

Category: Core Engine

ID	SF ID	Description	Bug Group/Status
216	212129	<p>Server doesn't support NFS/mapped paths</p> <p>If you try to access a database file that's mounted via NFS, but you specify its path in the filesystem as if it were local, InterBase would normally change that into a network request to prevent file corruption by multiple clients.</p> <p>Example: Pretend our system mounts /home from the root directory of a machine named simpsons. Now pretend you request to connect to the database /home/marge/groceries.gdb. InterBase <code>_should_</code> turn the request into a connection to <code>simpsons:/marge/groceries.gdb</code> and let simpsons handle reads and writes to the file on its own.</p> <p>The FreeBSD port of InterBase doesn't do that. So don't go connecting to databases mounted via NFS unless you know what you're getting into or file corruption doesn't bother you a bit.</p>	As Designed/Pitfall
220	212340	<p>Token unknown in simple SELECT with GROUP BY and ORDER BY</p> <p>This simple SELECT statement does not work</p> <pre>select country, count(country) from customer group by country order by count(country)</pre> <p>I'm getting SQL error code = -104, Token unknown count.</p>	Confirmed Bug
225	212899	<p>skywalker commit for jrd/cmp.c rev 1.3 buggy</p> <p>http://cvs.sourceforge.net/cgi-bin/cvsweb.cgi/interbase/jrd/cmp.c.diff?r1=1.2r2=1.3cvsroot=firebird</p>	Fixed v0.9
231	213708	<p>-502 Declared cursor already exists</p> <p>When connecting to IB60, microfocus COBOL programs receive:</p> <p>Dynamic SQL Error -SQL error code = -502 -Declared cursor already exists</p> <p>Host system AIX 4.2.1.0</p> <p>The program fails when connecting: - locally to IB60 on AIX (CLASSIC configuration) - remotely to IB60 on WINNT (SUPERSERVER configuration)</p>	Fixed v0.9
233	214298	<p>Select count(*) expression anomaly when table is empty</p> <p>Confirmed in v. 5.6, not tested in v. 6:</p> <p><code>select count(*) + 1</code> returns zero when table has no rows but returns correct result if <code>count(*)</code> 0.</p>	Fixed v0.9
235	216464	<p>IB6 can't connect database placed in non-ASCII directory</p> <p>Can't connect to database placed in directory contains non-ASCII characters, such as russian WIN1251.</p>	As Designed/Pitfall
236	216579	<p>generators in computed by columns will return wrong results</p> <p>Using generators in computed by columns will return wrong results and create an unusable database.</p> <pre>create table t0 (a integer, genid_field computed by (a + gen_id(gen1, 1))); show table t0; insert into t0(a) values(10); insert into t0(a) values(12); select * from t0;</pre>	Fixed v0.9
237	216733	<p>Too Many Generators Can Corrupt Database</p> <p>The number of generators you can have is dependant on (page size - unknown overhead) / size of generator. IB allows you to create generators past this limit with no complaint, but these generators will return random data and corrupt the database if incremented.</p> <p>IB seems to limit generators to one page, but no range checking is done. This is particularly bad on databases with small page sizes which migrate from ODS 9 to ODS 10, since the size of generates doubles from 32 bit to 64 bit, seriously reducing the limit. On a 1024 page size, this limit is somewhere less than 128 generators.</p>	Fixed v0.9
244	221589	<p>numeric fields and mathematical operations</p> <p>example: <code>select field1 * field2 from mytable</code> or <code>select field1 * (1+field2/100) from mytable</code></p> <p>and both fields are numeric type (interger) like <code>numeric(9,2)</code>, the result are incorrect.</p>	Fixed v1.0

Category: Core Engine

ID	SF ID	Description	Bug Group/Status
252	222476	<p>Avg and sum return empty field names in dialect 3</p> <pre>select avg(1), sum(1) from rdb\$database</pre> <p>If we go to dialect 1, things work as expected: the first field is named AVG and the second, SUM.</p>	Fixed v0.9
254	223056	<p>Blob-IDs are sometimes shared between more rows</p> <p>In SP, when you copy record (from/to the same table) by <pre>INSERT INTO tab SELECT ... FROM tab WHERE ...;</pre> or by <pre>SELECT ... FROM ... INTO _local_variables_;</pre> <pre>INSERT INTO tab VALUES (_local_variables_);</pre> and the record contains BLOB, then _sometimes_ newly created row will not contain its own copy of the BLOB, but instead it will use the same blob-id as the original record. (i.e. single blob is shared among more rows)</p> <p>This is really severe bug, because it will cause data lost - when you delete one of these rows, and then try to read the other one you will get BLOB not found error !!!</p> <p>Interestingly, when you execute exactly the same command directly, not as part of SP, blob-ids will be o.k.</p>	Confirmed Bug
256	223059	<p>Updating VARCHAR does not clear old data</p> <p>When IB updates VARCHAR string, it does not zero rest of the string. (when the row is decompressed in the buffer).</p> <p>e.g. you have table <pre>CREATE TABLE t (a VARCHAR(50));</pre> <pre>INSERT INTO t</pre> <pre>VALUES ('abcdefghijklmnopqrstuvwxy1234567890abcdefghijklmnop');</pre> <pre>COMMIT;</pre> <p>When you update it by <pre>UPDATE t SET a='XYZ';</pre> then IB creates new version of the row that should contain string: 'XYZ' + zero filled rest but it stores (in gdb file) this string instead: 'XYZabcdefghijklmnopqrstuvwxy1234567890abcdefghijklmnop'</p> <p>Because VARCHARs contain length of the string, client application will never notice any problem (i.e. it will always get correct result), but the gdb file can grow faster than expected (because such additional data can't be rle compressed), and database can get slower (because less useful data fit onto page).</p> </p>	Fixed v1.0
262	223512	<p>DROP VIEW shouldn't drop a table.</p> <p>Use CONNECT or CREATE DATABASE to specify a database SQL create database 'z0.fdb'; SQL create table t(a int); SQL create view v as select a from t;</p> <pre>/* Incorrect behavior, according to SQL standards.</pre> <pre>A view cannot be dropped as a table and vice versa. */</pre> <pre>SQL drop table v;</pre> <pre>SQL drop view t;</pre>	Fixed v1.0
267	223793	<p>isc_add_user() allows adding 32-char usernames</p> <p>In the file interbase\jrd\alt.c, function isc_add_user() (and also isc_modify_user, isc_delete_user) there is line <pre>if (strlen (user_data-user_name) 32)</pre> so this function allows adding usernames 32 characters long!</p>	Confirmed Bug
278	226456	<p>SELECT/PLAN does not understand delimited SQL index names</p> <p>When I have an index, for example: idx_asc_History_Stamp and try to use this index in my SELECT statement as shown below: <pre>... PLAN (History ORDER idx_asc_History_Stamp) ... then I get an error with ISC ERROR CODE = 335544343:</pre> <pre>invalid request BLR at offset 112</pre> <pre>there is no index IDX_ASC_HISTORY_STAMP for table History</pre> <p>But if I have an index named IDX_ASC_HISTORY_STAMP or both of them, this SELECT statement with my plan works fine, but uses the upper-cased one.</p> </p>	Fixed v1.0

Category: Core Engine

ID	SF ID	Description	Bug Group/Status
279	226614	<p>So it looks like PLAN cannot recognize an index name, which is SQL delimited identifier.</p> <p>temp buffer in FUNCTIONS_entrpoint too small</p> <p>The temp buffer in FUNCTIONS_entrpoint is 128 characters, but when it is called by ISC_lookup_entrpoint which is called by obj_init (intl.c) it is passed a path with MAX_PATH_LEN 128, and a name with no length limit (currently 20). If the length of the path and name were less than 128, everything was fine. If it was greater you get a segfault.</p>	Confirmed Bug
280	227375	<p>Grouping on derived fields processing NULL data kills IB</p> <p>The database has the following table:</p> <pre>CREATE TABLE TWODATE(TWODATEID INTEGER NOT NULL, DATEBEGIN DATE, DATEEND DATE, CONSTRAINT PK_TWODATE PRIMARY KEY (TWODATEID));</pre> <p>The table contains several records. Some of the values in the two date fields are null. Furthermore the database has the following view:</p> <pre>CREATE VIEW CALCDIFF (TWODATEID, DIFFYEAR) AS select TwoDateID, extract(year from datebegin) - extract(year from dateend) from twodate;</pre> <p>The following Select statement causes a lost connection to database - error:</p> <pre>SELECT DiffYear, count(*) FROM CalcDiff Group by DiffYear;</pre> <p>The error occurs only, when</p> <ul style="list-style-type: none"> - there are null values in the database - there is a subtraction between the two extract statements - the select has a group by statement 	Fixed v0.9-5
285	227758	<p>Field names with spaces cannot be used in VIEWS</p> <p>The following DDL exemplifies more completely the problem I've been encountering with accessing views.</p> <pre>CREATE DOMAIN IDINTEGER AS INTEGER NOT NULL; CREATE DOMAIN IDVARCHAR AS VARCHAR(31) NOT NULL; CREATE TABLE Company List (Company Name IDVARCHAR NOT NULL PRIMARY KEY, Company ID IDINTEGER NOT NULL UNIQUE); CREATE TABLE Vendor List (Company ID IDINTEGER NOT NULL PRIMARY KEY, Days to Quote Expiration SMALLINT); /* The following view is inaccessible.*/ CREATE VIEW Vendor Name List (Company ID, Company Name) AS SELECT Company ID, Company Name FROM Company List CL WHERE EXISTS (SELECT * FROM Vendor List VL WHERE VL.Company ID = CL.Company ID);</pre> <p>The message I got is:</p> <pre>Cannot access column Company Name in view Vendor Name List Statement: select * from Vendor Name List</pre> <p>And I verified that IB indeed stored the definition as it appears above, so this is not a problem of IBConsole or IB_WISQL. This is bug in IB: views based on tables whose fields carry spaces in dialect 3 don't work because they can't access those table's fields.</p> <p>/* However, the equivalent DDL with the spaces removed allows access to the view. */</p> <pre>CREATE TABLE CompanyList (CompanyName IDVARCHAR NOT NULL PRIMARY KEY, CompanyID IDINTEGER NOT NULL UNIQUE); CREATE TABLE VendorList (CompanyID IDINTEGER NOT NULL PRIMARY KEY, DaystoQuoteExpiration SMALLINT);</pre>	Confirmed Bug

Category: Core Engine

ID	SF ID	Description	Bug Group/Status
		CREATE VIEW VendorNameList (CompanyID, CompanyName) AS SELECT CompanyID, CompanyName FROM CompanyList CL WHERE EXISTS (SELECT * FROM VendorList VL WHERE VL.CompanyID = CL.CompanyID);	
286	227760	Zero-length db object names shouldn't be allowed Zero-length field names aren't valid and should be banned. Currently, IB in dialect 3 allows constructions like these: create table (int); commit; create index on (); set term ^; create procedure p returns(int) as begin for select from into : do begin = : * ; suspend; end end ^ set term ^; commit; create role ; grant select on to ; grant execute on procedure p to ; insert into values(1); insert into values(2); insert into values(3); commit; Then, select * from p works perfectly.	Fixed v1.0
291	228467	security bug with a hardcoded user with full rights to isc4 Look at this URL: https://www.kb.cert.org/vuls/id/247371	Fixed v0.9-4
295	229121	TEMP directory filling up This has started to happen since moving from Interbase 6 to firebird 0.94-1 on Linux. Same hardware, same data, same application. There are no visible files in the /tmp direcorey, but this is the output you get from lsdf: COMMAND PID USER FD TYPE DEVICE SIZE NODE NAME gds_inet_16911 interbase 5u REG 8,8 4160000 14 /tmp/gds_sort__0ckHHn (deleted) etc....	Fixed v0.9-4
296	229231	revoke is case sensitive If I create a user called admin and then grant all privileges on a table to them, the revoke command doesn't work unless I use the admin name in upper case. Here is an example: grant all on config to admin; revoke all on config from admin; You would think this would result is admin having no privilege on config but a select on the RDB\$USER_PRIVILEGE table shows the admin user still has full access! You need to do: revoke all on config from ADMIN in order to remove the privileges!	Fixed v0.9-5
1,350	421260	Character length not filled for UDFs The case for ODS10 was only written as an stub.	Fixed v0.9-5

Category: Core Engine

ID	SF ID	Description	Bug Group/Status
1,480	450405	<p>Tricky role defeats basic SQL security</p> <p>In this example t t is any table. Here, it doesn't matter that its name has an embedded space, it could be named T0 or mytable as well. SYSDBA does:</p> <pre>SQL create role for cvc; SQL create role for; SQL grant for cvc to user cvc; SQL grant select on table t t to for; SQL ^Z</pre> <p>Now, cvc has ONLY the role for cvc.</p> <p>However, sysdba assigned role for the right to select from the table. No user has been granted such role. Also, no rights have been given to the role granted to cvc, so it's useless.</p> <p>If cvc tries to use role for, it will be changed internally to NONE because that user hasn't been granted such role. However, cvc can play with the legitim assigned role and do:</p> <pre>H:\ibdev\fbbuild\interbase\dsqllsql Use CONNECT or CREATE DATABASE to specify a database SQL connect c:/proy/fbtest.gdb user cvc password pw role for cvc; Database: c:/proy/fbtest.gdb, User: cvc, Role: for cvc SQL select * from t t; SQL</pre> <p>Whoops! The engine allowed the request! Internally when the ACL is checked, the ACL and the role are compared at some time, and since the ACL is length-counted, it will stop at the third position, since the role that has legitim rights is for. At this time, check_string() queries the next character in the current role. In the case of for cvc, it's a blank, so the routine decides that both roles are the same and accept the credentials. It didn't check that there may be more valid characters after a couple of embedded blanks.</p> <p>The example is convoluted, but relying on the admin doing always the sensible decision is risky.</p>	Fixed v1.0
1,362	428889	<p>alter column col position pos zero-based</p> <p>The syntax</p> <pre>alter table tbl alter column col position pos;</pre> <p>is zero-based. Since it's SQL command, it should be in sync with SQL positions starting at one.</p>	Fixed v1.0
1,357	425799	<p>Renamed domain leaves behind dimensions</p> <p>With a domain being an array, renaming the domain causes rdb\$field_dimensions to be left unchanged; hence the connection between a domain and its dimensions specification is broken.</p> <p>Example:</p> <pre>create domain dunno int[0]; commit; alter domain dunno to ditto; commit; select * from rdb\$field_dimensions;</pre> <p>= will show that the referenced domain is still DUNNO instead of DITTO. System tables are out of sync.</p>	Fixed v1.0
1,346	419964	<p>buffer overflow in remote/interface.c li</p> <p>The version string is too long for the sprintf on line 933. Here is the string: UP-T0.9.4.101 Firebird Test1/tcp (cse-air-dhcp-153.ucsd.edu)/P10. The buffer allocated is only 64 bytes long. Looking at the string I feel the buffer needs to be bigger. The version includes the dns name, which is in no way constrained to 64 characters.</p>	Fixed v1.0
1,319	233124	<p>Connection lost during the bad SQL code execution</p> <p>If I try to alter domain with a bit wrong syntax like:</p> <pre>alter domain domain_name set type smallint;</pre> <p>where 'set' is illegal, the connection between server and client is lost with the following well-known error:</p> <pre>Unable to complete network request to host host_name. Error writing data to the connection. unknown Win32 error 10054</pre>	Fixed v1.0

Category: Core Engine

ID	SF ID	Description	Bug Group/Status
		The problem seems to be the client's one. Maybe it is not very serious bug, but any request should have predictable result,	
1,464	447377	GDS error ...can't find tip There is a bug in InterBase 5.6, 6.01, and the current Firebird 0.9-5 beta that causes the lookup of a transaction inventory page to fail if there are more than 32767 transaction pages. That makes the maximum safe transaction id for a database with: 1024 byte pages 131,596,287. 2048 byte pages 265,814,016. 4096 byte pages 534,249,472. 8192 byte pages 1,071,120,384. Although those are large numbers, there was a particular database exceeded 131 million transactions in six months. Attempts to attach to the database failed with the error gds internal consistency check, can't find tip. Suggestions: 1) don't use a 1024 byte page size. 2) do check your next transaction number from time to time. 3) if you see the next transaction number approaching the limit, backup and restore the database.	Fixed v1.0
1,478	450301	SUBSTRING doesnt work The syntax for SUBSTRING is: SUBSTRING(s FROM start [FOR length]) s is a string, start is 1-based. The command should extract from the string s starting at position start a substring with at most length chars. If length isn't given, then the command should extract the full string after the start position. The SUBSTRING keyword isn't available in early firebird-versions (e.g. 0.9.4) or in the interbase-versions. I have tested the command in FireBird 1.0.0 Beta 2, Build 324 Classic on Linux 2.2.16 (SuSE 7.0). When I use the command in the where-clause (e.g. with in or =) it doesn't work, when I use it with string-concatenation, it doesn't work, either.	Fixed v1.0
1,359	425949	Engine CRASH Error select count(*), adresy.rdb\$db_key from adresy adresy-this is any table	Fixed v0.9-5
1,340	414833	Join Procedure Bug Create a test Database using script //-----// SET NAMES WIN1251; CREATE DATABASE 'D:\TEST.GDB' USER 'SYSDBA' PASSWORD 'masterkey' PAGE_SIZE 2048 DEFAULT CHARACTER SET WIN1251; CREATE TABLE TEST_TABLE (ID INTEGER NOT NULL, DATEBEGIN DATE, DATEEND DATE); INSERT INTO TEST_TABLE (ID, DATEBEGIN, DATEEND) VALUES (1, '01/01/2001 00:00:00', '02/01/2000 00:00:00'); COMMIT WORK; SET TERM ^ ; CREATE PROCEDURE TEST1 (BEGIN_DATE DATE, END_DATE DATE) RETURNS (RESULT INTEGER) AS begin RESULT=0; suspend;	Initial Bug

Category: Core Engine

ID	SF ID	Description	Bug Group/Status
		end ^ SET TERM ; ^ and try to execute following statement Select h.datebegin,h.dateend from test_table h, TEST1(h.datebegin,h.dateend) g	
1,336	412201	Trigger update limit I have a trigger which autogenerate a reference. When I exchange data between two databases (replication) I need to deactivate this trigger when there is insert clauses . I reactivate it after the exchange. After 255 alter trigger XXX active/incative I have a two many version problem and I need to backup/restore database....	As Designed/Pitfall

Category: DSQL

ID	SF ID	Description	Bug Group/Status
292	228526	ambiguous statements return unpredictable results I noticed that IB happily executes an ambiguous query of the form: select ... from orders o left join customers c1 on (o.customer_id1 = c1.id) left join customers c2 on (o.customer_id2 = c2.id) where somefield = 'somevalue' Assuming somefield exists in both the customers and the orders table, the query is ambiguous unless the reference in the where clause is qualified.	Fixed v1.0
1,337	412417	altering from CHAR to VARCHAR Altering CHAR to VARCHAR column adds 2 bytes to field length. How to reproduce: CREATE TABLE TEST(N CHAR(40)); then ALTER TABLE TEST ALTER N TYPE VARCHAR(40); after that RDB\$FIELD_LENGTH will be 42, not 40 as it should be. RDB\$CHARACTER_LENGTH is OK, i.e. it stays 40 after ALTER. The '40' is not a magic number - same addition of 2 bytes will be for any altering from char to varchar. For example, altering char(40) to varchar(50) will give the resulting varchar field 52 characters length. Real varchar storage size is 2 bytes greater than char. But, column length definition (RDB\$FIELD_LENGTH) does not store physical field length - it stores column length. So, altering char to varchar must keep the same column size.	Fixed v1.0

Category: Data Types/On-Disk Struc(ODS)

ID	SF ID	Description	Bug Group/Status
1,316	231998	space before CASTed numeric expression in dialect 1 Andrew Velikoredchanin found interesting behavior in dialect 1, when numeric expression is being casted to a varchar. For example: select cast(22 / 7 as varchar(20)) from rdb\$database will give ' 3.142857142857143' as a result (note that there is a space before first digit).	Fixed v1.0

Category: Data Types/On-Disk Struc(ODS)

ID	SF ID	Description	Bug Group/Status
		If not an expression, but numeric value is casted, all is ok, i.e. no space before casted result.	
		Dialect 3 does not have this problem.	
		no difference between integer and numeric values - 22.0 / 7.0 will give space before first digit.	
1,320	233304	computed field and TIME datatype create table a (workstart time, workend time, duration computed by (workend - workstart)); field duration will be of numeric type, not TIME as it supposed to be.	As Designed/Pitfall

Category: Documentation

ID	SF ID	Description	Bug Group/Status
266	223789	Ib 6.0.1 and UDF ascii_char error If you declare the function ascii_char like stated in the documentation or in the example provided in InterBase\examples\Udf\ib_udf.sql, i.e.: DECLARE EXTERNAL FUNCTION ascii_char INTEGER RETURNS CHAR(1) FREE_IT ENTRY_POINT 'IB_UDF_ascii_char' MODULE_NAME 'ib_udf'; if you use it, i.e.: select ascii_char(65) from rdb\$database you will receive the error: -arithmetic exception, numeric overflow, or string truncation - That's because the declaration has to be different: DECLARE EXTERNAL FUNCTION ascii_char INTEGER RETURNS CSTRING(1) FREE_IT ENTRY_POINT 'IB_UDF_ascii_char' MODULE_NAME 'ib_udf';	Confirmed Bug

Category: FreeBSD port

ID	SF ID	Description	Bug Group/Status
238	216778	FreeBSD port should be linked with -ldescript FreeBSD port should be linked with -ldescript, not -lcrypt, because libcrypt.so can be symlink to libcrypt.so or it should have DES inside.	Fixed v0.9

Category: GPRE

ID	SF ID	Description	Bug Group/Status
284	227717	COBOL programs randomly return a -901 request sync. error COBOL programs randomly return a -901 request synchronization error. Alternate title: COBOL programs lose SQLCODE values during UPDATE. During 'UPDATE SET ... WHERE x=..' ('Mass Update') error codes returned by the update are not returned to the program, instead the program will see a -901 request synchronization error. The error is caused by bad code generated by GPRE.	Fixed v0.9

Category: IBGuardian

ID	SF ID	Description	Bug Group/Status
215	211790	Year 2000 incompliance in guardian Guardian window shows dates of 2000 year as dd/mm/100 instead of dd/mm/00.	Confirmed Bug
219	212328	iscGuard (ibguard) still leaks handles An internal test release of IBGuard with the name of iscGuard (by Mike) presents a decent date and time formatting in the crash report. No more Y2K problems in the report (IBGuard's properties). This test program closed two handle leaks, but there are still two more, probably produced when IB crashes and it's restarted by the Guardian, because the old handles are not released.	Fixed v0.9

Category: IBGuardian

ID	SF ID	Description	Bug Group/Status
		(IMHO, some parts of IBGuard that I reviewed through the web interface, seems to have been done by a sophomore and not by a professional developer.)	

Category: ISQL

ID	SF ID	Description	Bug Group/Status
218	212263	command line isql ignores -user / -password with -a or -x Workaround is to set ISC_USER / ISC_PASSWORD beforehand, or use unix authentication	Fixed v0.9-5

To reproduce (using windows):
unset ISC_USER and ISC_PASSWORD
isql -x {some database} -user {some valid user} -password {some valid password}

This appears to be a logic bug, rather than an option parsing bug. In isql.e -a and -x are handled differently from other cases which call newdb from do_isql

253	222563	isql extracts wrong sproc's parameters with UNICODE This is a bug involving metadata extraction. Given the following declaration:	Fixed v0.9
-----	--------	---	------------

```
CREATE PROCEDURE p
  (inparam CHAR(10) CHARACTER SET unicode_fss)
RETURNS
  (outparam CHAR(10) CHARACTER SET unicode_fss)
AS
DECLARE VARIABLE
  var CHAR(10) CHARACTER SET unicode_fss;
BEGIN EXIT; END
```

The macro-command SHOW PROCEDURE P will output:
Parameters:
INPARAM INPUT CHAR(30) CHARACTER SET UNICODE_FSS
OUTPARAM OUTPUT CHAR(30) CHARACTER SET UNICODE_FSS

As you can see, CHAR(10) becomes CHAR(30), because the tool is reading rdb\$field_length instead of rdb\$character_length for procedure's parameters. In the case of table's fields, the correct information is read and presented. The following command confirms that the engine itself is doing the right thing:

```
SELECT rdb$field_length, rdb$character_length
FROM rdb$fields
WHERE rdb$field_name IN (
  select rdb$field_source from rdb$procedure_parameters
  where rdb$procedure_name = 'P')
```

```
rdb$field_length rdb$character_length
=====
          30          10
          30          10
```

282	227473	isql always ignores charset NONE in metadata extraction	Confirmed Bug
-----	--------	--	---------------

This is a rather subtle bug that might cause a db to change its original meaning if the metadata is extracted as an script and later a new db is created from such script. It seems that isql was planned with charset NONE in mind.

Example, let's assume a database whose default character set is not NONE but WIN1250, namely, one whose rdb\$database's rdb\$character_set_name field is WIN1250 instead of NULL. Then given this declaration,

```
CREATE TABLE none(n char character set none);
```

ISQL extracts information as:

```
SET SQL DIALECT 3;
CREATE DATABASE 'D:\users\cvalde\proy\1252.gdb' PAGE_SIZE 1024 DEFAULT CHARACTER SET WIN1252;
[snip]
```

```
/* Table: NONE, Owner: SYSDBA */
CREATE TABLE NONE (
  N CHAR(1));
```

So, it's evident isql swallowed the charset specification, because it's NONE, so it assumed NONE is always the default. However, next time the script is submitted through the same isql.exe, since the db-wide charset is WIN1252, the table N will have a field

Category: ISQL

ID	SF ID	Description	Bug Group/Status
		named N whose charset will be WIN1252 and not NONE as initially created, because fields without charset specification assume the charset of the whole db. Solution: isql shouldn't discard charset NONE blindly when extracting metadata. The safe way is to query rdb\$database and skip (for being redundant) the fields that use the same charset than the whole db (and only if rdb\$database contains no charset specification, the utility can behave as it behaves currently) so the script can rebuild the db as it was originally, without subtle changes. This bug affects not only isql and IBConsole but third party tools, too.	
1,351	421262	ISQL reports UDF BLOB parameter BY VALUE It's obviously wrong. BLOBs can't be passed by value and a mechanism should not be specified in a script. Internally, BY DESCRIPTOR is used. This is only a bug when reporting metadata information.	Fixed v0.9-5
1,475	448613	ISQL corrupts memory Simple ISQL commands, such as: CREATE DATABASE c:\home\bar.gdb user builder password builder; cause the debug build to GPF in gds__free on the win32 platform. This seems to be caused by corruption of the client memory pool. There is a small chance that this is a client library bug.	Initial Bug

Category: Installation - Server

ID	SF ID	Description	Bug Group/Status
268	224037	Error in Install !!! I install Firebird in C:\WP\Firebird\ After install in registry: ERROR!!!! ERROR!!!!ERROR!!!!ERROR!!!!ERROR!!!! REGEDIT4 [HKEY_LOCAL_MACHINE\SOFTWARE\BORLAND\InterBase\CurrentVersion] RootDirectory=C:\Program Files\Firebird\ ServerDirectory=:\Program Files\Program Files\ PREDIC: PREDIC: PREDIC: PREDIC: PREDIC: PREDIC: PREDIC: REGEDIT4 [HKEY_LOCAL_MACHINE\SOFTWARE\BORLAND\InterBase\CurrentVersion] RootDirectory=C:\WP\Firebird\ ServerDirectory=:\WP\Firebird\ 	Confirmed Bug

Category: Linux ports

ID	SF ID	Description	Bug Group/Status
283	227647	the disk space is not freed after the sort file is deleted Concerns: Firebird 0.9 SS and CS The disk space is not freed (reported to the OS as free) after the sort file is deleted in the temp directory until the whole task finishes successfully (like gbak restore) or the IB server restarts. As a result running a gbak restore on 600MB db produces 1GB sort files and exhausts the available disk space.	Confirmed Bug
1,365	430311	Firebird under Redhat 7.x Firebird 9.4xxx cannot be run in redhat 7.0. i had try to install on a redhat 7 both copies did not fire up properly. it did start for a while, but then shutdown by itself. when I do a ps -ae i can see that there are a few ibserver been started. but all process did not run. because the cpu usage is 0%. is it not compactable with redhat7 or what??	As Designed/Pitfall

Category: Security Issues

ID	SF ID	Description	Bug Group/Status
1,342	416477	Malfunction of permissions/privileges In V6.0, we will grant permissions for a procedure and ONLY execution of the procedure to a user (users don't have direct access to data). In new V6.0.1, is not possible to the user's schema to execute procedures or select views, because the user must have ALL permissions on tables (and execute permissions on procedure) to execute a procedure. The documentation say in Data Definition Guide, page 205: 'Tip: As a security measure, privileges to tables can be granted to a	As Designed/Pitfall

Category: Security Issues

ID	SF ID	Description	Bug Group/Status
		procedure instead of to individual users. If a user has EXECUTE privilege on a procedure that accesses a table, then the user does not need privileges to the table.'. This is not possible in V6.0.1.	

Category: UDF/Built-In Functions

ID	SF ID	Description	Bug Group/Status
1,352	421263	UDF substr gives NULL if slice input The substr UDF that comes with FB returns NULL when the final position is greater than the last position in the input string argument. This causes any kind of problems. Hence, two solutions have been provided: - Fix substr so it will return NULL when provided with NULL, but it will report the full string if the final position of the slice is greater than the strig. - Add substrlen that will have the argument most people are used to: starting position plus length of the slice instead of starting and final position.	Fixed v0.9-5