

Category: Build Issues

ID	SF ID	Description	Priority
1380	440543	Cannot create tables 0.9.5.260	5 - Medium (next major release)

After installing Linux SS version 0.9.5.260 on Mandrake 8.0, I get the following error when creating new tables, both via local ISQL and thru Windows IBConsole/EMS QuickDesk clients:

COMMAND:
create table T1(F1 INTEGER NOT NULL PRIMARY KEY);

ERROR:
unsuccessful metadata update
-STORE RDB\$RELATIONS failed
-attempt to store duplicate value (visible to active transactions) in unique index RDB\$INDEX_1
-null segment of UNIQUE KEY

Category: Client/GDS32

ID	SF ID	Description	Priority
1376	436462	Rows affected incorrect with trigger	5 - Medium (next major release)

I'm having troubles with updatable views that have before update triggers. When i update just one record (using the PK), it returns 3 rows affected.

When i update this view using IBO it gives the error: 'MULTIPLE ROWS AFFECTED', because the number of rows affected is giving 3 rows!

PS: I can't put the before update in the table, because i have several view on one table and i expect that each view have an specific behavior.

Try this:

```
/* CREATE A TABLE */
create table test (
  test_id integer not null primary key,
  name varchar(30),
  address varchar(30) );

/* CREATE A VIEW */
create view v_test as
select test_id, name, address from test;

/* INSERT A RECORD */
insert into v_test values (1, 'name1', 'address1');
/* ROWS AFFECTED: 1 */

/* UPDATE THIS RECORD */
update v_test set name = 'name1' where test_id = 1;
/* ROWS AFFECTED: 1 */

/* CREATE A TRIGGER FOR THE VIEW */
set term !! ;
create trigger trg_v_test for v_test
before update
as
begin
  new.name = 'name1';
end !!
set term ; !!

/* UPDATE THE RECORD AGAIN - SAME AS ABOVE*/
update v_test set name = 'name1' where test_id = 1;
/* ROWS AFFECTED: 3 !!!!!!!!!!!!!!!!!!!!!!!!!!!!! */
```

Category: Core Engine

ID	SF ID	Description	Priority
260	223133	Ambiguous self join produce bizarre results	5 - Medium (next major release)

Doing a self join without assigning alias to both occurrences of the table can be a fatal trap for newbies. I will use a

Category: Core Engine

ID	SF ID	Description	Priority
		<p>contrives, useless example:</p> <pre>select rdb\$relation_name from rdb\$relations r join rdb\$relations on rdb\$relation_id = r.rdb\$relation_id order by 1 = produces a listing of tables, ordered as expected. The first table in the sample db is called A.</pre> <p>Observe that only one occurrence of rdb\$relations is qualified with an alias. Now, by simply assigning the alias to the second occurrence instead of the first one, we force the server to yield incredible results:</p> <pre>select rdb\$relation_name from rdb\$relations join rdb\$relations r on rdb\$relation_id = r.rdb\$relation_id order by 1 = produces A A A</pre>	
1477	449312	<p>Create External Table with relative path segfault</p> <p>The following: <pre>create table foo external file 'foo.dat' (bar char(10));</pre> Results in: Segmentation fault</p> <p>An external file without an absolute path doesn't make sense, but we should not segfault.</p>	6
273	224810	<p>DISTINCT propagates outside a VIEW</p> <p>Given these two views:</p> <pre>CREATE VIEW VDISTIDX (RDB\$RELATION_NAME) AS select distinct rdb\$relation_name from rdb\$indices</pre> <pre>CREATE VIEW VDISTIDX2 (RDB\$RELATION_NAME) AS select rdb\$relation_name from rdb\$indices group by rdb\$relation_name</pre> <p>Then</p> <pre>select count(*) from vdistidx v join rdb\$relation_fields rf on v.rdb\$relation_name = rf.rdb\$relation_name with the scheduler saying PLAN SORT (JOIN (V RDB\$INDICES NATURAL,RF INDEX (RDB\$INDEX_4)))</pre> <p>produces a different result than</p> <pre>select count(*) from vdistidx2 v join rdb\$relation_fields rf on v.rdb\$relation_name = rf.rdb\$relation_name with the scheduler saying PLAN MERGE (SORT (RF NATURAL),SORT (SORT (V RDB\$INDICES NATURAL)))</pre> <p>Notice both JOIN statements produce different query plans but both views alone produce the same plan. I think that the first case is a bug: the DISTINCT clause is propagating outside the VIEW; just change the first case to become</p> <pre>select rdb\$relation_name from VDISTIDX v join rdb\$relation_fields rf on v.rdb\$relation_name = rf.rdb\$relation_name</pre> <p>and you'll observe that effectively, the DISTINCT applied to the result of the JOIN, leaving each table name only once.</p>	5 - Medium (next major release)
1332	406355	<p>Deletion from referenced tables in SP</p> <p>For example, in ISQL:</p> <pre>SET SQL DIALECT 3; CREATE DATABASE 'c:\test.gdb' USER 'SYSDBA' PASSWORD '*****' PAGE_SIZE 1024; CREATE TABLE Table1 (id INTEGER NOT NULL, foreign_id INTEGER);</pre>	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority
		<pre> CREATE TABLE Table2 (id INTEGER NOT NULL); INSERT INTO Table2 (id) VALUES (1); INSERT INTO Table2 (id) VALUES (2); INSERT INTO Table1 (id, foreign_id) VALUES (1, 1); COMMIT WORK; ALTER TABLE Table1 ADD CONSTRAINT PK_TABLE1 PRIMARY KEY (id); ALTER TABLE Table2 ADD CONSTRAINT PK_TABLE2 PRIMARY KEY (id); ALTER TABLE Table1 ADD CONSTRAINT FK_TABLE1 FOREIGN KEY (foreign_id) REFERENCES Table2 (id); CREATE PROCEDURE CLEAR_TABLES AS begin delete from Table2; suspend; end Executing procedure CLEAR_TABLES silent, and no records deleted as result execution.</pre>	
1363	428903	<p>Exception Handling Bug</p> <p>If you raise user defined exception (by EXCEPTION command) and it is handled by WHEN ANY DO or WHEN EXCEPTION ... DO then changes made by that SP are NOT undone !!</p> <p>Here is example:</p> <pre> CREATE TABLE tab (i INTEGER); CREATE EXCEPTION e 'abcd'; /* Division by zero in this SP will cause INSERT to be undone */ CREATE PROCEDURE proc1 AS DECLARE VARIABLE x INTEGER; BEGIN INSERT INTO tab VALUES (123); x=1/0; WHEN ANY DO EXIT; END /* Calling EXCEPTION and handling it by WHEN ANY in this procedure does NOT undo INSERTed row !! */ CREATE PROCEDURE proc2 AS BEGIN INSERT INTO tab VALUES (456); EXCEPTION e; WHEN ANY DO EXIT; END /* Here is called EXCEPTION, but it is not handled by WHEN ANY, so inserted row will be correctly removed. */ CREATE PROCEDURE proc3 AS BEGIN INSERT INTO tab VALUES (789); EXCEPTION e; END EXECUTE PROCEDURE proc1; EXECUTE PROCEDURE proc2; EXECUTE PROCEDURE proc3; SELECT * FROM tab; ===== 456 table should be empty !!!</pre>	3 - Low (to be reviewed)

Category: Core Engine

ID	SF ID	Description	Priority
288	228030	Float problems. I found some problems with float type decimal part. for example i write to field: FLOAT 0.35 phisicaly it write 0.349999994039536.	5 - Medium (next major release)
251	222376	Horrible plan with a lot of OR conditions Consider: select distinct doc.id docID, doc.directory directoryID from documents doc, DOC_APPLICATION docAppl where (docAppl.DOC_ID=doc.ID) and (docAppl.SITE_ID=2 or docAppl.SITE_ID=7 or docAppl.SITE_ID=25 or) the corresponding plan is PLAN SORT (JOIN (DOC NATURAL,DOCAPPL INDEX (RDB\$FOREIGN122,RDB\$FOREIGN123,RDB\$FOREIGN123,RDB\$FOREIGN123,RDB\$FOREIGN123...)) And an RDB\$FOREIGN123 is added for each additionnal or docAppl.SITE_ID= in the query. In one particular circumstance the query gets particularly long and the corresponding plan as well... In that case I'm getting the unknown error in ISQL (the records are fetched and presented in the dbgrid, statistics are here, but plan is empty). And yes, IB does a very poor job at generating the plan in this particular case, the optimal plan is just PLAN SORT (JOIN (DOC NATURAL,DOCAPPL INDEX (RDB\$FOREIGN122,RDB\$FOREIGN123)) which speeds the request quite a lot...	3 - Low (to be reviewed)
264	223514	IB crashes with two procedures intermixed. These two procedures expect an order and an order-line number. The following will make IB crash: // Show al possible parties on stock, with the same characteristics select * from GET_ORDVRD(258,1) // Show the status of the stock (matched, alternatives etc) select * from INKORDVRDSTATUS(258,1) // First one select * from GET_ORDVRD(258,1)	5 - Medium (next major release)
239	217042	IB doesn't validate weird constructions 1) create table t0(check(0=0)) How can a table constraint be made without any field? I thought at least one field should be defined. 2) create table tgen2(a computed by(0)) create table tgen3(a computed by(0/(0))) create table tgen4(a computed by(rdb\$db_key)) Do you see any value in accepting this declaration? A computed field alone? I can never have a row. 3) create table tgen5(a int, b computed by(rdb\$db_key)) Insert values in field a and then try to select = string truncation or arithmetic overflow. 4) create domain dint int default 'hello' not null This is only a small example. IB never validates the default for any type of field. I can insert any kind of crap as default fc domains of any datatype and also for fields defined inside a table.	5 - Medium (next major release)
241	217138	JOIN including a complex view kills the server The SQL is as follows: Select distinct fm.code, fm.Description, fm.link from fullmenu fm join menu_groups mg on fm.code = mg.menu_id When I try and prepare this I get: "Unable to complete network request to host 10.0.0.200. Error reading data from the connection and the connection is lost." <For full details see original bug report>	5 - Medium (next major release)
1345	419065	Join on diffrent datatypes I have the same field, in one table as char(10) (named FZDNR) and in the second table as NUMERIC (HNR). I need the	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority
247	221925	<p>following select expression:</p> <pre>SELECT * FROM WA WHERE HNR NOT IN (SELECT FZDNR FROM H)</pre> <p>In the pre-6 Versions of Interbase it was working fine, but now I get the whole table WA, so I think the join on HNR-FZDNR doesnt work. Tell me pls, is it a bug or is it an feature?</p> <p>Left joining table to sproc: ORDER BY makes fields NULL</p> <pre>CREATE DOMAIN JOBCODETYPE AS VARCHAR(20); CREATE DOMAIN STOCKCODETYPE AS VARCHAR(30); create table Submit (jobid jobcodetype not null primary key, completeby timestamp); create table Faults (id stockcodetype not null, whenreported timestamp default 'Now' not null, whendone timestamp, jmsref jobcodetype not null, primary key (id,whenreported), foreign key (jmsref) references Submit(jobid) on update cascade); create view xAllFaults(jobid) as select distinct jmsref from Faults; create view xWaitingFaults(jobid) as select distinct jmsref from Faults where whendone is null; create view AllFaults(jobid,waiting) as select XAllFaults.jobid,XWaitingFaults.jobid from XAllFaults left join XWaitingFaults on XAllFaults.jobid = XWaitingFaults.jobid; /* All jobid's with faults, and 'f' or 'F' according to if any are undone */ set term !! ; create procedure JobFaults returns (jobid varchar(20), fault varchar(1)) as declare variable waiting varchar(20); begin for select XAllFaults.jobid,XWaitingFaults.jobid from XAllFaults left join XWaitingFaults on XWaitingFaults.jobid = XAllFaults.jobid into :jobid,:waiting do begin if (waiting is null) then fault = 'f'; else fault = 'F'; end suspend; end end!! set term ; !! insert into Submit values ('A',cast('Now' as timestamp)); insert into Submit values ('B',cast('Now' as timestamp)); insert into Faults(id,whenreported,jmsref) values ('item A',cast('Now' as timestamp),'A'); insert into Faults(id,whenreported,whendone,jmsref) values ('item B',cast('Now' as timestamp),cast('Now' as timestamp),'A'); /* NOW TRY THE FOLLOWING */ select * from Submit join JobFaults</pre>	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority																																																																																
1321	233602	<p>on Submit.jobid = JobFaults.jobid order by Submit.completeby; /* Here there appear to be no faults joined to the Submit table */</p> <p>select * from Submit join JobFaults on Submit.jobid = JobFaults.jobid /* Yet here the fault is shown correctly joined!! */</p> <p>Like operator doesn't work correctly with UPPER()</p> <p>select * from gd_contact where upper(name) like upper('%some_lowercase_belarusian_string%')</p> <p>this query returns empty cursor although there are records in the database which match the criteria.</p> <p>if change SQL to:</p> <p>select * from gd_contact where upper(name) like upper('%THE_SAME_STRING_IN_UPPERCASE%')</p> <p>all works fine, but I need to convert string to uppercase pragramattically, and of course, there is no need for second UPPER().</p> <p>Additional note: field NAME defined as VARCHAR(60) CHARACTER SET win1251 COLLATION PXW_CYRL.</p>	5 - Medium (next major release)																																																																																
255	223058	<p>Multi-hop server ability broken</p> <p>The feature that allows you to connect to IB server through several other IB servers, like alpha:beta:\gamma\delta:C:\InterBase\db.gdb does not work, if one of the relay servers (i.e. alpha, beta, gamma from above example) is IB6.</p> <p>IB6 will allow connection, but when trying to execute any command you will get Dynamic SQL Error -SQL error code = -901 -feature is not supported</p>	3 - Low (to be reviewed)																																																																																
289	228135	<p>NULL is returned as zero through a left join in simple VIEW</p> <p>Consider:</p> <pre>select pvb.project, pvb.mjr, pvb.mnr, pvb.rls, pvb.blid, b.id, b.priority, b.entered_by, b.closed, b.status from bug b left outer join project_version_bug pvb on (pvb.bug = b.id) where b.assigned_to = user;</pre> <table border="1"> <thead> <tr> <th></th> <th>PROJECT</th> <th>MJR</th> <th>MNR</th> <th>RLS</th> <th>BLD</th> <th>ID</th> <th>PRIORITY</th> </tr> <tr> <th>ENTERED_BY</th> <th colspan="3">CLOSED</th> <th colspan="2">STATUS</th> <th colspan="2"></th> </tr> </thead> <tbody> <tr> <td>N A</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> <td>SYSDBA</td> </tr> <tr> <td>N A</td> <td>1</td> <td>0</td> <td>1</td> <td>0</td> <td>1</td> <td>1</td> <td>SYSDBA</td> </tr> <tr> <td>N A</td> <td>null</td> <td>null</td> <td>null</td> <td>null</td> <td>2</td> <td>1</td> <td>SYSDBA</td> </tr> </tbody> </table> <p>This is correct, you can see the null values where there are no records in project_version_bug for bug id 2.</p> <p>However, using the BUG_ASSIGNED view (which is defined with exactly the same SQL statement) I get the following:</p> <pre>select * from bug_assigned;</pre> <table border="1"> <thead> <tr> <th></th> <th>PROJECT</th> <th>MJR</th> <th>MNR</th> <th>RLS</th> <th>BLD</th> <th>ID</th> <th>PRIORITY</th> </tr> <tr> <th>ENTERED_BY</th> <th colspan="3">CLOSED</th> <th colspan="2">STATUS</th> <th colspan="2"></th> </tr> </thead> <tbody> <tr> <td>N A</td> <td>1</td> <td>0</td> <td>0</td> <td>0</td> <td>1</td> <td>1</td> <td>SYSDBA</td> </tr> <tr> <td>N A</td> <td>1</td> <td>0</td> <td>1</td> <td>0</td> <td>1</td> <td>1</td> <td>SYSDBA</td> </tr> <tr> <td>N A</td> <td>0</td> <td>0</td> <td>0</td> <td>0</td> <td>2</td> <td>1</td> <td>SYSDBA</td> </tr> </tbody> </table>		PROJECT	MJR	MNR	RLS	BLD	ID	PRIORITY	ENTERED_BY	CLOSED			STATUS				N A	1	0	0	0	1	1	SYSDBA	N A	1	0	1	0	1	1	SYSDBA	N A	null	null	null	null	2	1	SYSDBA		PROJECT	MJR	MNR	RLS	BLD	ID	PRIORITY	ENTERED_BY	CLOSED			STATUS				N A	1	0	0	0	1	1	SYSDBA	N A	1	0	1	0	1	1	SYSDBA	N A	0	0	0	0	2	1	SYSDBA	5 - Medium (next major release)
	PROJECT	MJR	MNR	RLS	BLD	ID	PRIORITY																																																																												
ENTERED_BY	CLOSED			STATUS																																																																															
N A	1	0	0	0	1	1	SYSDBA																																																																												
N A	1	0	1	0	1	1	SYSDBA																																																																												
N A	null	null	null	null	2	1	SYSDBA																																																																												
	PROJECT	MJR	MNR	RLS	BLD	ID	PRIORITY																																																																												
ENTERED_BY	CLOSED			STATUS																																																																															
N A	1	0	0	0	1	1	SYSDBA																																																																												
N A	1	0	1	0	1	1	SYSDBA																																																																												
N A	0	0	0	0	2	1	SYSDBA																																																																												

Category: Core Engine

ID	SF ID	Description	Priority
		<p>N A</p> <p>or even this:</p> <pre>select * from bug_assigned where project is null;</pre> <pre> PROJECT MJR MNR RLS BLD ID PRIORITY ENTERED_BY CLOSED STATUS ===== 0 0 0 0 0 2 1 SYSDBA N A </pre> <p>Here all the null values as a result of the left outer join are presented as zeros.</p> <p>This is a case using one column to demonstrate even clearer the bug:</p> <pre>select project from bug_assigned where project is null;</pre> <p>Produces a zero as the result. The engine seems to recognize that the value is NULL so it honors the WHERE clause, but it returns zero. The view is very simple to be considered a boundary case.</p>	
1473	447535	<p>NUMERIC(18,0) as unique index problem</p> <p>If you try to execute this script on a firebird server on linux it will report you that there are duplicate values in the index.this also affects primary keys. if you put the CREATE UNIQUE INDEX statement before the INSERT INTO statements everything will work if you don't deactivate the index or backuprestore the database. it's extremely problematic if you try to backuprestore a database on linux, since you can't restore it if you have used a numeric(18,0) as primary key without deactivating all indexes. and they cannot be reactivated! It seems as the server is using DOUBLE values in some places resulting in rounding of the exact values. the same script, however, works without problems with the win32 version of firebird.</p> <pre>CREATE TABLE BACK_ME_UP (ID NUMERIC(18, 0) NOT NULL);</pre> <pre>INSERT INTO BACK_ME_UP (ID) VALUES (123456789012345678); INSERT INTO BACK_ME_UP (ID) VALUES (123456789012345679);</pre> <pre>CREATE UNIQUE INDEX IDX_BACK_ME_UP_ID ON BACK_ME_UP(ID);</pre>	5 - Medium (next major release)
243	219525	<p>No current record for fetch operation</p> <pre>SELECT * FROM view_det INNER JOIN view_inv ON view_det.invcod=view_inv.invcod</pre> <p>DOESN'T WORK - No current record for fetch operation.</p> <p>... and neither will:</p> <pre>SELECT * FROM view_det, view_inv WHERE view_det.invcod=view_inv.invcod</pre> <p>If you watch the query plan generated by the failing statement, you'll notice it's really big... no surprise since it involves a join of views that in turn are based on other joins. The workaround is to convert</p> <pre>SELECT * FROM view_det JOIN view_inv ON view_det.invcod=view_inv.invcod</pre> <p>to become</p> <pre>SELECT * FROM view_det JOIN view_inv ON view_det.invcod-view_inv.invcod=0</pre> <p>or</p> <pre>SELECT * FROM view_det JOIN view_inv ON view_det.invcod+0=view_inv.invcod</pre>	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority
		<p>or other variations that fool the optimizer by forcing it to a full scan to do the join. The original statement should work, so this is a bug and the same kind of bug that IB5.X suffered from.</p> <p><For full details see original bug report></p>	
246	221921	<p>ORDER BY has no effect</p> <p>Take the following example of a self-referential table and a sproc that returns the children of a specified item:</p> <pre> create table ExampleTable (code integer not null primary key, name varchar(100) not null unique, parent integer, foreign key (parent) references ExampleTable(code)); /* Children result is not null if this item has it's own children */ set term !! ; create procedure ChildrenOfItem(par integer) returns (code integer,children integer) as begin for select MainTypes.code, Min(ChildTypes.code) from ExampleTable MainTypes left join ExampleTable ChildTypes on MainTypes.code = ChildTypes.parent where MainTypes.parent = :par or (MainTypes.parent is null and :par is null) group by MainTypes.code into :code,:children do suspend; end!! set term ; !! insert into ExampleTable values (0,'A',null); insert into ExampleTable values (1,'AA',0); insert into ExampleTable values (3,'AB',0); insert into ExampleTable values (4,'AC',0); insert into ExampleTable values (2,'AD',0); insert into ExampleTable values (5,'B',null); insert into ExampleTable values (6,'BA',5); insert into ExampleTable values (7,'BB',5); insert into ExampleTable values (8,'BC',5); insert into ExampleTable values (9,'BD',5); insert into ExampleTable values (10,'BE',5); insert into ExampleTable values (11,'BF',5); select * from ChildrenOfItem(0); /* Gives 1,2,3,4 as you would expect */ select * from ChildrenOfItem(0) inner join ExampleTable on ChildrenOfItem.code = ExampleTable.code order by name; /* gives 'AA','AD','AB','AC' even though it is order on name!! Codes are still 1,2,3,4 HOWEVER, in this example, changing it to ORDER BY NAME DESC correctly returns AD,AC,AB,AA In my real system however, neither ORDER BY NAME or ORDER BY NAME DESC has any effect */ </pre>	5 - Medium (next major release)
276	225283	<p>ORDER BY on a VIEW turns values in fields into NULL</p> <p>Consider a database with 5 tables and a couple of views, one of those views is called v_observation.</p> <p>Consider two queries</p> <pre> select * from v_observation select * from v_observation order by iobservationagree desc </pre> <p>These two queries return different (!) results. And not just the order.... It seems Firebird has trouble dealing with NULL values that are returned and, in the second case, completely omit the information retrieved from the database.</p>	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority
1366	430509	<p>ORDER BY works wrong with collate PT_PT</p> <p>Consider,</p> <pre>CREATE TABLE STOCKS (MNEN INTEGER NOT NULL, ACTIVO VARCHAR(50) CHARACTER SET ISO8859_1 COLLATE PT_PT, CONSTRAINT PK_STOCKS PRIMARY KEY (MNEN));</pre> <p>having the values</p> <ol style="list-style-type: none"> 1, BA 2, BES 3, BCP 4, BA Pref. 5, Banif <p>doing</p> <pre>SELECT ACTIVO FROM STOCKS ORDER BY ACTIVO</pre> <p>gives:</p> <p>BA Banif BA Pref. BCP BES</p>	5 - Medium (next major release)
1456	447002	<p>Optimisation on subselects</p> <p>The InterBase Optimizer does not properly utilize indexes on sub-selects containing IN.</p> <p>Example: (Using - Employee.gdb)</p> <pre>select * from employee where dept_no in (select dept_no from department where department = Marketing);</pre> <p>Produces the plan:</p> <pre>PLAN (DEPARTMENT INDEX (RDB\$4,RDB\$PRIMARY5)) PLAN (EMPLOYEE NATURAL)</pre> <p>You can fool the optimizer into using the appropriate index by using the following trick.</p> <pre>select * from employee where dept_no 0 /* something that is always true */ and dept_no in (select dept_no from department where department = Marketing);</pre> <p>Produces the plan:</p> <pre>PLAN (DEPARTMENT INDEX (RDB\$4,RDB\$PRIMARY5)) PLAN (EMPLOYEE INDEX (RDB\$FOREIGN8))</pre> <p>The second plan is correct behaviour</p>	3 - Low (to be reviewed)
229	213460	<p>Registering Events w/certain configuration crashes IB Server</p> <p>Overview</p> <p>With a certain server network configuration, registering an event will lock up the client and lockup the server (requiring a reboot if it runs as a service, or a process kill if it runs as an application (the guardian does not detect it's lockup)) *if* the client is accessing the server through the internet (if it comes through the local lan, then it's fine).</p> <p>Configuration Description To Cause Bug</p>	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority
		<p>The above system has only one nic card in it. This card was assigned two IP's, the first being a 192.168.0.x (local network), and the second being it's internet address. A gateway for the NIC card was also set, in addition to IP Forwarding (which was removed later).</p> <p>This configuration, while not considered the best, is perfectly legal, and works effectively with all the different servers run on it (MS SQL Server, ftp and web servers, etc.), both accessed from the internet and the local intranet.</p> <p>It is only when a client attempts to register events when connected through the internet to the server system that the lockups occur. Otherwise, if the system is on the intranet, the access runs fine.</p> <p>Fix</p> <p>By setting the internet IP first, then having the intranet ip set in the advanced screen (under the internet IP), the problem is resolved (a reboot will be required after changing the configuration).</p> <p>Comments</p> <p>This was a strange bug to hunt down, and stems from the fact that the register events connection to the server through the internet is somehow 'different' enough to cause the crashes (the client will come back to life if the server is shutdown (in NT the server process has to be manually terminated).</p> <p>While the problem requires a certain network configuration to occur (which tends to be common for small internal networks with a dsl or cablemodem connection to the internet), but given the anticipated spreading of InterBase with new applications, it becomes possible that some of our customers may receive this frustrating error.</p> <p>This is made even more dangerous by the fact that their network configurations may have been set by a third party consultant, or department IT person.</p> <p>At the minimum, I would recommend that the IB server at have an exception written into it to handle this type of error gracefully, and return an exception to the connecting client indicating the network incompatibility.</p>	
1317	233025	<p>Server hangs when executing Stored Proc more than once</p> <p>When executing a SP inside a loop that accepts input parameters, the server hangs on the second iteration if the statement is not unprepared and re-prepared between iterations. This is an old bug which goes back to v. 5.x as well.</p>	9 - High Priority/Immediate (next
293	228716	<p>Tbl left join sproc + order by = no records returned</p> <p>Create the following procedure with the employee.gdb that comes with IB6, still in dialect 1:</p> <pre> set term ^; create procedure del_old_salary_history as declare variable dbk char(8); declare variable dbk_prev char(8); declare variable emp_no smallint; declare variable emp_no_prev smallint; begin emp_no_prev = NULL; FOR select sh.rdb\$db_key, emp_no from salary_history sh order by emp_no, change_date INTO :dbk, :emp_no DO BEGIN IF (emp_no = emp_no_prev) then delete from salary_history where rdb\$db_key = :dbk_prev; emp_no_prev = emp_no; dbk_prev = dbk; END end ^ set term ;^ </pre>	5 - Medium (next major release)
		<p>But for testing that without deleting records, I decided to show which one I would delete, so I converted the procedure; only the records that are shown are deleted in the real case:</p> <pre> set term ^; create procedure show_del_old_salary_history returns (out smallint, out2 timestamp) </pre>	

Category: Core Engine

ID	SF ID	Description	Priority
		<pre> as declare variable dbk char(8); declare variable dbk_prev char(8); declare variable emp_no smallint; declare variable emp_no_prev smallint; declare variable change_date timestamp; declare variable change_date_prev timestamp; begin emp_no_prev = NULL; FOR select sh.rdb\$db_key, emp_no, change_date from salary_history sh order by emp_no, change_date INTO :dbk, :emp_no, :change_date DO BEGIN out = emp_no_prev; out2 = change_date_prev; IF (emp_no = emp_no_prev) then /* delete from salary_history where rdb\$db_key = :dbk_prev */ suspend; dbk_prev = dbk; emp_no_prev = emp_no; change_date_prev = change_date; END end ^ set term ;^ </pre>	
		<p>Then, since I didn't want to do a tedious comparison by hand between the affected table and the procedure, this sentence should compare records in the table with the ones that the real procedure would delete: the ones with NULL in the procedure side aren't deleted:</p>	
		<pre> select s.emp_no, s.change_date, p.out, p.out2 from salary_history s left join show_del_old_salary_history p on s.emp_no = p.out and s.change_date = p.out2 </pre>	
		<p>It works, but since the mandatory table is not ordered, it makes very difficult to tell if the shown records are the correct. Hence, I added a sorting clause to the above LEFT JOIN and...</p>	
		<pre> *** - order by 1 = plan but no data (no records) - order by 1,2 = plan but no data (no records) - order by 2 = plan but «Conversion error from string 1» - order by 3,4 = plan but «Conversion error from string 1» - group by s.emp_no, s.change_date, p.out, p.out2 = plan but no data (no records) - group by p.out, p.out2, s.emp_no, s.change_date = plan but «Conversion error from string 1» </pre>	
		<pre> *** Changing LEFT JOIN to JOIN displays data even with an ORDER BY, but that's not useful for solving the need I had to check my original procedure's actions. </pre>	
		<pre> *** Trying to narrow down the conversion error problem, I used only a sort on the procedure: select p.out, p.out2 from show_del_old_salary_history p order by 1,2 = no problem </pre>	
		<pre> select s.emp_no from salary_history s left join show_del_old_salary_history p on s.emp_no = p.out order by 1 = no records </pre>	
		<pre> select s.emp_no from salary_history s left join show_del_old_salary_history p on s.emp_no = p.out order by p.out = «Conversion error from string 1» and since I'm not showing the two timestamp fields, this is a problem with the sort routine. </pre>	
		<pre> *** Workaround: we need to order the salary_history in advance, using a procedure, since VIEWS doesn't allow an ORDER BY </pre>	

Category: Core Engine

ID	SF ID	Description	Priority
		<p>clause inside their definition:</p> <pre> set term ^; create procedure save_me returns(emp_no smallint, change_date timestamp) as begin for select s.emp_no, s.change_date from salary_history s order by 1,2 into :emp_no, :change_date do suspend; end ^ set term ;^ </pre> <p>Now, we do our rogue left join, but using the two procedures:</p> <pre> select s.emp_no, s.change_date, p.out, p.out2 from save_me s left join show_del_old_salary_history p on s.emp_no = p.out and s.change_date = p.out2 </pre> <p>This proved that my routine is okay: the original procedure deletes exactly what I want, but I stumbled across a sorting bug in the road. The last sentence returns 49 records but the flawed ones return zero records.</p>	
1335	409769	<p>UDF argument can't be query-parameter</p> <pre> SELECT Field1, Field2, ... , UDF1(Field1, ?Param1) FROM TABLE1 </pre>	5 - Medium (next major release)
245	221649	<p>don't work - message is unknown data type</p> <p>Unique index allowed on NULLABLE field</p> <pre> create table mau(a int); commit; create UNIQUE index idx_mau on mau(a); commit; </pre> <p>Whether this was left as a transition from Paradox (that accepts NULL PKs) or it's an oversight is what I want to confirm.</p> <p>As a comparison, you cannot declare a PK or UNIQUE constraint if you don't include the NOT NULL clause on the affected field(s), so the underlying automatic index for such constraint will be always on non-nullable column(s).</p>	5 - Medium (next major release)
230	213462	<p>Win 32 Open database file exclusive accept/accepts ambiguous connect stri</p> <p>Win 32 allows clients to connect with either servername:c:\localpathname or servername:c:\localpathname.</p> <p>If connected clients have connected with a mixture of these string formats, the database is treated as if it were multiple databases and corruptions occur.</p> <p>Opening the database file using exclusive access will ensure that access is controlled.</p>	8
214	211781	<p>Win32: Server don't close thread handles</p> <p>Allocation location is in SVC_start, last param 'service-svc_handle'.</p>	5 - Medium (next major release)
1322	233644	<p>cannot specify PLAN in UPDATE statement</p> <p>I need to set FULLNAME=NAME for all records in company where fullname IS NULL or fullname = "", so I wrote the following statement:</p> <pre> UPDATE company cp SET fullname= (SELECT name FROM contact ct WHERE ct.id=cp.contactkey) WHERE cp.fullname IS NULL OR cp.fullname = ""; </pre> <p>Both tables contains ~40K records so the request hangs my machine. After reboot I checked the plan and it was:</p> <pre> PLAN(company NATURAL, contact NATURAL) </pre> <p>I tried to set PLAN for entire UPDATE statement or only for SELECT sub-query but in both cases I received error message Invalid SQL or something like this.</p> <p>In the end of all, I just wrote STORED PROC what did all the work in 3 seconds.</p>	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority
217	212177	<p>error with non-english column defaults</p> <pre>create table bug_bugbase (s varchar(20) character set win1251 collate win1251 default 'XXXX' not null);</pre> <p>Note: win1251 code page, is for belarusian language and the default value above 'XXXX' is a cyrillic (belarusian) text.</p> <p>Database works well, it can insert records, update, and delete them. After backup, the restore raises:</p> <pre>gbak: restore failed for record in table BUG_BUGBASE gbak: ERROR: arithmetic exception, numeric overflow, or string truncation gbak: ERROR: Cannot transliterate character between character sets</pre> <p>it seems that IB supports default values only in standard (english) code page.</p>	9 - High Priority/Immediate (next
234	215169	<p>stored procedure gets off sync on the server</p> <p>Sometimes a prepared stored procedure appears to get off sync on the serverThis code is meant to try to work around the problem simply by retrying. It is in IBSQL.PAS, in the TIBSQL.ExecQuery:</p> <pre>... SQLExecProcedure: begin fetch_res := Call(isc_dsqli_execute2(StatusVector, TRHandle, @FHandle, Database.SQLDialect, FSQParams.AsXSQlda, FSQRecord.AsXSQlda), False); if (fetch_res = 0) and (fetch_res = isc_deadlock) then begin { Sometimes a prepared stored procedure appears to get off sync on the serverThis code is meant to try to work around the problem simply by retrying. This need to be reproduced and fixed. } isc_dsqli_prepare(StatusVector, TRHandle, @FHandle, 0, PChar(FProcessedSQL.Text), 1, nil); Call(isc_dsqli_execute2(StatusVector, TRHandle, @FHandle, Database.SQLDialect, FSQParams.AsXSQlda, FSQRecord.AsXSQlda), True); end; end ... </pre>	5 - Medium (next major release)
1339	414420	<p>wrong order by in table join storedproc</p> <p>InterBase will not sort correctly the result set of a table joined with a stored procedure (inner join), if an ascending sort on an indexed field is requested.</p> <p>To expose the bug we need to join a table with a stored procedure; here's the test metadata (tested with dialect 1):</p> <pre>CREATE TABLE MY_TABLE (K VARCHAR(10) NOT NULL, D1 INTEGER, D2 INTEGER, V1 VARCHAR(10), PRIMARY KEY (K)); CREATE PROCEDURE SELECT_ME RETURNS (DATA VARCHAR(10)) AS begin data = 'one'; suspend; data = 'two'; suspend; data = 'three'; suspend; end </pre>	5 - Medium (next major release)

Category: Core Engine

ID	SF ID	Description	Priority
----	-------	-------------	----------

Let's populate my_table with some test data:

```
insert into my_table values ('one', 1, 99, 'zz');
insert into my_table values ('two', 2, 98, 'yy');
insert into my_table values ('three', 3, 97, 'xx');
```

The following query exposes the bug:

```
select *
from my_table t join select_me p
on (t.k = p.data)
order by t.d1
```

The ordering is correct on every field we put in the ORDER BY clause, UNLESS that field is also indexed. Here are the correct results of the query:

K	D1	D2 V1	DATA
one	1	99 zz	one
two	2	98 yy	two
three	3	97 xx	three

If we now do:

```
CREATE INDEX I1 ON MY_TABLE(D1);
```

the same query returns the following results:

K	D1	D2 V1	DATA
one	1	99 zz	one
three	3	97 xx	three
two	2	98 yy	two

In short, if an inner join between a table and a stored procedure is performed, ordering will work only on non-indexed fields. Please note that:

- if we do a left join instead the problem disappears.
- if we remove the index the problem also disappears.
- if we order by t.k (the table's primary key) the results are always correct.
- if we request a descending order the problem also disappears.

Category: DSQL

ID	SF ID	Description	Priority
----	-------	-------------	----------

1385 444763 **Empty column names with aggregate funcs** 5 - Medium (next major release)

If you execute this statement the column name of the inner select is empty:

```
select (select count(1) from rdb$database) from rdb$database
```

I don't know if this is a bug or just a limitation but it seems that this bug(limitation) is somehow related to #222476. The difference is that it happen with all aggregate functions(count, min, max, sum, avg) and also in dialect 1.

1383 444081 **FOR SELECT firs SUB SELECT is wrong** 5 - Medium (next major release)

```
FOR SELECT
DO BEGIN
SELECT xxx FROM WHERE yyy in (SELECT xxx)...; -
It's wrong
END
```

265 223516 **Missing types in rdb\$types.** 5 - Medium (next major release)

LangRef describes among values that rdb\$fields.rdb\$field_type can assume the following two:
D_FLOAT - 11
INT64 - 16

Both values appear in the documentation but not in rdb\$types. I don't have any idea what D_FLOAT is but at least INT64 should be listed in rdb\$types for consistency reasons. Also, putting NUMERIC (1) and DECIMAL (2) as registered sub_types would help programs that cope with metadata.

Category: DSQL

ID	SF ID	Description	Priority
1382	442140	<p>ROLES GRANT FOR VIEWS NOT WORKING</p> <p>Please check why I'm getting error that i don't have select rights on a table when i select from a view who has select rights on the table.</p> <p>To Check te problem :</p> <ol style="list-style-type: none"> 1. Restore the DB 2. Create a new user 3. Make the user member off CALS_USER 4. Login with that user and the role CALS_USER 5 Execute the following statement. <p>SELECT * FROM VW_OBJECTS.</p> <p>I'm getting the following error : No Select/Read rights on tablecatalog\$lookup_info.</p>	5 - Medium (next major release)
1476	449011	<p>Unique key and char type in SQLDialect 3</p> <p>Unique key and char type in SQL Dialect 3 If to create a database with use SQL Dialects 3, to the line data there is a following:</p> <p>Example of the table: CREATE TABLE LABELS (CODE VARCHAR (13) NOT NULL, VALUT SMALLINT NOT NULL, PRIMARY KEY (CODE));</p> <p>Both columns should not be equaled NULL and primary (unique!) key - CODE.</p> <p>At addition in a database of a line (for example CODE = ABC and VALUT=1) all is normal,</p> <p>But if to remember changes, then it is possible to appropriate(give) to a field CODE=NULL!!! - and it NOT a ZERO INDEX, and primary!</p>	5 - Medium (next major release)
1379	440541	<p>insert into... select</p> <p>one table. create table t1 (f1 varchar(5), f2 integer)</p> <p>insert into t1 (f1, f2) values ('ss',1); insert into t1 (f1, f2) values ('qq',2);</p> <p>bug ===== insert 27(!!!) records instead 1 insert into t1 select * from t1 where f2=2; same insert into t1 select t2.* from t1 t2 where t2.f2=2;</p> <p>trap server if execute insert into t1 (f1,f2) select f1,f2 from t1 where t2.f2=2;</p>	5 - Medium (next major release)

Category: GBAK

ID	SF ID	Description	Priority
1344	417567	<p>GBAK can't restore database after backup</p> <p>Sometimes GBAK can't restore database from successfull backup. It can't resolve dependencies between triggered views.</p>	5 - Medium (next major release)
1355	424307	<p>GBAK overwrites BK file</p> <p>GBAK overwrites BK file before starting back up process. So if you have database with valuable data and backup copy of it. And your database file gets corrupted so GBAK cannot proceed and you will execute command gbak -b database backup</p> <p>gbak will erase backup file then checks database file and after encountering error will stop process so you will loose your backup file forever.</p> <p>In my example I had GDB file and BK. GDB file was in 5.6 format and server was of 6.0. I forgot for version of GDB file and run gbak. It deleted my BK file and raises error message that GDB of wrong ODS version.</p>	5 - Medium (next major release)

Category: GBAK

ID	SF ID	Description	Priority
290	228431	<p>The right behavior would be to overwrite BK file only after successful process is done.</p> <p>gbak cannot restore backup made by IB v5</p> <p>gbak under Linux is coredumping while loading a database backup made by gbak v5. the bug is deterministic and does not concern the official borland gbak for v6 and is related to the database content - only one of 3 different database backups generated the bug.</p> <p>Details: System - Redhat 7 Firebird - 0.9-4 CS</p> <p>Database in question: about 150MB, no triggers, no stored procedures - pure data but about a dozen of generators.</p> <p>Returning to official gbak the bug disappeared.</p>	5 - Medium (next major release)

Category: GPRE

ID	SF ID	Description	Priority
1341	416228	<p>gpre generates invalid isc_vtov calls</p> <p>gpre generates the following incorrect code for tan.e in TCS (line 6303+):</p> <pre> isc_vtov (isc_342.isc_345, run_name, sizeof (run_name)); run_count = isc_342.isc_343; isc_vtov (isc_342.isc_346, max_date, sizeof (max_date)); isc_vtov (isc_342.isc_347, min_date, sizeof (min_date)); </pre> <p>The second parameter to isc_vtov is the incorrect type. This causes the compiler in Darwin to generate an error message and not compile the tan.c file.</p>	7 - Medium-High (next minor rele

Category: ISQL

ID	SF ID	Description	Priority
263	223513	<p>Ambiguity between tables and views in isql's SHOW commands.</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> /* Notice both T and V are shown, but only T should be shown. One can argue that a view is a virtual table, etc. */ SQL show tables; T V SQL show views; V /* This shouldn't work */ SQL show table v; A INTEGER Nullable View Source: ===== select a from t /* This shouldn't work */ SQL show view t; A INTEGER Nullable </pre>	5 - Medium (next major release)
1353	422240	<p>Left join with stored procedure</p> <p>We have two stored procedures and want to do select *</p>	5 - Medium (next major release)

Category: ISQL

ID	SF ID	Description	Priority
		from proc1('AAA') p1 left join proc2('AAA') p2 on (p1.id=p2.id) if execute this query - server crashes.	
258	223126	Misplaced collation when extracting metadata with isql There is a bug that if you override a domain's constraints or default and also override the collate the collate will be extracted before the default or constraint.	5 - Medium (next major release)
275	225219	isql -a: wrong order with domains based on table's fields According to the documentation: «The CHECK constraint in a domain definition sets a dom_search_condition that must be true for data entered into columns based on the domain. The CHECK constraint cannot reference any domain or column.» So, this should be banned: create domain d_ontable int check(value(select rdb\$relation_id from rdb\$database)) Now what happens with a restore if the domain's check constraint is based on a user table? (And not a system table, where it could work eventually.) Markus Kemper discovered this issue initially but he wrote that isql was unable to extract info from it, because he thought it was an isql limitation. This is only geeky detail or joke, a bit unrelated with the previous: please see how useful is this domain: create domain value_gt_value int check(valuevalue)	5 - Medium (next major release)
1481	450667	isql needs update to cope with FB enhancements We forgot ISQL with some modifications to the core engine. For example, I previously coded a workaround to inform that BLOBs are returned by descriptor in the SHOW TABLE command, since they were reported as returned by VALUE, that's impossible. This by descriptor info is strictly speaking, an error. BLOBs are returned normally by blob struct, unless they really were requested by descriptor. Second, ISQL knows nothing in the SHOW TABLE and in the [full] metadata script command (-x -a) about descriptors in DSQL, so once a UDF declaration using the new syntax parameter BY DESCRIPTOR has been issued, isql won't be able to rebuild properly the script from the live database.	5 - Medium (next major release)
1479	450404	isql uppercases role in the command line I think example should be enough clear: when isql is invoked and the role is passed in the command line, it always uppercases the role: H:\ibdev\fbbuild\interbase\dsqllsql c:/proy/fbtest.gdb -u cvc -p pw -r for cvc Database: c:/proy/fbtest.gdb, User: cvc, Role: FOR CVC SQL As you can see, isql puts the message with the role in uppercase and indeed the engine received it uppercased. Notice the double quotes weren't interpreted as don't change the input parameter. However, when doing the connection as a command, the role is processed correctly: 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 Now, isql reflects the unchanged role and the engine receives it correctly. Whether this anomaly exists for some purpose I don't know, but I consider it a bug. Both ways to work with isql should be equivalent when facing identifiers surrounded by double quotes.	5 - Medium (next major release)

Category: Installation - Server

ID	SF ID	Description	Priority
----	-------	-------------	----------

Category: Installation - Server

ID	SF ID	Description	Priority
261	223327	<p>FireBird 0.9 (fbinst_prod_110300-2.exe) reg. error</p> <p>Installed fbinst_prod_110300-2.exe on a clean machine (WIN2000) in D:\Program Files\FireBird.</p> <p>Tried to connect via IBConsole, got "couldn't attach to password db".</p> <p>Checked with regedit and found:</p> <p>HKEY_LOCAL_MACHINE\SOFTWARE\Borland\InterBase\CurrentVersion\RootDirectory = :\Program Files\FireBird\ and HKEY_LOCAL_MACHINE\SOFTWARE\Borland\InterBase\CurrentVersion\ServerDirectory = c:\Program Files\FireBird\ corrected both to the installed path (d:\Program Files\FireBird\), restarted the InterBase Server service and everything was ok.</p>	5 - Medium (next major release)
270	224130	<p>Install conflict with Interbase</p> <p>I had previously installed the ib_client package. I installed firebird. Things got very confused. For example here are the registry entries:</p> <pre>[HKEY_LOCAL_MACHINE\Software\Borland\InterBase\CurrentVersion] UseCount=dword:00000001 RootDirectory=C:\Program Files\Borland\InterBase\ Version=W1-V6.0.0 DefaultMode=-r ServerDirectory=:\Program Files\firebird\ GuardianOptions=1</pre> <p>Notice that it was installed into C:\Program Files\Firebird but the root directory is C:\Program Files\Borland\Interbase. The RootDirectory entry was apparently left over from my install of ib_client. After I thoroughly uninstall Firebird and Interbase and re-install Firebird, the RootDirectory is C:\Program Files\firebird\ which is probably correct.</p> <p>The normal procedures of installing software under Windows is that either [a] Firebird is Interbase so I should be able to install over an old version, or [b] Firebird is not Interbase so I should have no problem installing both on the same machine. But in may ways Firebird still thinks it is named InterBase and gets confused.</p> <p>We should anyone installing Firebird to uninstall any version of InterBase, server or client, before installing Firebird. If possible, the Firebird install should check to see if InterBase is installed, and halt with an error message if so.</p>	5 - Medium (next major release)
269	224129	<p>Windows 98 Install Fails</p> <p>I downloaded fbinst_prod_110300-2.exe from SourceForge. I ran the .exe it on my computer, which is Windows 98. It is unsuccessful; it has errors. I suspect that the cause of the error is that whoever packaged fbinst_prod_110300-2.exe assumed Windows NT.</p> <p>(Previously I had installed ib_client on my machine, and found out the hard way that one must uninstall any InterBase package before installing Firebird; Firebird has an identity problem and gets confused. The errors here refer to after I thoroughly removed both Interbase and Firebird from my system and then tried to install Firebird again).</p> <p>The install left an error message n a DOS window named Finished - net: Error 2185: The service name is invalid. Make sure you are specifying a valid service name, and then try again.</p> <p>I ran IBConsole. When I try to contact the local server, it gives this error message: Error logging into the requested server Cannot attach to services manager. Ctrl+Alt+Del does not show a server running.</p> <p>The documentation says To start a server that has been shut down, run InterBase Server from the InterBase 5 Start menu. There is no Interase Server on the Windows 98 Start menu; under Programs / Firebird there are only IBConsole and License Agreement. MSCONFIG does not show any database server to be run upon startup. I presume that the command to start the server has been lost. I go to C:\Program Files\Firebird\bin and double click on ibserver.exe and now I have an icon on my task bar.</p> <p>I run IBConsole again and try to log in to the local server; now things seem to be OK.</p> <p>So the problem is that Firebird, in order to set the server to start automatically at boot, is trying to contact a services manager, which fails under Windows 98. The Firebird install needs to check the Windows version, and use an alternative method - e.g. add C:\Program Files\Firebird\bin\ibserver.exe to the StartUp menu - for Windows 95/98/ME.</p>	5 - Medium (next major release)
1374	436085	<p>isql internal error on SS0.9.4</p> <p>After installing Firebird 0.9.4 on a Linux RHat 7.1 i386 system isql gave the foloowing error message: Use CONNECT or CREATE DATABASE to specify a database SQL connect examples/employee.gdb; Internal error: Unexpected isc_info_value 0</p>	5 - Medium (next major release)

Category: Installation - Server

ID	SF ID	Description	Priority
272	224333	<p>Database: examples/employee.gdb, User: sysdba Server version too old to support the isql command SQL</p> <p>xinitd in Mandrake 7.2</p> <p>Mandrake 7.2 can be setup to use xinetd instead of the regular inetd in which case the /etc/inetd.conf file is not the right way to setup the server. Instead a file needs to be placed in /etc/xinet.d with basically the same information as goes in inetd.conf.</p>	5 - Medium (next major release)

Category: InterClient

ID	SF ID	Description	Priority
281	227414	<p>Sometimes Interbase/Interserver stop grant connections</p> <p>Sometimes I have this problem while asking for Interbase connection, sometimes releasing it. I cannot understand the exact way to reproduce it, but I attached a simple java program that loops until the error presents. If you want to try to reproduce it execute three of four instances of the program at the same time.</p> <p>The error is java.lang.VerifyError.</p> <p>I'm using:</p> <ul style="list-style-type: none"> - Interbase 6.0.1 - Interclient 2.0 beta (I tried IC1.6 but with the same result...) - Windows 2000 Professional - Sun JVM 1.3 <p>The same program works fine, without any error, for hours and hours, if I use the latest trial version of Easysoft ODBC driver</p>	5 - Medium (next major release)
222	212651	<p>Statement.setMaxRows has no effect if fetchSize==0 (default)</p> <p>- applies to InterClient1.6 and InterClient2.0</p> <p>- Statement.setMaxRows (PreparedStatement.setMaxRows respectively) have no effect if no fetchSize has been defined (fetchSize is 0 by default). In this case a ResultSet will return all rows of a query, even if the rowcount exceeds maxRows.</p> <p>- how to fix: <sniped, see bug report for details></p>	5 - Medium (next major release)
1472	447462	<p>memory leak problem in Statement.java</p> <p>There is a showstopping bug in the interclient.jar files that are distributed with all binaries. it's also in the sources in the cvs i hope they can be updated. it results in exhaustive memory usage when opening statements, then inserting data and then closing statements. it can result in such memory blow up that interbase simply stops working and even starts to corrupt the database.</p> <p>some code that will let you reproduce the bug:</p> <pre>for(int i=0;i1000;i++) { Statement st=aConn.createStatement(); st.executeUpdate(insert into myTable values (+i+)); st.close(); }</pre> <p>you'll see memory growing on ibserver.exe and it is never freed. however, there is a very simple solution to this problem. in executeUpdate() in Statement.java there is a variable called openOnServer_ set to false, but it should be set to true! after i made this change and recompiled the interclient.jar file everything works perfectly as expected.</p>	5 - Medium (next major release)

Category: Security Issues

ID	SF ID	Description	Priority
297	229237	<p>Blank passwords poorly supported</p> <p>The tools like isql poorly support having a user with a blank password. You are allowed to create one but then you can't ever connect using isql if you create the user that way. The obvious option of trying -p for the password doesn't work and is thrown out by the argument parser. Not supplying a password also doesn't work.</p> <p>The isql tool should allow -p for the password and try to connect using a blank password in that case.</p>	5 - Medium (next major release)
300	229894	<p>Client program can log on as any user.</p> <p>The client program by manipulation of the user_string property can log onto the server as any user including sysdba.</p>	5 - Medium (next major release)

Category: Security Issues

ID	SF ID	Description	Priority
250	222375	<p>The usage of the user_string passed from the client to the server in the form name:uid:gid and is used as the default logon userid if no explicit user name/password is supplied. This field is accepted without question by the server.</p> <p>If the user_string is blank when passed to the server (as apparently is the case for BDE), then (linux server anyway) defaults the user to the geteuid() of the server process.</p> <p>Grants overwrite previous rdb\$security_classes entries</p> <p>First, let me say that a clarification is need. If Firebird can rely solely on rdb\$user_privileges, then this report is not critical. Otherwise, the results of records being overwritten silently on rdb\$security_classes is an issue. The following is an almost unedited copy from an article I wrote on the NG:</p> <ul style="list-style-type: none"> - Create a table with 31-byte name. - Create a second table with the same name, except that it differs only in the last, 31th character. - Grant permissions to userA on the first table and to userB on the second table. - Go to see rdb\$user_privileges and you'll see your entries as they should be. - Go to rdb\$security_classes: with rdb\$ as a prefix, 31-4=27, oh what a brilliant person I am. <p>Here's your 27-byte limitation: what happened to this table? Both entries were mapped to the same record, same security class! As a result of that, the second GRANT overwrote the information for the first one, because rdb\$security_classes.rdb\$security_class used sql\$table_name hence the last 4 bytes were ignored. Feature? Undocumented limitation? Bug?</p> <p>So, what LangRef says could be refurbished as: Some objects, such as security class names, are restricted in practice to 27 bytes in length.</p> <p>The final part comes from Ivan Prenosil:</p> <p>Or perhaps something like Table names should not exceed 27 bytes, otherwise you can expect problems with privileges....</p> <p>I shortly played with these things and found out that -when creating new table, two entries are inserted into rdb\$security_classes, one SQL\$tab_name, one SQL\$DEFAULTx. -so I tried to create new table with name DEFAULT1 and it overwrote already existing row in rdb\$security_classes; I do not know what exact consequences it will cause though.</p> <p>C. (On behalf of Ivan, too.)</p>	5 - Medium (next major release)
259	223128	<p>SYSDBA can grant non existent roles</p> <p>IB doesn't check the user in a GRANT statement probably because the db can be moved to another server where such user is defined, since the information is stored in isc4.gdb only.</p> <p>But why would IB allow SYSDBA to grant no existent roles to users? For example, this is accepted:</p> <pre>grant anything to alice</pre> <p>However, anything doesn't exist in rdb\$roles but rdb\$user_privileges logs a role granted to alice.</p> <p>In contrast, a non-privileged user can't grant a role that doesn't exists. So I wonder is this is a bug or a feature. Why would SYSDBA need to bypass role checking when granting roles to users?</p>	5 - Medium (next major release)
1354	423810	<p>User Password length</p> <p>try to connect to a database with user=SYSDBA password=masterke It seems that password length is truncated to 8 char. Looking in isc4.gdb, PASSWD (in table users) is defined as PASSWD AS VARCHAR(32) CHARACTER SET OCTETS;.</p>	5 - Medium (next major release)