

ORACLE®




ORACLE®

Oracle Rdb Support's Examples and Experiences

Wolfgang Kobarg-Sachsse
Principal Technical Support Engineer

Preamble

- Selection of some Service Requests
- Some have an interesting solution
- From some even Rdb Support learned something
- Some have an easy solution that surprised the customer





Case #1

SR #:3-nnnnnnnnnnn

Summary: Poll command in thincontroller can break after bad regex filter qualifiers.

Description: Using bad regex input to the #name poll qualifier breaks polling in JDBC Thin Controller.



Case #1

- Regex? What is it?
- Found this in the Rdb JDBC User Guide Version 7.3:
 - See your JAVA documentation for more information on Regular Expressions.
- Feature in Java:
 - Java 4 (JDK 1.4) and later have comprehensive support for regular expressions through the standard `java.util.regex` package.
- In short: it is a feature of using wildcards in Java strings.
- We have implemented that in Rdb JDBC 7.3.



Case #1

- It can be used with the Rdb Thin Controller. This is a valid command:
 - `rdbthincontrol> poll #name:.*Srv1.*`
- The wildcard (regex) is “.*”, and the above command shows all JDBC servers where the name contains the substring “Srv1”.
- The following happens if there is an error in the regex, e.g. only the “*” instead of “.*” is used:
 - `rdbthincontrol> poll #name:*Srv1*`
`rdbthincontrol> Polling servers (using qualifiers : #name:*Srv1*)`
Dangling meta character '*' near index 0
`*Srv1*`
^



Case #1

- Until here, everything works as expected
- But then came the problem: Any following poll command returned NOTHING. So the incorrect regex broke something.
- Bug opened, bug fixed by Rdb Engineering in Rdb JDBC 7.3.0.1 and learned something that we did not know before: Regex features in Java and they are supported.

Service Request closed.



Case #2

SR #:3-nnnnnnnnnnn

Summary: JDBC server process stuck in LEF.

Description: We have a problem where JDBC server processes get stuck in LEF. When this happens any user connected to that server will get locked up and cannot do anything until we do a \$ STOP/ID on the server process. This problem happens several times per week.



Case #2

- One further symptom was that while one process stuck in LEF, one other access violated in a loop.
 - That is clearly buggy, but where?
 - We were not able to reproduce the issue.
 - Rdb Engineering and also HP support looked at this case with no result.
 - Happens only on Itanium.
-
- Eventually another customer logged this Service Request:



Case #2a

SR #:3-nnnnnnnnnnn

Summary: "ASTP1" routine in RDMSHRP72 module modifies R13 incorrectly.

Description: Looping Processes. The problem was reported and analyzed by OpenVMS Engineering. Their Finding is , "ASTP1" routine in RDMSHRP72 module modifies R13 incorrectly.



Case #2a

- Does not look similar to case #2 at the first glance.
- But here are more details of what is happening:
 - Process A tries constantly to insert a record into an Oracle Rdb table by violating a Unique Index (returning an appropriate error code).
 - When process B tries to do the same thing, process A reports access violation and jumps into the infinite exception handling loop, while process B comes to a hang.
 - The program must be linked with PTHREAD\$RTL, otherwise the problem does not happen.
 - Happens only on Itanium.
- So process A is access violating in a loop, and process B hangs – same issue as case #2.



Case #2a

- JAVA\$JAVA.EXE is always linked with PTHREAD\$RTL, so it is predestinated to show the symptoms.
- This time a reproducer was provided – one without Rdb JDBC involved.
- With this in our hands we went to Rdb Engineering.
- Rdb Engineering worked hand in hand with OpenVMS Engineering.
- Finally it came out that the bug was in OpenVMS, and HP fixed it!
- Both customers tested and confirmed the fix.

Both Service Requests closed.



Case #3

SR #:3-nnnnnnnnnnn

Summary: Rdb unicode, how to get Japanese characters via Oracle Rdb ODBC and display in Windows?

Description: In our database some fields are stored in different formats. One of the format is unicode. We know that the content is "correct Japanese". But when moving the characters from Rdb via the Oracle ODBC driver into Windows .NET the characters shown incorrectly.



Case #3

- A little bit more detailed:
- The application picks the string from the database and converts it to UTF8 when displayed.
- Now they want to do the same using the Rdb ODBC driver in a .Net application, but then the result is garbage.

Case #3

- We asked for more information, and received this:

- SQL> show character sets;
Default character set is DEC_MCS
National character set is DEC_MCS
Identifier character set is DEC_MCS
Literal character set is DEC_MCS
Display character set is UNSPECIFIED

- Columns for table TABLE:

Column Name	Data Type	Domain
-----	-----	-----
COL1	CHAR (8)	COL1
COL2	CHAR (35)	COL2
COL2_UNICODE	CHAR (105)	COL2_UNICODE



Case #3

- Conclusion from this information:
- Just calling a column COL2_UNICODE does not make it unicode, it is still DEC_MCS.
- You can store random bit patterns using DEC_MCS but these characters have no special meaning in the database (just binary bit patterns).
- If the application is knowing how to insert/select this bit pattern, then you get always the correct interpretation of the bit pattern.
- ODBC wants to interpret this bit pattern as real unicode, and that does not fit.

Case #3

- Solution would be to put the unicode characters into real unicode fields:

Column Name	Data Type	Domain
-----	-----	-----
COL1	CHAR (8)	COL1
COL2	CHAR (35)	COL2
COL2_UNICODE	CHAR (36)	COL2_UNICODE
UNICODE 18 Characters, 36 Octets		

- But that failed also. Why?



Case #3

- From the Rdb ODBC developer: "Technically speaking we don't support unicode in ODBC yet."
- Until there is support, what are the alternatives?
- Use the Rdb Data Provider for .Net instead of the ODBC Data Provider for .Net. That one is using the Rdb JDBC driver and supports unicode.
 - That was not possible at this point in time because there is not yet an official ORDP version available that works with Visual Studio 2010, but customer uses only that VS version.
 - Customer uses a self-built workaround that the .Net application interprets the bit pattern like their other application does.

Service Request closed.



Case #4

SR #:3-nnnnnnnnnnn

Summary: How to convert CHAR column to UTF8 CHAR?

Description: Existing Rdb database, with table TABTEST, and column COL1 CHAR(79). We would like to convert this column to UTF8 79 Characters, 237 Octets, and keep the data already in it. How to proceed ?



Case #4

- This one looks related to case #3, but it is different.
- As the steps are not obvious at the first glance, we showed the customer the example that follows.



Case #4

- At first, we created a table comparable to the one the customer has:

```
SQL> create database filename testdb;  
SQL> create table t1 (i integer, c char(10));  
SQL> insert into t1 values(1,'abcdef');  
1 row inserted  
SQL> commit;
```

Case #4

```
SQL> show table (col) t1;
```

```
Information for table T1
```

```
Columns for table T1:
```

Column Name	Data Type	Domain
-----	-----	-----
I	INTEGER	
C	CHAR(10)	

```
SQL> select * from t1;
```

```
      I      C
```

```
      1      abcdef
```

```
1 row selected
```

```
SQL>
```

Case #4

```
$ rmu/unload testdb t1 t1
```

```
%RMU-I-DATRECUNL,    1 data records unloaded
```

```
$ sql
```

```
SQL> att 'file testdb';
```

```
SQL> create global temporary table tt1 (i integer, c char(10));
```

```
SQL> alter table t1 alter column c char(30) character set utf8;
```

```
SQL> commit;
```

```
SQL> show table (col) t1;
```

```
Information for table T1
```

```
Columns for table T1:
```

Column Name	Data Type	Domain
-----	-----	-----
I	INTEGER	
C	CHAR(30)	
	UTF8 10 Characters,	30 Octets



Case #4

```
SQL> select * from t1;
```

```
%RDB-E-CONVERT_ERROR, invalid or unsupported data conversion  
-RDMS-E-CSETBADASSIGN, incompatible character sets prohibit the  
requested assignment
```

```
SQL> truncate table t1;
```

```
SQL> commit;
```

```
SQL> create trigger trttl after insert on ttl
```

```
cont> (insert into t1 values
```

```
cont> (ttl.i,
```

```
cont> translate(ttl.c using rdb$utf8))
```

```
cont> ) for each row;
```

```
SQL> commit;
```




Case #4

```
$ rmu/load testdb t1 t1
```

```
%RMU-F-FLDMUSMAT, Specified fields must match in number and  
datatype with the unloaded data
```

```
%RMU-I-DATRECSTO,    0 data records stored
```

```
%RMU-F-FTL_LOAD, Fatal error for LOAD operation
```

```
$ rmu/load testdb tt1 t1
```

```
%RMU-I-DATRECREAD,   1 data records read from input file.
```

```
%RMU-I-DATRECSTO,    1 data records stored
```



Case #4

```
$ sql
SQL> att 'file testdb';
SQL> select * from t1;
           I      C
          1  abcdef
1 row selected
SQL>
```

Service Request closed.



Case #5

SR #:3-nnnnnnnnnnn

Summary: Try to encrypt data in Rdb.

Description: Is there an easy way to encrypt sensitive data in some tables? We store confidential data, and we need to give a test DB to external developers. We want to create a copy of our production DB, with unrecognizable data in some columns in some tables, so persons cannot be identified any more.



Case #5

- The first investigation went into using the Caesar cipher (too simple), or the Vigenère cipher.
- But as this encryption method needed to be programmed this would mean some effort.
- It did not matter if the encrypted string remain inside the characters A-Z, thus the solution was to use the encryption method that OpenVMS provides.



Case #5

- We are providing a sample BLISS program in SQL\$SAMPLE called RDB_CYPHER.B32 that uses OpenVMS ENCRYPT.
- There is a note in My Oracle Support that describes how to use this sample program (it contains the sample program also in C).
- Customer was happy to use this encryption program (as an external procedure).

Case #5

- Example:

```
SQL> select coll from test_string;
```

```
COL1
```

```
My_Test_String
```

```
SQL> BEGIN
```

```
cont>     DECLARE :STRING_CYPHER CHAR (24);
```

```
cont>     DECLARE :STRING_PLAIN CHAR (24);
```

```
cont>     SELECT coll into :STRING_PLAIN FROM TEST_STRING;
```

```
cont>     CALL ENCRYPT_STRING ('THIS_KEY', :STRING_PLAIN, :STRING_CYPHER);
```

```
cont>     UPDATE TEST_STRING SET coll = :STRING_CYPHER;
```

```
cont> END;
```

```
SQL> select coll from test_string;
```

```
COL1
```

```
HÚ ³ñÒæĬh.:+2 @"...Ö;¥$ò
```



Case #5

- The next question was: We will also need to encrypt birthdays (date).
- This cannot be done by the same algorithm because the encrypted date must still fulfill the rules of a DATE field.



Case #5

- But that was easy to do:

```
SQL> select col1 from test_date;
```

```
COL1
```

```
21-SEP-1975 00:00:00.00
```

```
SQL> UPDATE TEST_DATE SET col1 = SCRAMBLE_DATE(col1);
```

```
1 row updated
```

```
SQL> select col1 from test_date;
```

```
COL1
```

```
11-JUN-1976 00:00:00.00
```




Case #5

- The SCRAMBLE_DATE function looks like:
 - function scramble_date (in :bday date vms)
returns date vms;
begin
declare :doy char(3);
declare :rc date vms;
set :doy = cast(extract(julian from :bday) as char(3));
set :rc = cast(cast(:bday as date ansi)+cast(:doy as interval day) as date vms);
return :rc;
end;

Service Request closed.



Case #6

SR #:3-nnnnnnnnnnn

Summary: # of Reserved Journals Go up Everytime when AIJ is re-enabled.

Description: After executing \$rmu /set after /aij_option=aij.opt, the number of reserved journals was added to the number of existing journals. Why?



Case #6

- To clarify this question:

```
$ type aij.opt
```

```
JOURNAL IS ENABLED -  
RESERVE 8 -
```

```
...
```

```
$ rmu /set after /aij_option=aij.opt mydb.rdb
```

```
$ rmu /show after mydb.rdb
```

```
JOURNAL IS ENABLED -  
RESERVE 17 -
```

```
...
```



Case #6

- That is the behaviour as it is described. So it is expected.
- Creating the AIJ options file using \$RMU/SHOW AFTER command, automatically puts the RESERVE n clause into the output.
- And if that options file is applied several times (after each DROP of the journals) then the number of reserved journal slots is increasing each time.
- Solution is easy: just drop the RESERVE clause from the options file, or set the number to 0.
- But:



Case #6

- The following happens if there is no RESERVE clause in the options file, or if the RESERVE value is 0:

\$ `rmu/set after /aij_option=aij.opt mydb.rdb`

%RMU-F-VALLSSMIN, value (0) is less than minimum allowed value (1) for RESERVE

%RMU-F-FTL_SET, Fatal error for SET operation



Case #6

- Problem has been logged as a bug.
- Workaround is to either specify a non zero reserve clause in the AIJ options file or do not use the AIJ options file with the "RMU/SET AFTER" command but use the "/ADD" qualifiers on the command line instead.
- Bug has been fixed in Rdb 7.2.5 and later.

Service Request closed.



Case #7

SR #:3-nnnnnnnnnnn

Summary: Differences between a View definition and the result of RMU extract.

Description: RMU adds a distinct clause to the create statement.



Case #7

- To clarify: A view is created as such:

```
SQL> create view test_view as  
cont> select col1 from test1  
cont> union  
cont> select col1 from test2;
```




Case #7

- Issue an RMU/EXTRACT. The following results appear:

```
create view TEST_VIEW  
  (COL1) as  
  select C2.COL1  
    from TEST1 C2  
 union distinct  
  select C3.COL1  
    from TEST2 C3;
```



Case #7

- The problem is that RMU/EXTRACT turns UNION into UNION DISTINCT.
- Technically, this is the SAME! UNION DISTINCT is the Default, in opposite to UNION ALL.
- RMU/EXTRACT always writes the complete clause into its output, because defaults might change over time (happens seldom, but happens).
- So there is effectively no problem at all...



Case #7

- An argument was mentioned about inconsistencies:
“There are problems with changing database from Oracle Rdb to 11g. We are not happy with this situation that it is sometimes UNION and sometimes UNION DISTINCT. But we will live with this situation.”
- It is important to note that this is an SQL default, not an Rdb default. For all DBMS. Rdb will not change this behavior.

Service Request closed.



Case #8

SR #:3-nnnnnnnnnnn

Summary: Data stored in different storage area than stated by SQL SHOW STORAGE AREA.

Description: SHOW STORAGE AREA myArea tells that data from table myTable are stored in this area. RMU/ANA/AREA tells that this area is empty. But the table has many rows. RMU/ANA/AREA of a different area finds that this area contains the data from myTable and from another table.



Case #8

- That is obviously strange. How shall SQL SHOW STORAGE AREA and RMU/ANALYZE/AREA return contradictory information?
- How shall we reproduce this behavior? This looks buggy, but we cannot log a bug with the symptoms only.
- Fortunately, we were able to acquire the details that resulted in this behavior so that a bug could be logged.



Case #8

- To sketch out what need to happen to run into this problem:
- Open the database on node A.
- Open the database on node B. Add two storage areas. Create two tables. For each one create a storage map that points into the separate storage area. Then close the database.
- Close the database on node A.
- Open the database on node B.



Case #8

- Open the database on node A. Drop the two tables and commit. Do not disconnect from the database.
- On node B, drop and recreate the storage areas with the same names.
- On node A, recreate the two tables and its two storage maps with the same names, then disconnect from the database.
- Now RMU/ANALYZE/AREA shows both tables in the same storage area.



Case #8

- When working with issues such as this acquiring information from the end user is critical to solving the problem!
- After acquiring the reproducer we were able to log a bug.
- Bug is fixed in Rdb 7.2.4 and later.

Service Request closed.



Case #9

SR #:3-nnnnnnnnnnn

Summary: Assignment of NULL not detected.

Description: Assignment of NULL not detected as expected using variable defined as NOT NULL.



Case #9

- To clarify: Create a domain, use it in the following way and the following expected error message is returned:

```
SQL> CREATE DOMAIN C_DOM AS CHAR(1)
```

```
cont> CHECK (VALUE IS NOT NULL) NOT DEFERRABLE;
```

```
SQL> commit;
```

```
SQL> declare :mychar c_dom;
```

```
SQL> begin
```

```
cont> select null into :mychar from rdb$database;
```

```
cont> end;
```

```
%RDB-E-NULL_NO_INDICAT, null can't be reported due to missing  
indicator
```



Case #9

- The next step is to create the following function:

```
SQL> create module TEST_MOD
cont> language SQL
cont> function GET_CHAR (in :inchar c_dom) returns C_DOM;
cont> begin
cont> declare :mychar c_dom;
cont> set :mychar = null;
cont> return :mychar;
cont> end;
cont> end module;
SQL> commit;
```



Case #9

- The result is unexpectedly a NULL and not an error message:

```
SQL> select get_char('A') from rdb$database;
```

NULL

1 row selected



Case #9

- We spoke with development because it was suspected this was expected behavior and not a bug.
- Assumptions were correct, this was the reply:
 - The CHECK is only used for columns based on the domain, not variables.
 - From the documentation for DECLARE variable statement

domain-name

Specifies the domain name assigned to the variable.

The domain supplies the data type and, for interactive SQL, the edit string of the variable.



Case #9

- So the solution is to add the CHECK constraint to the DECLARE variable statement:

...

```
declare :mychar c_dom DEFAULT :inchar  
        CHECK (VALUE IS NOT NULL) NOT DEFERRABLE;
```

...

- That solved the problem, now an error message is returned on a NULL assignment.
- Did this solve this issue? Unfortunately... No.



Case #9

- The customer returned with this information: “We have tested several examples, and all end up in bugcheck dumps.”
- Searching in our knowledge base for the exception that was shown in the bugcheck dump revealed that the customer hit a known bug that is fixed in Rdb 7.2.4.1.
- The behavior was viewed in Rdb 7.2.3.2.
- Upgrade necessary.

Service Request closed.



Case #10

SR #:3-nnnnnnnnnnn

Summary: Compatibility between RDB 7.0-5 and RDBMS 11R2.

Description: I would like to use DBlinks in a new project between RDBMS 11R2 and RDB 7.0-5. Is it compatible? Is there any sheet with a compatibility matrix?



Case #10

- What does this customer expect? Does his old TV display HD programs? 😊
- So the answer is NO, this combination is not compatible. Because:
 - Customer's version of SQL/Services is 7.1.5.7.3, and that is using the Oracle Link Version 8.0.5.1.0, which means that the OCI Services emulate an 8.0.5 database.
 - A server version 8.0.5 is not compatible to a client version 11.2.0.
 - Yes, we have compatibility matrixes for that. Article ID: 73657.1



Case #10

- Upgrade SQL/Services and Rdb to more modern versions is the solution. But the only one?
- No, there is a workaround that normally no customer would think about.
- On a Windows machine you can install
 - Rdb ODBC Driver (latest version)
 - Oracle Database Gateway for ODBC (latest version)
 - Then the database link in RDBMS 11.2 can connect to Rdb, using the old versions that the customer has. That has even been tested with Rdb 6.1.

Service Request closed.



Case #11

SR #:3-nnnnnnnnnnn

Summary: %SECURITY_LOGFAI, User xxxxxx failed to create a detached process.

Description: When client applications connect to Oracle Rdb databases via ODBC, on the server we receive the following error on the console:

Status: %SYSTEM-F-INVLOGIN, login information invalid at remote node



Case #11

- To clarify what has happened:
 - Customer uses MS Access and uses linked tables, using Rdb ODBC.
 - Each time when a linked table is used, it asks for the password, that is as expected and it works fine. Nevertheless there is the following message in the dispatcher log file:
 - %SQLSRV-E-AUTH_FAILURE, Authentication/authorization failure
 - After several successful table links, the SECURITY_LOGFAI error starts to happen.
 - This generates an intrusion record and the user becomes an intruder and the user gets DISUSER.
 - This started to happen after upgrading to SQL/Services 7.3.0.3.



Case #11

- This is expected behavior. From the SQL/Services Release Notes 7.3.0.3:

1.2.5 Scan Intrusion Security Now Supported

Oracle SQL/Services release 7.3.0.3 now supports intrusion detection via the OpenVMS Security Services. See the "OpenVMS Guide to System Security" manual for more information on this feature.



Case #11

- Going back to SQL/Services 7.3.0.2 is not a good solution. All future versions of SQL/Services will have this security checking, Oracle won't change that back because Rdb itself has this intrusion detection also included.
- Disabling intrusion detection on OpenVMS level would affect all connections to Rdb. This would open a security hole.
- Modifying the logicals that control the intrusion counters might relax the situation, but that is also not a good solution.
- Are there options?

Case #11

- We investigated why we have these additional %SQLSRV-E-AUTH_FAILURE errors, though the connection is done.
- This is the usual popup where the user enters the password:





Case #11

- We found out that Microsoft is doing a strange thing here.
- BEFORE the popup screen appears to ask for the password, ODBC tries to connect to the database WITHOUT a password.
- The popup window appears after an unsuccessful connection without the password.
- And that is causing the %SQLSRV-E-AUTH_FAILURE errors. And those are causing the intrusions for the user in UAF.



Case #11

- Can we ask Microsoft to change this ODBC behavior? We can, but there is no control of when/if it would be addressed.
- What are the options?
- We found one workaround. It is possible to leave out the username in the DSN.
- In this way we don't have an intrusion, just because there is no given username.
- But there is more effort for the user because it is necessary to enter username and password each time when the window pops up.



Case #11

- Again: can we ask Microsoft to change this ODBC behavior? We can try, but in the meantime, ODBC development can work around this behavior!
- And we did!
- The latest release of Rdb ODBC (version 3.3.2.1) contains an option which either allows to keep Microsoft's way of trying to connect at first without a password, or to deny this.

Service Request closed.



Case #12

SR #:3-nnnnnnnnnnnn

Summary: RDMSHRP72\RDMS\$\$DETACH_DATABASE +
00001DD0

Description: ***** Exception at 00000000810B3C90 :
RDMSHRP72\RDMS\$\$DETACH_DATABASE + 00001DD0
%RDB-F-WRONGRDB, RDB\$SHARE image is wrong
-RDMS-E-SIGNATURE_MISMA, invalid parameter signature on
procedure call
Saved PC = 00000000810B3DE0 :
RDMSHRP72\RDMS\$\$PARSE_INTCOM_BUFFER + 00000060



Case #12

- Since Rdb 7.0.0.0 there were various reports of this error:
RDB-F-WRONGRDB, RDB\$SHARE image is wrong
RDMS-F-OBSVER, obsolete version of database
Exception occurred at
RDMSHRP72\RDMS\$\$PARSE_INTCOM_BUFFER +
00000740
- We had many SRs with that error message, also some bugs.
- No customer was able to reproduce the problem by will. Rdb Engineering was also not able to build a reproducer.



Case #12

- In all cases, just restarting the application solved the issue.
- Some customers observed that problem only once or twice. Others had the problem more frequently, it was possible to reproduce it in customer's environment, but no customer was able to strip down their application for a reproducer.
- Based on the available information, engineering knew the problem was caused by a timing issue in the attach/detach code.



Case #12

- Rdb Engineering has rewritten the attach/deattach code in Rdb 7.2.4.
- That did not completely solve the problem.
- The error message changed to:
RDB-F-WRONGRDB, RDB\$SHARE image is wrong
RDMS-E-SIGNATURE_MISMA, invalid parameter signature on
procedure call
Exception occurred at
RDMSHRP72\RDMS\$\$DETACH_DATABASE + 00001DD0
Called from RDMSHRP72\RDMS\$\$PARSE_INTCOM_BUFFER +
00000060



Case #12

- The following short notice from June 2010 was the breakthrough in this long running and ugly issue:

“Oracle Rdb Engineering has created a reproducer.”

- The problem has been fixed in Rdb 7.2.4.2 and later.
- We had no further report of this problem if using Rdb 7.2.4.2 so we are very sure that it is really fixed.

Service Request closed.



Case #13

SR #:3-nnnnnnnnnnn

**Summary: Restored database from
RMU/BACKUP/ONLINE verifies with warnings**

**Description: Customer restored database to test
system and RMU/VERIFY returned warnings. Restored
database again but to different disks and RMU/VERIFY
returned the same warnings.**



Case #13

- Warnings like:

...

```
%RMU-W-DATNOTIDX, Row in table TABLE_NAME is not in any indexes.  
                Logical dbkey is 65:36632:0.
```

...

```
%RMU-W-BADIDXREL, Index INDEX_NAME either points to a non-existent record or  
                has multiple pointers to a record in table TABLE_NAME.  
                The logical dbkey in the index is 69:12932:92.
```

...

- Production database verified cleanly
- RMU/DUMP from production database of some data pages with warnings showed the pages were updated during backup



Case #13

- RMU/DUMPs of bad pages from backup file and restored database showed the same corruption
- Examination of AIJ that was active during backup looking for TSNs of updating transactions showed many lines erased and many more added to pages in question
- Wrote a C program attempting to simulate the work done; create array of rows for a key, delete all rows for that key, insert rows from the array
- After several attempts was able to reliably reproduce



Case #13

- Rdb Engineering made changes to RMU backup code
- This problem has existed for a long time
- Corrected in Oracle Rdb 7.2.4.2 update 1 and 7.2.5

Service Request closed.



Case #14

SR #:3-nnnnnnnnnnn

Summary: Database from RMU/COPY/ONLINE verifies with warnings

Description: Large DELETE transaction performed during copy while other normal database activity was ongoing



Case #14

- Modified reproducer from prior case to use RMU/COPY/ONLINE rather than RMU/BACKUP/ONLINE
- Problem reproduced with Oracle Rdb 7.2.4.2
- Did not reproduce with Oracle Rdb 7.2.4.2 update 1

Service Request closed.



Case #15

SR #:3-nnnnnnnnnnn

Summary: RMU/BACKUP/ONLINE bugchecks

Description: ***** Exception at 0000000080815060 :
RMU72\DI OBCK\$GET_PAGE + 00001820
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 000000008031D130 :
RMU72\RMUBCK\$STAREA_PAGES_ONLINE + 00000EB0



Case #15

- A subsequent attempt to perform backup may succeed
- Problem reproduces with same reproducer as the prior two cases with the exception that the database TSN must overflow out of the low 32 bits
- Rdb Engineering continues to work this problem

Service Request ongoing.

Conclusion

- The work on customer's problems is in most cases an interesting work, often a challenge, even with 10+ years of experience in Rdb Support.
- The cooperation within our Rdb Support team is great.
- The cooperation with the customer is in most cases very enjoyable because of the strong Oracle-to-Customer relationship in Rdb Support.
- When support has offered a solution, please be sure to confirm it so that we can share the news with ALL customers!!





Q&A

QUESTIONS
ANSWERS

ORACLE®