeves

Logging of SIGPIPE errors has been disabled.

100% CPU usage was exhibited by the cache_writer thread in some rare cases (reported by Adrianos dos Santos Fernandes). To reproduce, open two command prompts. — Not logged

prompt1:

```
isql
CREATE DATABASE 'test.fdb';
CREATE TABLE T (N INTEGER);
EXIT;

gbak -B test.fdb test.fbk
del test.fdb
gbak -C test.fbk test.fdb
```

prompt2:

```
isql test.fdb
```

prompt1:

```
gbak -B test.fdb test.fbk
```

The server would consume 99% of CPU until the isql t prompt2 was disconnected. The bug didn't occur when passing -GARBAGE_COLLECT in the last command.

Solution — V. Horsun

Fixed

A possible source of server crash was discovered in the op_connect handler. When a TCP/IP packet lacking user information) was received on the server port, the server could crash. Because it was the first packet (op_connect) in the client-server protocol, it exposed the server to any kind of DoS attack. Anyone could kill the server with just one TCP packet.	Not logged
Solution Fixed	A. Karyakin, D. Yemanov

The server could crash with complex queries where lots of streams were used in a sort/merge. A complex union with many aggregations and merge joins could crash the server because of a streams buffer overflow. Although the current limit is 255 streams per request, the temporary buffer could accommodate only 128 items.	Not logged
Solution Fixed	D. Yemanov

The server was blocking when events were used with Network Address Translation (NAT) gateways. Auxiliary connections (for events) were established by the client library using the server-reported TCP/IP address. But the address returned by the server may be incorrect if it is behind a NAT box.	Not logged
Solution The fix was to use the address that was used to connect the main socket, not the address reported by the server.	C. Waters, D. Yemanov

---

Sweeper would not release its lock when database shutdown was executed. A server crash could occur when a database shutdown was initiated while the sweep is being in progress.	Not logged
Solution Fixed	V. Horsun

The least significant bits of a floating-point value would be lost when rounding the value to an integer or int64 value. Dialect 3 database:	Not logged

```
SELECT CAST(CAST( 1.005E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.015E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.025E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.035E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.045E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
```

```
UNION ALL
SELECT CAST(CAST( 1.055E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.065E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.075E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.085E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
UNION ALL
SELECT CAST(CAST( 1.095E0 AS NUMERIC(15,2))
  AS VARCHAR(30)) FROM RDB$DATABASE
```

FB 1.5.1 returns

```
F_1
----
1.00
1.01
1.03
1.04
1.05
1.05
1.06
1.08
1.09
1.10
```

Solution — V. Horsun

Fixed. FB 1.5.2 returns

```
F_1
----
1.01
1.02
1.03
1.04
1.05
1.06
1.07
1.08
1.09
1.10
```

A few memory access problems were detected when testing HEAD under Valgrind. — Not logged

Solution — N. Samofatov

HEAD fixes were back-ported to v.1.5.2

64-bit SuperServer builds on platforms such as Linux/AMD64/NPTL, which use the high-order bits of a 64-bit thread ID, were exhibiting run-time errors. — Not logged

Solution — N. Samofatov

Fixed

CURRENT_TIMESTAMP was yielding unpredictable results on 64-bit platforms.	Not logged
Solution Fixed	N. Samofatov

M I N O R E N H A N C E M E N T S

ISSUE	SF Bug # Fixed by
Performance improvement for permissions checking	n/a
Solution Resource lists to check permissions are now computed on the fly as needed. For complex schemas, this significantly reduces memory and CPU time consumption.	N. Samofatov, D. Urban
POSIX build and packaging changes	n/a
Solution 1. Work around bugs in GCC 3.3.2 and 3.3.3 2. Support GCC 3.4 build 3. Limit exports of Firebird libraries using version script 4. Link client library and UDF libraries with POSIX threads. This cures problems with single-threaded hosts like PHP linking with libfbclient.so from CS packages	N. Samofatov
More POSIX build and packaging changes	1027636
Solution ● To prevent the startup status from being overwritten by the next status message, the /etc/init.d/firebird script needed to have a line consisting only of "echo" after RETVAL=$?. ● Erik LaBianca extended the Firebird build system to generate source bundles in a generic fashion and without autoconf dependency. He uses this facility for his Fedora Core packages.	D. Mullins E. S. LaBianca
Changes to the standard ib_udf library declaration script	n/a
Solution The default declarations of the string manipulation routines in ib_udf.sql were altered to accept strings with lengths up to 255 characters	N. Samofatov

64-BIT RELEASES

by Nickolay Samofatov

A minimal number of changes needed to produce fully functional 64-bit builds has been back-ported to the Firebird 1.5.2 series from the Firebird 2 development tree. The resulting builds are fast and reported to be usable in production environments. Wire protocol compatibility is fully retained: 32-bit clients can talk with 64-bit servers and vice versa.

> **However, the general recommendation is to proceed with caution when deploying these builds.**

Known issues

1. The types of Public API handles are expected to change in the 2.0 series from 64-bit pointers to 32-bit integers. This means that 64-bit client applications will have to be recompiled to work with Firebird 2.0 client libraries.
2. Because the 64-bit builds must use 64-bit data alignment, the On-Disk Structure (ODS10) is currently not the same for the 64-bit and 32-bit builds.

   This situation is expected to be resolved in Firebird 2.0 which is expected to have the same, 64-bit On-Disk Structure (ODS11) for both 32-bit and 64-bit versions of the engine.
3. Some things are known to be still pretty flaky in the Linux x64 (aka AMD64) world. Many 2.6 kernels, including the latest at the time of this writing (2.6.9) have 64-bit Interprocess Communications (IPC) facilities broken in one way or another. This is especially true if you have a SMP or NUMA machine.
4. Before reporting a problem related to server hangs or transient lock-ups, try
   - downgrading the kernel to the 2.4 version (SMP kernels are known to work fine)

     or
   - test with the non-preemptible 2.6 kernel, with SMP/NUMA support disabled.
