

Oracle® Rdb for OpenVMS

Table of Contents

<u>Oracle® Rdb for OpenVMS</u>	1
<u>Release Notes</u>	2
<u>April 2004</u>	3
<u>Contents</u>	4
<u>Preface</u>	5
<u>Purpose of This Manual</u>	6
<u>Intended Audience</u>	7
<u>Document Structure</u>	8
<u>Chapter 1 Installing Oracle Rdb Release 7.0.7.2</u>	9
<u>1.1 Requirements</u>	10
<u>1.2 Invoking VMSINSTAL</u>	11
<u>1.3 Stopping the Installation</u>	12
<u>1.4 After Installing Oracle Rdb</u>	13
<u>1.5 Maximum OpenVMS Version Check Added</u>	14
<u>Chapter 2 Software Errors Fixed in Oracle Rdb Release 7.0.7.2</u>	15
<u>2.1 Software Errors Fixed That Apply to All Interfaces</u>	16
<u>2.1.1 Wrong Index Retrieval is Selected in Query with Range List Predicates</u>	16
<u>2.1.2 Applications That Use \$HIBER/\$WAKE Hang in HIB</u>	17
<u>2.1.3 Bugchecks at PIOABW\$SYNCH PAGE + 00000564</u>	18
<u>2.1.4 Query Bugcheck After Logical RDMS\$SET FLAGS Set to SELECTIVITY(2)</u>	18
<u>2.1.5 Bugchecks or Corruption With Indexes of Type is Sorted Ranked</u>	19
<u>2.1.6 Query with Constant Column in UNION/GROUP BY Returns Wrong Results</u>	19
<u>2.1.7 Left Outer Join Query with Sub-select Overflows the Stack</u>	22
<u>2.1.8 Signal Failing in Compound Statement</u>	23
<u>2.1.9 Bugcheck in KOD\$ROLLBACK and PSII2REMOVEDUPBBC with Sorted Ranked Indexes</u>	23
<u>2.1.10 Database Corruption When LOCKING IS PAGE LEVEL Enabled</u>	24
<u>2.1.11 Table Constraint Should Fail on Updating a Row</u>	24
<u>2.1.12 Zigzag Match Query with Descending Index Segment Returns Wrong Results</u>	27
<u>2.2 SQL Errors Fixed</u>	30
<u>2.2.1 COMMIT or ROLLBACK Bugchecks in Conditional Expression</u>	30
<u>2.2.2 Unexpected Failure When Using ALTER ... THRESHOLDS Clause</u>	30

Table of Contents

<u>2.2 SQL Errors Fixed</u>	
<u>2.2.3 Unexpected Bugcheck when Oracle-style Outer Join Used with Subselect</u>	31
<u>2.2.4 %SQL-F-NODBFIL. Alias Missing a Declaration With SQLMOD</u>	31
<u>2.3 Oracle RMU Errors Fixed</u>	33
<u>2.3.1 RMU/CHECKPOINT Slow When Number of Nodes is One</u>	33
<u>2.3.2 RMU/LOAD/DEFER INDEX UPDATES ACCVIO with a Hashed Partitioned Index</u>	34
<u>2.3.3 RMU/RECOVER/ONLINE Without /JUST CORRUPT or /AREA Qualifiers Bugchecks</u>	35
<u>2.3.4 Invalid RMU-E-INASPAREA Message in RMU/VERIFY/ROOT</u>	36
<u>2.3.5 RMU Support For /DENSITY = SDLT320</u>	37
<u>2.3.6 Incorrect Verification of Invalid Reference DBKEYS in Ranked Indices</u>	37
<u>2.4 LogMiner Errors Fixed</u>	38
<u>2.4.1 RMU /UNLOAD /AFTER JOURNAL Incorrect NULL Bit Setting When VARCHAR is Last Column</u>	38
<u>2.5 Oracle Trace Errors Fixed</u>	40
<u>2.5.1 Oracle TRACE COLLECT FORMAT Command Could Not Create a Database After Rdb Upgrade</u>	40
<u>Chapter 3 Software Errors Fixed in Oracle Rdb Release 7.0.7.1</u>	41
<u>3.1 Software Errors Fixed That Apply to All Interfaces</u>	42
<u>3.1.1 Recovery of Empty Optimized AIJ Does Not Update the Sequence Number</u>	42
<u>3.1.2 Left Outer Join Query With OR Predicate Returns Wrong Results</u>	43
<u>3.1.3 Left Outer Join Query With OR Predicate Returns Wrong Results</u>	44
<u>3.1.4 Query Bugchecks When IN Clause Contains More Than Two DBKEYS</u>	45
<u>3.1.5 Processes Loop at IPL2 When VLM Feature Used</u>	46
<u>3.1.6 DBR Bugchecks at DBR\$DDTM RESOLVE + 000003F4</u>	46
<u>3.1.7 Replication Option and LogMiner Features Active at the Same Time</u>	47
<u>3.1.8 RMU /UNLOAD /AFTER JOURNAL Created .RRD Content Clarification</u>	47
<u>3.1.9 Page Locks Not Released When LOCKING IS PAGE LEVEL</u>	47
<u>3.1.10 Looping or Bugchecks in DIO\$FETCH DBKEY for SORTED RANKED Indexes</u>	48
<u>3.1.11 RMU/CLOSE/WAIT Hangs Waiting for ALS to Terminate</u>	49
<u>3.1.12 Query With EXISTS Clause and COMPUTED BY Column Returns Wrong Results</u>	49
<u>3.1.13 Bugcheck in DIO\$FREE CURRENT LOCK for Sorted Ranked Indexes</u>	53
<u>3.1.14 Ranked Index Overflow Node Corruption on Insert</u>	54
<u>3.1.15 Illegal Page Count Error in the Dynamic Optimizer</u>	55
<u>3.1.16 Bugcheck During Create Index with Mapping Values</u>	56
<u>3.1.17 Error %RDMS-E-NOSOL FOUND in Full Outer Join Query</u>	56
<u>3.1.18 TRUNCATE TABLE and RMU /REPAIR Corruption Corrected</u>	58
<u>3.1.19 UNION Query With Two Left Outer Joins in First Leg Returns Wrong Results</u>	59
<u>3.1.20 Logical Area Record Erasure Count Not Updated for Cached Rows</u>	62
<u>3.1.21 Query With Sum Function of Two Select Counts Bugchecks</u>	62
<u>3.1.22 Left Outer Join Query With CONCAT Function Returns Wrong Results</u>	63
<u>3.1.23 RCS Bugchecks at DIOCCH\$UNMARK GRIC ENT</u>	64
<u>3.1.24 Wrong Sort Order for Query with Aggregate, Group By, Order By</u>	65
<u>3.1.25 Left Outer Join Query with SUBSTRING and CHAR LENGTH Bugchecks</u>	66

Table of Contents

<u>3.1 Software Errors Fixed That Apply to All Interfaces</u>	
<u>3.1.26 Join Query with GROUP BY/ORDER BY Returns Wrong Order</u>	67
<u>3.1.27 Query Joining Two Derived Tables of a View with UNION Overflows the Stack</u>	68
<u>3.1.28 Query with Shared Expression in Two Predicates Returns Wrong Results</u>	69
<u>3.1.29 Wrong Index Retrieval is Selected in Query with GTR Predicate</u>	71
<u>3.1.30 Memory Corruption When Using Explicit 2PC</u>	72
<u>3.1.31 Query Applying Zero Shortcut Returns Wrong Results</u>	72
<u>3.2 SQL Errors Fixed</u>	74
<u>3.2.1 Unexpected DATEEQILL Error During IMPORT With CREATE INDEX or CREATE STORAGE MAP</u>	74
<u>3.2.2 Incorrect Unit for the DETECTED ASYNC PREFETCH THRESHOLD Option</u>	74
<u>3.2.3 DECLARE LOCAL TEMPORARY TABLE Limited to 10 Tables Per Session</u>	75
<u>3.2.4 IVP or Other Failure with Dynamic SQL if SQL\$INT is Installed /RESIDENT</u>	75
<u>3.2.5 Bugcheck at PSIINDEX\$FIND ENTS EXACT + 54</u>	75
<u>3.2.6 Multistatement Procedures Used with Connections Resulted in %RDB-E-OBSOLETE METADA Error Message</u>	76
<u>3.2.7 IMPORT May Generate ACCVIO Exception During Import of a Module</u>	76
<u>3.2.8 Unexpected ACCVIO When Reporting Incompatible Character Set Assignments</u>	76
<u>3.2.9 SQL-F-NODBFIL When SQL Modules are Compiled With /CONNECT</u>	77
<u>3.2.10 GET DIAGNOSTICS Keyword CONNECTION NAME Returned Incorrect Value</u>	77
<u>3.2.11 Errors Not Reported by SQL</u>	78
<u>3.2.12 Variable Updated Though No Rows Found</u>	78
<u>3.2.13 Partitioning Clause Not Working in Embedded SQL</u>	79
<u>3.3 RDO and RDML Errors Fixed</u>	82
<u>3.3.1 RDML /DATE TYPE Qualifier Default is Now NOEMPTY RECORDS</u>	82
<u>3.4 Oracle RMU Errors Fixed</u>	83
<u>3.4.1 RMU /COLLECT May Default to a READ WRITE Transaction</u>	83
<u>3.4.2 RMU/VERIFY/CONSTRAINTS Problems With Named Tables And Constraints</u>	84
<u>3.4.3 RMU Unload Incorrectly Using DBKEY SCOPE IS ATTACH</u>	86
<u>3.5 LogMiner Errors Fixed</u>	87
<u>3.5.1 LogMiner Elimination of Processing Unneeded AIJ Files</u>	87
<u>3.5.2 /TRANSACTION TYPE Qualifier for RMU /UNLOAD /AFTER JOURNAL</u>	87
<u>3.6 RMU Show Statistics Errors Fixed</u>	89
<u>3.6.1 RMU /SHOW STATISTICS Writes Invalid Configuration File</u>	89
<u>3.7 Hot Standby Errors Fixed</u>	90
<u>3.7.1 Starting LRS on Master Database Caused Shutdown</u>	90
<u>Chapter 4 Enhancements</u>	91
<u>4.1 Enhancements Provided in Oracle Rdb Release 7.0.7.2</u>	92
<u>4.1.1 Rdb Optional Site-Specific Startup Procedure</u>	92
<u>4.1.2 Oracle Rdb SGA API</u>	92

Table of Contents

<u>4.1 Enhancements Provided in Oracle Rdb Release 7.0.7.2</u>	
<u>4.1.3 CHRONO FLAG Replaces Older CRONO FLAG Keyword</u>	93
<u>4.2 Enhancements Provided in Oracle Rdb Release 7.0.7.1</u>	94
<u>4.2.1 RDMSBIND SNAP QUIET POINT Logical No Longer Used</u>	94
<u>4.2.2 Determining Which Oracle Rdb Options Are Installed</u>	94
<u>4.2.3 New Procedure RDB\$IMAGE VERSIONS.COM</u>	95
<u>Chapter 5 Documentation Corrections</u>	96
<u>5.1 Documentation Corrections</u>	97
<u>5.1.1 Database Server Process Priority Clarification</u>	97
<u>5.1.2 Waiting for Client Lock Message</u>	97
<u>5.1.3 Clarification of PREPARE Statement Behavior</u>	98
<u>5.1.4 SQL EXPORT Does Not Save Some Database Attributes</u>	99
<u>5.1.5 RDMSBIND LOCK TIMEOUT INTERVAL Overrides the Database Parameter</u>	99
<u>5.1.6 New Request Options for RDO, RDBPRE and RDB\$INTERPRET</u>	100
<u>5.1.7 Missing Descriptions of RDB\$FLAGS from HELP File</u>	102
<u>5.1.8 A Way to Find the Transaction Type of a Particular Transaction Within the Trace Database</u>	104
<u>5.1.9 Clarification of SET FLAGS Option DATABASE PARAMETERS</u>	104
<u>5.1.10 Additional Information About Detached Processes</u>	105
<u>5.1.11 The Halloween Problem</u>	106
<u>5.1.12 RDMSBIND MAX DBR COUNT Documentation Clarification</u>	107
<u>5.1.13 RMU /UNLOAD /AFTER JOURNAL NULL Bit Vector Clarification</u>	108
<u>5.1.14 Location of Host Source File Generated by the SQL Precompilers</u>	111
<u>5.1.15 Suggestion to Increase GH RSRVPGCNT Removed</u>	112
<u>5.1.16 Clarification of the DDLDONOTMIX Error Message</u>	112
<u>5.1.17 Compressed Sorted Index Entry Stored in Incorrect Storage Area</u>	113
<u>5.1.18 Partition Clause is Optional on CREATE STORAGE MAP</u>	115
<u>5.1.19 Oracle Rdb Logical Names</u>	115
<u>5.1.20 Documentation Error in Oracle Rdb Guide to Database Performance and Tuning</u>	115
<u>5.1.21 SET FLAGS Option IGNORE OUTLINE Not Available</u>	116
<u>5.1.22 SET FLAGS Option INTERNALS Not Described</u>	116
<u>5.1.23 Documentation for VALIDATE ROUTINE Keyword for SET FLAGS</u>	117
<u>5.1.24 Documentation for Defining the RDBSERVER Logical Name</u>	117
<u>5.1.25 Undocumented SET Commands and Language Options</u>	118
<u>5.1.25.1 QUIET COMMIT Option</u>	118
<u>5.1.25.2 COMPOUND TRANSACTIONS Option</u>	119
<u>5.1.26 Undocumented Size Limit for Indexes with Keys Using Collating Sequences</u>	120
<u>5.1.27 Changes to RMU/REPLICATE AFTER/BUFFERS Command</u>	120
<u>5.1.28 Change in the Way RDMAIJ Server is Set Up in UCX</u>	121
<u>5.1.29 CREATE INDEX Supported for Hot Standby</u>	122
<u>5.1.30 Dynamic OR Optimization Formats</u>	122
<u>Chapter 6 Known Problems and Restrictions</u>	123

Table of Contents

6.1 Oracle Rdb Considerations	124
6.1.1 AIJ Log Server Process May Loop Or Bugcheck.....	124
6.1.2 Optimization of Check Constraints.....	124
6.1.3 Dynamic Optimization Estimation Incorrect for Ranked Indices.....	127
6.1.4 Running Rdb Applications With the VMS Heap Analyzer.....	127
6.1.5 RMU/RECOVER/AREA Needs Area List.....	128
6.1.6 PAGE TRANSFER VIA MEMORY Disabled.....	128
6.1.7 RMU/VERIFY Reports PGSPAMENT or PGSPMCLST Errors.....	128
6.1.8 Behavior Change in 'With System Logical Name Translation' Clause.....	129
6.1.9 Carry-Over Locks and NOWAIT Transactions Clarification.....	130
6.1.10 Strict Partitioning May Scan Extra Partitions.....	130
6.1.11 Exclusive Access Transactions May Deadlock With RCS Process.....	131
6.1.12 Oracle Rdb and OpenVMS ODS-5 Volumes.....	131
6.1.13 Clarification of the USER Impersonation Provided by the Oracle Rdb Server.....	132
6.1.14 Index STORE Clause WITH LIMIT OF Not Enforced in Single Partition Map.....	132
6.1.15 Unexpected NO META UPDATE Error Generated by DROP MODULE ... CASCADE When Attached by PATHNAME.....	133
6.1.16 Application and Oracle Rdb Both Using SYSSHIBER.....	133
6.1.17 IMPORT Unable to Import Some View Definitions.....	134
6.1.18 AIJSERVER Privileges.....	135
6.1.19 Lock Remastering and Hot Standby.....	136
6.1.20 RDB_SETUP Privilege Error.....	136
6.1.21 Starting Hot Standby on Restored Standby Database May Corrupt Database.....	136
6.1.22 Restriction on Compound Statement Nesting Levels.....	137
6.1.23 Back Up All AIJ Journals Before Performing a Hot Standby Switchover Operation.....	138
6.1.24 Concurrent DDL and Read-Only Transaction on the Same Table Not Compatible.....	138
6.1.25 Oracle Rdb and the SRM_CHECK Tool.....	139
6.1.26 Oracle RMU Checksum Verification Qualifier.....	139
6.1.27 Do Not Use HYPERSORT with RMU/OPTIMIZE/AFTER JOURNAL (Alpha).....	140
6.1.28 Restriction on Using /NOONLINE with Hot Standby.....	140
6.1.29 SELECT Query May Bugcheck with PSII2SCANGETNEXTBBCDUPLICATE Error.....	141
6.1.30 DBAPack for Windows 3.1 is Deprecated.....	141
6.1.31 Determining Mode for SQL Non-Stored Procedures.....	141
6.1.32 DROP TABLE CASCADE Results in %RDB-E-NO META UPDATE Error.....	143
6.1.33 Bugcheck Dump Files with Exceptions at COSI CHF SIGNAL.....	144
6.1.34 Interruptions Possible when Using Multistatement or Stored Procedures.....	145
6.1.35 Row Cache Not Allowed on Standby Database While Hot Standby Replication Is Active.....	146
6.1.36 Hot Standby Replication Waits when Starting if Read-Only Transactions Running.....	146
6.1.37 Error when Using the SYS\$LIBRARY:SQL_FUNCTIONS70.SQL Oracle Functions Script.....	146
6.1.38 DEC C and Use of the /STANDARD Switch.....	147
6.1.39 Excessive Process Page Faults and Other Performance Considerations During Oracle Rdb Sorts.....	147
6.1.40 Performance Monitor Column Mislabeled.....	149
6.1.41 Restriction Using Backup Files Created Later than Oracle Rdb Release 7.0.1.....	149
6.1.42 RMU Backup Operations and Tape Drive Types.....	149
6.1.43 Use of Oracle Rdb from Shared Images.....	150
6.1.44 Restriction Added for CREATE STORAGE MAP on Table with Data.....	150

Table of Contents

<u>6.1 Oracle Rdb Considerations</u>	151
6.1.45 Oracle Rdb Workload Collection Can Stop Hot Standby Replication.....	151
6.1.46 RMU Convert Command and System Tables.....	152
6.1.47 Converting Single-File Databases.....	152
6.1.48 Restriction when Adding Storage Areas with Users Attached to Database.....	153
6.1.49 Support for Single-File Databases to be Dropped in a Future Release.....	153
6.1.50 DECdtm Log Stalls.....	153
<u>6.1.51 Cannot Run Distributed Transactions on Systems with DECnet/OSI and OpenVMS Alpha Version 6.1 or OpenVMS VAX Version 6.0.....</u>	154
6.1.52 Multiblock Page Writes May Require Restore Operation.....	155
6.1.53 Replication Option Copy Processes Do Not Process Database Pages Ahead of an Application.....	155
6.1.54 SQL Does Not Display Storage Map Definition After Cascading Delete of Storage Area.....	156
6.1.55 ARITH EXCEPT or Incorrect Results Using LIKE IGNORE CASE.....	156
6.1.56 Different Methods of Limiting Returned Rows from Queries.....	156
6.1.57 Suggestions for Optimal Usage of the SHARED DATA DEFINITION Clause for Parallel Index Creation.....	158
6.1.58 Side Effect when Calling Stored Routines.....	159
6.1.59 Considerations when Using Holdable Cursors.....	160
6.1.60 INCLUDE SQLDA2 Statement Is Not Supported for SQL Precompiler for PL/I in Oracle Rdb Release 5.0 or Higher.....	161
6.1.61 SQL Pascal Precompiler Processes ARRAY OF RECORD Declarations Incorrectly.....	161
6.1.62 RMU Parallel Backup Command Not Supported for Use with SLS.....	162
<u>6.2 Oracle CDD/Repository Restrictions.....</u>	163
6.2.1 Oracle CDD/Repository Compatibility with Oracle Rdb Features.....	163
6.2.2 Multischema Databases and CDD/Repository.....	164
6.2.3 Interaction of Oracle CDD/Repository Release 5.1 and Oracle RMU Privileges Access Control Lists.....	165
6.2.3.1 Installing the Corrected CDDSHR Images.....	166
6.2.3.2 CDD Conversion Procedure.....	167

Oracle® Rdb for OpenVMS

Release Notes

Release 7.0.7.2

April 2004

Oracle Rdb Release Notes, Release 7.0.7.2 for OpenVMS

Copyright © 1984, 2004 Oracle Corporation. *All rights reserved.*

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent and other intellectual and industrial property laws. Reverse engineering, disassembly or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of Oracle Corporation.

If the Programs are delivered to the U.S. Government or anyone licensing or using the programs on behalf of the U.S. Government, the following notice is applicable:

Restricted Rights Notice Programs delivered subject to the DOD FAR Supplement are "commercial computer software" and use, duplication, and disclosure of the Programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, Programs delivered subject to the Federal Acquisition Regulations are "restricted computer software" and use, duplication, and disclosure of the Programs shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software – Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and Oracle Corporation disclaims liability for any damages caused by such use of the Programs.

Oracle is a registered trademark, and Hot Standby, LogMiner for Rdb, Oracle CDD/Repository, Oracle CODASYL DBMS, Oracle Expert, Oracle Rdb, Oracle RMU, Oracle RMUwin, Oracle SQL/Services, Oracle Trace, and Rdb7 are trademark or registered trademarks of Oracle Corporation. Other names may be trademarks of their respective owners.

Contents

Preface

Purpose of This Manual

This manual contains release notes for Oracle Rdb Release 7.0.7.2. The notes describe changed and enhanced features; upgrade and compatibility information; new and existing software problems and restrictions; and software and documentation corrections. These release notes cover both Oracle Rdb for OpenVMS Alpha and Oracle Rdb for OpenVMS VAX, which are referred to by their abbreviated name, Oracle Rdb.

Intended Audience

This manual is intended for use by all Oracle Rdb users. Read this manual before you install, upgrade, or use Oracle Rdb Release 7.0.7.2.

Document Structure

This manual consists of six chapters:

<u>Chapter 1</u>	Describes how to install Oracle Rdb Release 7.0.7.2.
<u>Chapter 2</u>	Describes software errors corrected in Oracle Rdb Release 7.0.7.2.
<u>Chapter 3</u>	Describes software errors corrected in Oracle Rdb Release 7.0.7.1.
<u>Chapter 4</u>	Describes enhancements introduced in Oracle Rdb Release 7.0.7.2.
<u>Chapter 5</u>	Provides information not currently available in the Oracle Rdb documentation set.
<u>Chapter 6</u>	Describes problems, restrictions, and workarounds known to exist in Oracle Rdb Release 7.0.7.2.

Chapter 1

Installing Oracle Rdb Release 7.0.7.2

This software update is installed using the standard OpenVMS Install Utility.

NOTE

Beginning with Release 7.0.6.2 of Oracle Rdb, all new Oracle Rdb kits released are full kits. Oracle will no longer ship partial kits (known as ECOs in the past). Therefore, there is no need to install any prior release of Oracle Rdb when installing new Rdb kits.

1.1 Requirements

The following conditions must be met in order to install this software update:

- Oracle Rdb must be shutdown before you install this update kit. That is, the command file `SY$STARTUP:RMONSTOP(70).COM` should be executed before proceeding with this installation. If you have an OpenVMS cluster, you must shutdown all versions of Oracle Rdb on all nodes in the cluster before proceeding.
- The installation requires approximately 130,000 free blocks on your system disk for OpenVMS VAX systems; 240,000 blocks for OpenVMS Alpha systems.

1.2 Invoking VMSINSTAL

To start the installation procedure, invoke the VMSINSTAL command procedure:

```
@SYS$UPDATE:VMSINSTAL variant-name device-name OPTIONS N
```

variant-name

The variant names for the software update for Oracle Rdb Release 7.0.7.2 are:

- RDBSE2G070 for Oracle Rdb for OpenVMS VAX standard version.
- RDBASE2G070 for Oracle Rdb for OpenVMS Alpha standard version.
- RDBMVE2G070 for Oracle Rdb for OpenVMS VAX multiversion.
- RDBAMVE2G070 for Oracle Rdb for OpenVMS Alpha multiversion.

device-name

Use the name of the device on which the media is mounted.

- If the device is a disk drive, such as a CD-ROM reader, you also need to specify a directory. For CD-ROM distribution, the directory name is the same as the variant name. For example:

```
DKA400:[RDBAMVE2G070.KIT]
```

- If the device is a magnetic tape drive, you need to specify only the device name. For example:

```
MTA0:
```

OPTIONS N

This parameter prints the release notes.

The following example shows how to start the installation of the Alpha multiversion kit on device MTA0: and print the release notes:

```
$ @SYS$UPDATE:VMSINSTAL RDBAMVE2G070 MTA0: OPTIONS N
```

1.3 Stopping the Installation

To stop the installation procedure at any time, press Ctrl/Y. When you press Ctrl/Y, the installation procedure deletes all files it has created up to that point and exits. You can then start the installation again.

If VMSINSTAL detects any problems during the installation, it notifies you and a prompt asks if you want to continue. You might want to continue the installation to see if any additional problems occur. However, the copy of Oracle Rdb installed will probably not be usable.

1.4 After Installing Oracle Rdb

This update provides a new Oracle Rdb Oracle TRACE facility definition. Any Oracle TRACE selections that reference Oracle Rdb will need to be redefined to reflect the new facility version number for the updated Oracle Rdb facility definition, "RDBVMSV7.0-72".

If you have Oracle TRACE installed on your system and you would like to collect for Oracle Rdb, you must insert the new Oracle Rdb facility definition included with this update kit.

The installation procedure inserts the Oracle Rdb facility definition into a library file called EPC\$FACILITY.TLB. To be able to collect Oracle Rdb event-data using Oracle TRACE, you must move this facility definition into the Oracle TRACE administration database. Perform the following steps:

1. Extract the definition from the facility library to a file (in this case, RDBVMS.EPC\$DEF).

```
$ LIBRARY /TEXT /EXTRACT=RDBVMSV7.0-72 -  
_ $ /OUT=RDBVMS.EPC$DEF SYS$SHARE:EPC$FACILITY.TLB
```

2. Insert the facility definition into the Oracle TRACE administration database.

```
$ COLLECT INSERT DEFINITION RDBVMS.EPC$DEF /REPLACE
```

Note that if you are installing the multiversion variant of Oracle Rdb, the process executing the INSERT DEFINITION command must use the version of Oracle Rdb that matches the version used to create the Oracle TRACE administration database or the INSERT DEFINITION command will fail.

1.5 Maximum OpenVMS Version Check Added

As of Oracle Rdb Release 7.0.1.5, a maximum OpenVMS version check has been added to the product. Oracle Rdb has always had a minimum OpenVMS version requirement. With 7.0.1.5 and for all future Oracle Rdb releases, we have expanded this concept to include a maximum VMS version check and a maximum supported processor hardware check. The reason for this check is to improve product quality.

OpenVMS Version 7.3–x is the maximum supported version of OpenVMS.

As of Oracle Rdb Release 7.0.3, the Alpha EV6 processor is supported. As of Oracle Rdb Release 7.0.5, the Alpha EV67 processor is supported. As of Oracle Rdb Release 7.0.6, the Alpha Wildfire processor is supported (see <http://metalink.oracle.com> for specifics on which Wildfire configurations are supported). As of Oracle Rdb Release 7.0.6.2, the Alpha EV68 processor is supported. As of Oracle Rdb Release 7.0.7, the Alpha EV7 processor is supported.

The check for the OpenVMS operating system version and supported hardware platforms is performed both at installation time and at runtime. If either a non–certified version of OpenVMS or hardware platform is detected during installation, the installation will abort. If a non–certified version of OpenVMS or hardware platform is detected at runtime, Oracle Rdb will not start.

Chapter 2

Software Errors Fixed in Oracle Rdb Release 7.0.7.2

This chapter describes software errors that are fixed by Oracle Rdb Release 7.0.7.2.

2.1 Software Errors Fixed That Apply to All Interfaces

2.1.1 Wrong Index Retrieval is Selected in Query with Range List Predicates

Bug 3243452

The following query slows down significantly when the range list index retrieval [0:1,1:0] is wrongly applied instead of [3:3] for greater and less than predicates.

```
set flags 'strategy,detail';
select count(*)
  from product_division pd
       inner join product_warehouse pw on
           pw.company_no = pd.company_no and pw.division_no =
           pd.division_no and pw.product_no = pd.product_no
       left outer join wm_location wml on
           wml.company_no = pd.company_no
           and wml.division_no = pd.division_no
           and wml.warehouse_no = pw.warehouse_no
           and wml.slot_id = pw.slot_id
where pd.company_no = 1
and pw.division_no = 1
and pd.product_no not between 700000 and 799999
and pd.product_no < 900000
and pd.product_status_cd not in ('D','O','R');
Tables:
  0 = PRODUCT_DIVISION
  1 = PRODUCT_WAREHOUSE
  2 = WM_LOCATION
Aggregate: 0:COUNT (*)
Conjunct: (0.PRODUCT_NO < 700000) OR (0.PRODUCT_NO > 799999)
Cross block of 2 entries          (Left Outer Join)
Cross block entry 1
  Cross block of 2 entries
    Cross block entry 1
      Conjunct: 0.COMPANY_NO = 1
      Conjunct: (0.PRODUCT_NO < 700000) OR (0.PRODUCT_NO > 799999)
      Conjunct: 0.PRODUCT_NO < 900000
      Conjunct: (0.PRODUCT_STATUS_CD <> 'D') AND (0.PRODUCT_STATUS_CD <> 'O')
      Conjunct: 0.PRODUCT_STATUS_CD <> 'R'
      Leaf#01 NdxOnly 0:PRODUCT_DIVISION Card=153560
        Bool: 0.DIVISION_NO = 1
        FgrNdx  PRODUCT_DIVISION_CMP_NDX [2:2] Fan=33
          Keys: (0.DIVISION_NO = 1) AND (0.COMPANY_NO = 1)
        BgrNdx1  PRODUCT_DIVISION_PRD_NDX [0:1] Fan=33
          Keys: 0.PRODUCT_NO < 900000
    Cross block entry 2
      Conjunct: 1.DIVISION_NO = 1
      Leaf#02 BgrOnly 1:PRODUCT_WAREHOUSE Card=153560
        Bool: (1.COMPANY_NO = 0.COMPANY_NO) AND (1.DIVISION_NO = 0.DIVISION_NO
              ) AND (1.PRODUCT_NO = 0.PRODUCT_NO)
        BgrNdx1  PRODUCT_WAREHOUSE_PRIMARY [0:1,1:0] Fan=33 <= Should be [3:3]
          Keys: r0: 1.PRODUCT_NO > 799999
              r1: 1.PRODUCT_NO < 700000
        Bool: (1.PRODUCT_NO < 900000) AND (1.COMPANY_NO = 1) AND (
```

Oracle® Rdb for OpenVMS

```
1.DIVISION_NO = 1)
Cross block entry 2
Conjunct: (2.COMPANY_NO = 0.COMPANY_NO) AND (2.DIVISION_NO = 0.DIVISION_NO)
          AND (2.WAREHOUSE_NO = 1.WAREHOUSE_NO) AND (2.SLOT_ID = 1.SLOT_ID)
Index only retrieval of relation 2:WM_LOCATION
Index name WM_LOCATION_PMRV [4:4]
Keys: (2.SLOT_ID = 1.SLOT_ID) AND (2.WAREHOUSE_NO = 1.WAREHOUSE_NO) AND
      (2.DIVISION_NO = 0.DIVISION_NO) AND (2.COMPANY_NO = 0.COMPANY_NO)

0
1 row selected
```

The query applies [3:3] index retrieval if the SQL flag 'SELECTIVITY' is enabled.

The strategy output is similar to the above except for the following lines pointed to by <= in the inner Cross block entry 2.

```
Cross block entry 2
Conjunct: 1.DIVISION_NO = 1
Leaf#02 BgrOnly 1:PRODUCT_WAREHOUSE Card=153560
  Bool: (1.COMPANY_NO = 0.COMPANY_NO) AND (1.DIVISION_NO = 0.DIVISION_NO
        ) AND (1.PRODUCT_NO = 0.PRODUCT_NO)
BgrNdx1 PRODUCT_WAREHOUSE_PRIMARY [3:3] Fan=33          <=
  Keys: (1.DIVISION_NO = 0.DIVISION_NO) AND (1.PRODUCT_NO <=
        0.PRODUCT_NO) AND (1.COMPANY_NO = 0.COMPANY_NO) <=
  Bool: (1.PRODUCT_NO < 900000) AND ((1.PRODUCT_NO < 700000) OR ( <=
        1.PRODUCT_NO > 799999)) AND (1.COMPANY_NO = 1) AND ( <=
        1.DIVISION_NO = 1)
```

This is a regression caused by the fix for Bug 2634849 in Oracle Rdb Release 7.1.2.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.2 Applications That Use \$HIBER/\$WAKE Hang in HIB

Bug 2881846

User applications that utilize the OpenVMS \$HIBER/\$WAKE system services to process asynchronous events could hang in "HIB" state if the database had the FAST COMMIT feature enabled. The hung process would resume executing normally if a \$WAKE was issued against it by another process.

For example, the Oracle SQL/Services product utilizes \$HIBER/\$WAKE to coordinate events between the server processes (dispatcher and executor). The SQL/Services executor processes could sometimes hang in "HIB" state.

This problem would sometimes occur when a global checkpoint request was issued. A global checkpoint will occur whenever Oracle Rdb switches to another journal, or when a checkpoint is manually requested by issuing an RMU/CHECKPOINT command. A race condition between the database checkpoint activity and the application's usage of \$WAKE could cause a \$WAKE intended for the application to be consumed by the database checkpoint activity, preventing the application from properly waking from its hibernate state.

This problem can be avoided by disabling the fast commit feature. Note that this can have a significant impact on performance.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.3 Bugchecks at PIOABW\$SYNCH_PAGE + 00000564

Bug 3264272

If a process did not have sufficient quota it was possible to encounter bugchecks like the following:

```
***** Exception at 0109556C : PIOABW$SYNCH_PAGE + 00000564
%RDMS-F-CANTWRITEDBS, error writing pages 2:367-369
-SYSTEM-F-EXQUOTA, process quota exceeded
```

While in general it is necessary to have sufficient quotas to support the total number of potential concurrent disk I/Os possible, certain operations, like asynchronous batch-writes, can safely ignore these errors. If an asynchronous batch-write encounters an EXQUOTA error, the disk write can be retried later when another attempt is made to write the buffer. The asynchronous batch-write operations have been modified to tolerate EXQUOTA errors.

Note that there are many other operations that may still fail with a bugcheck if an EXQUOTA error is encountered when attempting to read or write database disk files. For example, EXQUOTA errors encountered when attempting I/O to the recovery-unit journal (.RUJ) or after-image journal (.AIJ) will still result in a fatal bugcheck. Also, if an attempt to write out all modified buffers is unable to start any I/Os successfully, then a bugcheck will still occur.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.4 Query Bugcheck After Logical RDMS\$SET_FLAGS Set to SELECTIVITY(2)

Bug 3242615

Before the problem was fixed, a query resulted in a bugcheck when the query was executed after the sampled selectivity capability was enabled. Sampled selectivity can be enabled in any of several ways, one of which is to define the VMS logical name RDMS\$SET_FLAGS to be SELECTIVITY(2). Below is an example of a failing query:

```
select count (*) from xxx
  where (c in (-1, 0) and b > 48955) or
        (a > 30000 and b > 49990) or
        (a > 49900 and b > 20000);
```

The problem occurs when there is an OR condition in the WHERE clause of the query.

As a workaround, enable sampled selectivity only for specific queries rather than for all queries within an application. This can be done by using the OPTIMIZE WITH SAMPLED SELECTIVITY clause on specific queries rather than defining RDMS\$SET_FLAGS as SELECTIVITY(2).

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.5 Bugchecks or Corruption With Indexes of Type is Sorted Ranked

Bug 3009262

When updating an index of *TYPE IS SORTED RANKED*, it was possible that a bugcheck or index corruption could occur. This problem could occur if the index being updated allowed duplicates and was most likely when there existed a large number of duplicates for a particular key value and the duplicate rows were widely distributed in the database.

When a key value was changed or a row was deleted requiring the removal of a dbkey from a duplicates chain, it was possible that a bugcheck dump could be generated.

The following example shows an update that causes a key value to change. During the attempt to remove the dbkey for the affected row from the appropriate duplicates chain, Rdb generates a bugcheck dump.

```
SQL> update tt11 set f1 = 'b' where f1 = 'a' and f3 = 5 and f2 < 125;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file
USER1:[BUGCHECK_DIR]RDSBUGCHK.DMP;
%COSI-F-BUGCHECK, internal consistency failure
```

The bugcheck could contain either of the two exceptions shown in the following example.

```
***** Exception at 00A32E98 : PSII2SCANGETNEXTBBCDUPLICATE + 00000120
%COSI-F-BUGCHECK, internal consistency failure

***** Exception at 00C53E4C : PSII2REMOVEDUPBBC + 0000130C
%COSI-F-BUGCHECK, internal consistency failure
```

Under rare circumstances, it was possible that the insertion of a dbkey into a duplicates chain could cause the duplicates chain to become corrupt. The corruption could be removed by dropping and recreating the offending index.

As a workaround, the problem can be avoided by using alternate index types, or by adding columns to the index to make it more unique.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.6 Query with Constant Column in UNION/GROUP BY Returns Wrong Results

Bug 3336416

The following query containing a constant column in one of the UNION legs with a GROUP BY clause returns the wrong result.

```
set flags 'strategy,detail';

SELECT CTEGEN
FROM
  (SELECT CODSOC,CTEGEN,CDEV,SUM(SIGNE) AS SIGNE
   FROM (
```

Oracle® Rdb for OpenVMS

```
(SELECT CODSOC,CTEGEN,CDEV,SUM(MONTANT_SIGNE) AS SIGNE
FROM T1
WHERE
    CTEGEN NOT LIKE '7%' AND PERIODE = '2001'
GROUP BY CODSOC,CTEGEN,CDEV)
UNION
(SELECT CODSOC,'59100000',CDEV,SUM(MONTANT_SIGNE) AS SIGNE
FROM T1
WHERE
    CTEGEN LIKE '7%' AND PERIODE = '2001'
GROUP BY CODSOC,CTEGEN,CDEV)
) AS BBB
GROUP BY CODSOC,CTEGEN,CDEV) as
BBB(CODSOC,CTEGEN,CDEV,SIGNE)
WHERE
    SIGNE=0 AND
    CODSOC = 'MUT' AND CTEGEN LIKE '121%';
Tables:
    0 = T1
    1 = T1
Conjunct: <mapped field> = 0
Conjunct: <mapped field> LIKE '121%'
Merge of 1 entries
Merge block entry 1
Aggregate: 0:SUM (<mapped field>)
Sort: <mapped field>(a), <mapped field>(a), <mapped field>(a)
Conjunct: '59100000' LIKE '121%'          <== See Note
Merge of 1 entries
Merge block entry 1
Reduce: <mapped field>, <mapped field>, <mapped field>, <mapped field>
Sort: <mapped field>(a), <mapped field>(a), <mapped field>(a),
    <mapped field>(a)
Merge of 2 entries
Merge block entry 1
Aggregate: 1:SUM (0.MONTANT_SIGNE)
Conjunct: (0.CODSOC = 'MUT') AND (0.CTEGEN LIKE '121%')
Conjunct: NOT (0.CTEGEN LIKE '7%') AND (0.PERIODE = '2001')
Get      Retrieval by index of relation 0:T1
Index name T1_NDX [3:3]
Keys: (0.PERIODE = '2001') AND (<mapped field> = 'MUT') AND (
    <mapped field> LIKE '121%')
Bool: (0.CTEGEN LIKE '121%') AND (NOT (0.CTEGEN LIKE '7%'))
Merge block entry 2
Aggregate: 2:SUM (1.MONTANT_SIGNE)
Conjunct: 1.CODSOC = 'MUT'
Conjunct: '59100000' LIKE '121%'
Conjunct: (1.CTEGEN LIKE '7%') AND (1.PERIODE = '2001')
Get      Retrieval by index of relation 1:T1
Index name T1_NDX [3:3]
Keys: (1.PERIODE = '2001') AND (<mapped field> = 'MUT') AND (1.CTEGEN
    LIKE '7%')
Bool: ('59100000' LIKE '121%') AND (1.CTEGEN LIKE '7%')
0 rows selected
```

Note:: The conjunct containing the constants in both operands is incorrectly placed outside of the merge (UNION query).

This problem occurs when the query contains the following:

1. The main table is derived from a subquery of UNION legs.

2. Each UNION leg contains a GROUP BY clause.
3. The subquery SELECT statement contains an aggregate function SUM on a column which is NOT part of the GROUP BY columns.
4. One of the UNION legs contains a constant value as the SELECT column.
5. One of the WHERE predicates of the main query references the column that is mapped to the constant value column of the UNION leg.

As a workaround, the query works if the constant column is wrapped inside a CASE statement, as in the following example.

```

SELECT CTEGEN
FROM
  (SELECT CODSOC,CTEGEN,CDEV,SUM(SIGNE) AS SIGNE
    FROM (
      (SELECT CODSOC,CTEGEN,CDEV,SUM(MONTANT_SIGNE) AS SIGNE
        FROM T1
        WHERE
          CTEGEN NOT LIKE '7%' AND PERIODE = '2001'
        GROUP BY CODSOC,CTEGEN,CDEV)
      UNION
      (SELECT CODSOC,
!         '59100000',          ! <== wrap this constant in CASE statement
        CASE CTEGEN WHEN NULL THEN '*****' ELSE '59100000' END
        CDEV,SUM(MONTANT_SIGNE) AS SIGNE
        CDEV,
        SUM(MONTANT_SIGNE) AS SIGNE
        FROM T1
        WHERE
          CTEGEN LIKE '7%' AND PERIODE = '2001'
        GROUP BY CODSOC,CTEGEN,CDEV)
      ) AS BBB
    GROUP BY CODSOC,CTEGEN,CDEV) as
  BBB(CODSOC,CTEGEN,CDEV,SIGNE)
WHERE
  SIGNE=0 AND
  CODSOC = 'MUT' AND CTEGEN LIKE '121%';

```

Tables:

```

0 = T1
1 = T1

```

Conjunct: <mapped field> = 0

Merge of 1 entries

Merge block entry 1

Aggregate: 0:SUM (<mapped field>)

Sort: <mapped field>(a), <mapped field>(a), <mapped field>(a)

Merge of 1 entries

Merge block entry 1

Reduce: <mapped field>, <mapped field>, <mapped field>, <mapped field>

Sort: <mapped field>(a), <mapped field>(a), <mapped field>(a),
<mapped field>(a)

Conjunct: 0.CTEGEN LIKE '121%' <== See Note1

Merge of 2 entries

Merge block entry 1

Aggregate: 1:SUM (0.MONTANT_SIGNE)

Conjunct: (0.CODSOC = 'MUT') AND (0.CTEGEN LIKE '121%')

Conjunct: NOT (0.CTEGEN LIKE '7%') AND (0.PERIODE = '2001')

Get Retrieval by index of relation 0:T1

Index name T1_NDX [3:3]

Keys: (0.PERIODE = '2001') AND (<mapped field> = 'MUT') AND (
<mapped field> LIKE '121%')

Oracle® Rdb for OpenVMS

```
      Bool: (0.CTEGEN LIKE '121%') AND (NOT (0.CTEGEN LIKE '7%'))
Merge block entry 2
Aggregate: 2:SUM (1.MONTANT_SIGNE)
Conjunct: 1.CODSOC = 'MUT'
Conjunct: (1.CTEGEN LIKE '7%') AND (1.PERIODE = '2001')
Get      Retrieval by index of relation 1:T1
      Index name T1_NDX [3:3]
      Keys: (1.PERIODE = '2001') AND (<mapped field> = 'MUT') AND (1.CTEGEN
            LIKE '7%')
      Bool: 1.CTEGEN LIKE '7%'
CTEGEN
12111090
1 row selected
```

Note1:: The correct conjunct is now generated instead of the wrong one.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.7 Left Outer Join Query with Sub-select Overflows the Stack

Bugs 3357593 and 2649215

The following left outer join query with a sub-select overflows the stack.

```
SELECT
  (select ct.name from
    (select firm_id, cust_id, addr_id,
      (select name from t1 where
        t2.addr_id = t1.addr_id) as name
      FROM t2) as ct
    where ct.firm_id = t5.firm_id and
          ct.cust_id = t5.cust_id) as cust_name
FROM
  stock t3
  left outer join t4
    on t4.firm_id = t3.firm_id and
       t4.trade_id = t3.trade_id and
       t4.trade_pos = t3.trade_pos
  left outer join t5
    on t5.firm_id = t4.firm_id and
       t5.trade_id = t4.trade_id;
%RDB-F-IMP_EXC, facility-specific limit exceeded
-RDMS-F-XPR_STACK_OFLO, expression forces too many levels of recursion
```

The problem occurs when the query selects from three tables (t3, t4 and t5) of left outer join and contains a sub-select clause with two tables (t1 and t2) joined by the equality predicates referencing the columns from table t5.

This problem is caused by the fix made for Bug 2649215 which did not cover this particular query.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.8 Signal Failing in Compound Statement

Bug 3364468

A signal in a compound statement could cause SQL/RDB to go into an infinite loop, especially if a label with a "leave label" statement was encountered.

The following example shows this behavior.

```
begin
declare :cnt integer = 0;
declare :tdbkey char(8);

while (exists (select c0 from update_db.module.dp$_2)) loop
    rollback;

    set transaction read write isolation level read committed
    reserving update_db.job_history for shared write;

    label_loop:

    for :dp as each row of cursor dp$_2 for
        select c0 from update_db.module.dp$_2
    do
        set :tdbkey = :dp.c0;

        delete from update_db.module.dp$_2 where current of dp$_2;

        update update_db.job_history set
            supervisor_id = supervisor_id where dbkey = :tdbkey;

        set :cnt = :cnt + 1;
        if (:cnt >= 30) then leave label_loop;
        end if;
    end for;
    commit;
    if (:cnt >= 1) then
        signal 'JDLAY';
    end if;
end loop;
end;
```

One possible workaround would be to avoid using a "leave label" statement.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.9 Bugcheck in KOD\$ROLLBACK and PSII2REMOVEDUPBBC with Sorted Ranked Indexes

Bug 3408974

In Oracle Rdb Release 7.0.7.1, a problem was introduced in the handling of sorted ranked indexes by the optimizer.

Depending on the structure and data distribution of index nodes in the sorted ranked index, it was possible that one of the following exceptions would be raised.

```
COSI-F-BUGCHECK, internal consistency failure  
Exception occurred at KOD$ROLLBACK + 00000328
```

```
COSI-F-BUGCHECK, internal consistency failure  
Exception occurred at PSII2REMOVEDUPBBC + 0000142C
```

These problems only occur when a sorted ranked index is being used by the optimizer to filter out possible candidate dbkeys for selection.

A workaround for this problem is to use normal sorted indexes.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.10 Database Corruption When LOCKING IS PAGE LEVEL Enabled

Storage areas created using the "LOCKING IS PAGE LEVEL" option would sometimes have missing updates. This problem would occur when there was a lot of contention for pages in the storage area.

For example, if an index update were lost, IDXDATMIS errors may have been displayed by RMU/VERIFY.

```
%RMU-W-IDXDATMIS, Index ALL_DEPT does not point to a row in table  
DEPARTMENT_REC.  
Logical dbkey of the missing row is 149:302:5.
```

This problem can be avoided by using the "LOCKING IS ROW LEVEL" option, which is the default.

Note that the "LOCKING IS PAGE LEVEL" option works best with applications that have little contention between users for database pages. Also, ideally, transactions should be short in duration and most if not all pages used by the transaction should be able to simultaneously reside in the buffers allocated to the process.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.11 Table Constraint Should Fail on Updating a Row

Bug 3434648

A customer updates a row in the table, changing the value in one column from NULL to some text string. Under Oracle Rdb Release 7.0.6.3, a constraint on that table correctly reports that the update is in error. However in Oracle Rdb Release 7.0.7.1, the constraint failure does not occur.

A simple reproducer using a SELECT statement is created from the UPDATE/COMMIT statement of the original BUG report where the CHECK clause defined in the table constraint contains four OR predicates, such as the following.

```
predicate-1 OR predicate-2 OR predicate-3 OR predicate-4
```

which is inverted using a NOT operator, such as:

Oracle® Rdb for OpenVMS

NOT (predicate-1 OR predicate-2 OR predicate-3 OR predicate-4)

The query selects "Constraint violation" on the ALIAS_LOC table where the dbkey is the target row that UPDATE performs on. It should find the row that gets the constraint violation when the conditions in the predicates are not met.

```
set flags 'strategy,detail';
```

```
Simple reproducer for the problem
```

```
-----  
  
declare :dbk bigint;  
select dbkey into :dbk from ALIAS_LOC A limit to 1 row;  
  
select 'Constraint Failure' from ALIAS_LOC A  
  where a.dbkey = :dbk and  
        NOT (  
          ! Part 1  
          (A.LOCATION_1_NAME is null and  
           A.LOCATION_2_NAME is null and  
           A.LOCATION_3_NAME is null)  
        OR  
          ! Part 2  
          ((A.LOCATION_2_NAME is null and           ! <== missing in Cross block 3  
           A.LOCATION_3_NAME is null)              ! <== of the strategy  
           and exists  
           (select * from LOCATION C10  
            where  
              C10.ORGANIZATION_NAME = A.ORGANIZATION_NAME and  
              C10.UNIT_NAME         = A.UNIT_NAME             and  
              C10.PARENT_LOCATION_NAME is null                and  
              C10.LOCATION_NAME     = A.LOCATION_1_NAME  
            )           ! end of select  
          )           ! end of or  
        OR  
          ! Part 3  
          ((A.LOCATION_3_NAME is null)  
           and exists  
           (select * from LOCATION C11, LOCATION C12  
            where  
              C11.ORGANIZATION_NAME = A.ORGANIZATION_NAME and  
              C11.UNIT_NAME         = A.UNIT_NAME             and  
              C11.PARENT_LOCATION_NAME is null                and  
              C11.LOCATION_NAME     = A.LOCATION_1_NAME     and  
              C12.PARENT_LOCATION_NAME = C11.LOCATION_NAME and  
              C12.LOCATION_NAME     = A.LOCATION_2_NAME  
            )           ! end of select  
          )           ! end of or  
        OR  
          ! Part 4  
          (exists  
           (select * from LOCATION C13, LOCATION C14, LOCATION C15  
            where  
              C13.ORGANIZATION_NAME = A.ORGANIZATION_NAME and  
              C13.UNIT_NAME         = A.UNIT_NAME             and  
              C13.PARENT_LOCATION_NAME is null                and  
              C13.LOCATION_NAME     = A.LOCATION_1_NAME     and  
              C14.PARENT_LOCATION_NAME = C13.LOCATION_NAME and  
              C14.LOCATION_NAME     = A.LOCATION_2_NAME     and  
              C15.PARENT_LOCATION_NAME = C14.LOCATION_NAME and
```


Oracle® Rdb for OpenVMS

```

                C15.LOCATION_NAME          = A.LOCATION_3_NAME
            )
            )
        )
    limit to 1 row;
Tables:
0 = ALIAS_LOC
1 = LOCATION
2 = LOCATION
3 = LOCATION
4 = LOCATION
5 = LOCATION
6 = LOCATION
Firstn: 1
Cross block of 4 entries
Cross block entry 1
  Conjunct: NOT MISSING (0.LOC_NAME_1) OR NOT MISSING (0.LOC_NAME_2) OR NOT
            MISSING (0.LOC_NAME_3)
  Conjunct: 0.DBKEY = <var0>
  Firstn: 1
  Get      Retrieval by DBK of relation 0:ALIAS_LOC
Cross block entry 2
  Conjunct: NOT MISSING (0.LOC_NAME_3) OR (<agg0> = 0)
  Aggregate-F1: 0:COUNT-ANY (<subselect>)
Cross block of 2 entries
  Cross block entry 1
    Index only retrieval of relation 2:LOCATION
    Index name  LOCATION_NDX [4:4]          Direct lookup
    Keys: (2.ORG_NAME = 0.ORG_NAME) AND (2.UNIT_NAME = 0.UNIT_NAME) AND
          (MISSING (2.P_LOC_NAME)) AND (2.LOC_NAME = 0.LOC_NAME_1)
  Cross block entry 2
    Conjunct: (3.P_LOC_NAME = 2.LOC_NAME) AND (3.LOC_NAME = 0.LOC_NAME_2)
    Index only retrieval of relation 3:LOCATION
    Index name  LOCATION_NDX [0:0]
Cross block entry 3
  Conjunct: <agg1> = 0          ! <=== See Note below.
  Aggregate-F1: 1:COUNT-ANY (<subselect>)
  Index only retrieval of relation 1:LOCATION
  Index name  LOCATION_NDX [4:4]          Direct lookup
  Keys: (1.ORG_NAME = 0.ORG_NAME) AND (1.UNIT_NAME = 0.UNIT_NAME) AND (
        MISSING (1.P_LOC_NAME)) AND (1.LOC_NAME = 0.LOC_NAME_1)
Cross block entry 4
  Conjunct: <agg2> = 0
  Aggregate-F1: 2:COUNT-ANY (<subselect>)
Cross block of 3 entries
  Cross block entry 1
    Index only retrieval of relation 4:LOCATION
    Index name  LOCATION_NDX [4:4]          Direct lookup
    Keys: (4.ORG_NAME = 0.ORG_NAME) AND (4.UNIT_NAME = 0.UNIT_NAME) AND
          (MISSING (4.P_LOC_NAME)) AND (4.LOC_NAME = 0.LOC_NAME_1)
  Cross block entry 2
    Conjunct: (5.P_LOC_NAME = 4.LOC_NAME) AND (5.LOC_NAME = 0.LOC_NAME_2)
    Index only retrieval of relation 5:LOCATION
    Index name  LOCATION_NDX [0:0]
  Cross block entry 3
    Conjunct: (6.P_LOC_NAME = 5.LOC_NAME) AND (6.LOC_NAME = 0.LOC_NAME_3)
    Index only retrieval of relation 6:LOCATION
    Index name  LOCATION_NDX [0:0]
0 rows selected

```

NOTE:: Note that the following conjunct under the Cross block entry 3 should contain the following predicates, but that is missing.

```
Conjunct: NOT MISSING (0.LOC_NAME_2) OR NOT MISSING (0.LOC_NAME_3) OR
          (<aggl> = 0)
```

This problem was caused by the fix for Bug 2285818 where the query contains predicates shared by other parts of the OR expression tree.

This query also contains the following similar predicates shared by other parts of the OR expression tree.

```
NOT (
  ! Part 1
  (A.LOCATION_1_NAME is null and
   A.LOCATION_2_NAME is null and
   A.LOCATION_3_NAME is null)
OR
  ! Part 2
  ((A.LOCATION_2_NAME is null and           ! <== shared in PART 1
   A.LOCATION_3_NAME is null)             ! <== shared in PART 1
   and
   ...etc...)
)
! end of or
OR
  ! Part 3
  ((A.LOCATION_3_NAME is null) and         ! <== shared in PART 1 and 2
   ...etc...)
OR
  ! Part 4
  (...etc...)
);
```

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.1.12 Zigzag Match Query with Descending Index Segment Returns Wrong Results

Bug 3514413

The following zigzag match query with a descending index segment should select 6 rows but instead just one row is selected.

```
set flags 'strategy,detail';
SELECT  MBR_ID, ACCNT_TYPE, ACCNT_NO, CNTR_PRCE, TRB.SERI_CODE
FROM    TRD_BOOK_TB TRB, SERIES_DAILY_TB SED, SERIES_TB SER
WHERE   TRB.DRV_MKT_ID = 'BF'
        AND TRB.UJLY_ID = '61'
        AND SED.YYMMDD   = '20040309'
        AND TRB.SERI_CODE = SED.SERI_CODE
        AND SED.SERI_CODE = SER.SERI_CODE
        AND SER.ACT_DATE  <= '20040309'
        AND SER.DEL_DATE  >= '20040309'
;
Tables:
```

Oracle® Rdb for OpenVMS

```

0 = TRD_BOOK_TB
1 = SERIES_DAILY_TB
2 = SERIES_TB
Conjunct: 1.SERI_CODE = 2.SERI_CODE
Match
Outer loop
  Conjunct: 0.SERI_CODE = 1.SERI_CODE
  Match
    Outer loop
      Sort: 0.SERI_CODE(a)
      Conjunct: (0.DRV_MKT_ID = 'BF') AND (0.ULY_ID = '61')
      Get      Retrieval sequentially of relation 0:TRD_BOOK_TB
    Inner loop      (zig-zag)
      Index only retrieval of relation 1:SERIES_DAILY_TB
      Index name  SERIES_DAILY_IDX [1:1]
      Keys: 1.YYMMDD = '20040309'
    Inner loop
      Temporary relation
      Sort: 2.SERI_CODE(a)
      Conjunct: 2.ACT_DATE <= '20040309'
      Leaf#01 BgrOnly 2:SERIES_TB Card=1002
      Bool: 2.DEL_DATE >= '20040309'
      BgrNdx1 SERIES_IDX [0:1] Fan=11
      Keys: 2.ACT_DATE <= '20040309'
      Bool: 2.DEL_DATE >= '20040309'
TRB.MBR_ID  TRB.ACCNT_TYPE  TRB.ACCNT_NO  TRB.CNTR_PRCE  TRB.SERI_CODE
016          3              001000027      106.70        KR4161460000
1 row selected

```

As a workaround, the query works if the zigzag match strategy is disabled, as in the following example.

```

SQL> set flags 'nozigzag_match';
OR
$DEFINE RDMS$$DISABLE_ZIGZAG_MATCH 2

```

Here is the new strategy of the new run:

Tables:

```

0 = TRD_BOOK_TB
1 = SERIES_DAILY_TB
2 = SERIES_TB
Conjunct: 1.SERI_CODE = 2.SERI_CODE
Match
Outer loop
  Conjunct: 0.SERI_CODE = 1.SERI_CODE
  Match
    Outer loop
      Sort: 0.SERI_CODE(a)
      Conjunct: (0.DRV_MKT_ID = 'BF') AND (0.ULY_ID = '61')
      Get      Retrieval sequentially of relation 0:TRD_BOOK_TB
    Inner loop
      Index only retrieval of relation 1:SERIES_DAILY_TB
      Index name  SERIES_DAILY_IDX [1:1]
      Keys: 1.YYMMDD = '20040309'
    Inner loop
      Temporary relation
      Sort: 2.SERI_CODE(a)
      Conjunct: 2.ACT_DATE <= '20040309'
      Leaf#01 BgrOnly 2:SERIES_TB Card=1002
      Bool: 2.DEL_DATE >= '20040309'

```

Oracle® Rdb for OpenVMS

```

BgrNdx1 SERIES_IDX [0:1] Fan=11
  Keys: 2.ACT_DATE <= '20040309'
  Bool: 2.DEL_DATE >= '20040309'
TRB.MBR_ID  TRB.ACCNT_TYPE  TRB.ACCNT_NO  TRB.CNTR_PRCE  TRB.SERI_CODE
029          3              001000025     103.90        KR4161430000
017          0              001000004     103.90        KR4161430000
029          1              001000023     106.45        KR4161460000
016          3              001000027     106.45        KR4161460000
029          2              001000099     106.70        KR4161460000
016          3              001000027     106.70        KR4161460000
6 rows selected

```

The query also works with the zigzag match strategy if the index column YMMDD of SERIES_DAILY_IDX index is re-defined as ASCENDING instead of DESCENDING.

```

drop index SERIES_DAILY_IDX;
create unique index SERIES_DAILY_IDX
  on SERIES_DAILY_TB (
!  YMMDD desc,
  YMMDD asc,
  SERI_CODE asc);

```

Tables:

```

0 = TRD_BOOK_TB
1 = SERIES_DAILY_TB
2 = SERIES_TB

```

Conjunct: 1.SERI_CODE = 2.SERI_CODE

Match

Outer loop

Conjunct: 0.SERI_CODE = 1.SERI_CODE

Match

Outer loop

Sort: 0.SERI_CODE(a)

Conjunct: (0.DRV_MKT_ID = 'BF') AND (0.ULY_ID = '61')

Get Retrieval sequentially of relation 0:TRD_BOOK_TB

Inner loop (zig-zag)

Index only retrieval of relation 1:SERIES_DAILY_TB

Index name SERIES_DAILY_IDX [1:1]

Keys: 1.YMMDD = '20040309'

Inner loop

Temporary relation

Sort: 2.SERI_CODE(a)

Conjunct: 2.ACT_DATE <= '20040309'

Leaf#01 BgrOnly 2:SERIES_TB Card=1002

Bool: 2.DEL_DATE >= '20040309'

BgrNdx1 SERIES_IDX [0:1] Fan=11

Keys: 2.ACT_DATE <= '20040309'

Bool: 2.DEL_DATE >= '20040309'

```

TRB.MBR_ID  TRB.ACCNT_TYPE  TRB.ACCNT_NO  TRB.CNTR_PRCE  TRB.SERI_CODE
029          3              001000025     103.90        KR4161430000
017          0              001000004     103.90        KR4161430000
029          1              001000023     106.45        KR4161460000
016          3              001000027     106.45        KR4161460000
029          2              001000099     106.70        KR4161460000
016          3              001000027     106.70        KR4161460000
6 rows selected

```

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.2 SQL Errors Fixed

2.2.1 COMMIT or ROLLBACK Bugchecks in Conditional Expression

Bug 2350880

In prior versions of Oracle Rdb, attempts to COMMIT or ROLLBACK a transaction within the scope of a conditional expression (IF or WHILE) would cause a bugcheck.

- Alpha OpenVMS 7.3-1
- Oracle Rdb Server V7.0-70
- COSI-F-BUGCHECK, internal consistency failure
- Exception occurred at KOD\$PREPARE + 00000228
- Called from KOD\$COMMIT + 0000018C
- Called from RDMS\$\$INT_COMMIT_TRANSACTION + 000002BC

The following shows one example query and the resulting bugcheck dump.

```
SQL> attach 'filename sql$database';
SQL>
SQL> set flags 'trace';
SQL>
SQL> begin
cont> set transaction read only;
cont>
cont> if exists
cont>     (select *
cont>       from rdb$relation_fields t1, rdb$relations t2
cont>       where t2.rdb$relation_name = t1.rdb$relation_name)
cont> then
cont>     trace 'rows exist';
cont>     commit;
cont>     set transaction read only;
cont> end if;
cont>
cont> commit;
cont> end;
~Xt: rows exist
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DISK1:[TEST]RDSBUGCHK.DMP;
```

This problem has been corrected in Oracle Rdb Release 7.0.7.2. Rdb now closes the index scan from the IF and WHILE condition prior to executing the body of the conditional expression.

2.2.2 Unexpected Failure When Using ALTER ... THRESHOLDS Clause

Bug 3283407

In prior releases of Oracle Rdb, the THRESHOLDS clause for ALTER INDEX or ALTER STORAGE MAP was rejected if the logical area already existed.

The following example shows the reported error.

```
SQL> alter index ul_partitioned_table store using (attr1, attr2)
cont> in sa1 ( thresholds are (90,90,90) ) with limit of (10,15)
cont> in sa2 ( thresholds are (90,90,90) ) with limit of (10,25)
cont> in sa3 ( thresholds are (90,90,90) ) with limit of (10,35);
SQL> commit;
SQL> alter index ul_partitioned_table store using (attr1, attr2)
cont> in sa1 ( thresholds are (90,90,90) ) with limit of (10,15)
cont> in sa2 ( thresholds are (90,90,90) ) with limit of (10,25)
cont> in sa3 ( thresholds are (90,90,90) ) with limit of (10,35);
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-THRESHAREEXI, illegal thresholds usage - area SA1 exists,
and cannot have THRESHOLDS respecified
```

The problem was that the new threshold values were the same as those in that index or storage map partition and could have been ignored by Rdb.

This problem has been corrected in Oracle Rdb Release 7.0.7.2. Rdb now ignores the THRESHOLDS clause if it respecifies the existing thresholds for the index or storage map partition.

2.2.3 Unexpected Bugcheck when Oracle-style Outer Join Used with Subselect

Bug 3329186

When the Oracle style outer join modifier (+) was used in a WHERE clause that also contained a subselect with an aggregate function (SUM, AVG, MIN, MAX, or COUNT), a bugcheck would be generated by SQL.

The following shows a simple example of the type of query.

```
SQL> select *
cont> from employees e, salary_history sh
cont> where e.employee_id (+) = sh.employee_id
cont> and sh.salary_amount = (select max (salary_amount)
cont> from salary_history);
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DISK1:[DATABASE]SQLBUGCHK.DMP;
%SQL-F-BUGCHK, There has been a fatal error. Please contact your Oracle
support representative. SQL$SEMRSE - 71
```

The only workaround is to recode the query using ANSI/ISO Standard SQL syntax using the LEFT OUTER JOIN or RIGHT OUTER JOIN operators.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.2.4 %SQL-F-NODBFIL, Alias Missing a Declaration With SQLMOD

Bug 1258536

This problem can arise when a SQL\$MOD module is compiled with /CONNECT and has a database alias declared with a run-time resolution. In this case, if another SQL\$MOD module or a Dynamic SQL module

Oracle® Rdb for OpenVMS

declares the same alias but with a compile-time resolution, the first call to any procedure in the module with the run-time resolution will fail with:

```
SQL-F-NODBFILE, ALIAS <alias-name> IS MISSING A DECLARATION  
where <alias-name> will be the alias that fails.
```

For example, suppose SQL\$MOD module1 declares an alias with run-time resolution like this:

```
declare base alias for compiletime filename 'mf_personnel'  
        runtime filename db_name  
Note: procedures in module1 will have "db_name" as a parameter
```

Suppose also that SQL\$MOD module2 declares the same alias but with a compile-time resolution like this:

```
declare alias base filename 'mf_personnel'
```

Now both modules are linked with a main program that calls a procedure from module1. The call to the module1 procedure will result in:

```
SQL-F-NODBFILE, ALIAS BASE IS MISSING A DECLARATION
```

There is no known workaround for this problem. There is no way to achieve run-time resolution of the alias.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.3 Oracle RMU Errors Fixed

2.3.1 RMU/CHECKPOINT Slow When Number of Nodes is One

Bug 3240378

If a database was set to have NUMBER OF CLUSTER NODES 1, and the AIJ log server (ALS) was not enabled, and there was no update activity occurring at that point in time, the RMU/CHECKPOINT command could take many minutes to complete. For example:

```
$ SQL$
CREATE DATABASE FILENAME CKPT_TEST
  NUMBER OF CLUSTER NODES 1
  CREATE STORAGE AREA RDB$SYSTEM FILENAME CKPT_TEST;
DISCONNECT ALL;

ALTER DATABASE FILE CKPT_TEST
  JOURNAL IS ENABLED
  (FAST COMMIT IS ENABLED,
   -- This test only fails if there is no ALS.
   LOG SERVER IS MANUAL)
  ADD JOURNAL AIJ_1 FILENAME 'SYS$DISK:[ ]CKPT_TEST.AIJ';

EXIT;
$ CREATE SUB1.COM
$ SQL$
ATTACH 'FILENAME CKPT_TEST';
INSERT INTO T1 VALUES (1);
COMMIT;
-- Wait till we are killed.
$WAIT 1:0:0
$
$ SPAWN /NOWAIT/OUT=SUB1.LOG/PROCESS=CKPT_TIMELY_SUB @SUB1.COM
%DCL-S-SPAWNED, process CKPT_TIMELY_SUB spawned
$ WAIT 0:1:0
$ START_TIME = F$CVTIME (,,"MINUTEOFYEAR")
$ RMU/CHECKPOINT CKPT_TEST
$ END_TIME = F$CVTIME (,,"MINUTEOFYEAR")
$ CKPT_DURATION = END_TIME - START_TIME
$ IF CKPT_DURATION .GT. 2
$ THEN WRITE SYS$OUTPUT -
  "'CKPT_DURATION' minutes is too long for an RMU/CHECKPOINT command"
5 minutes is too long for an RMU/CHECKPOINT command
$ ENDIF
$ STOP CKPT_TIMELY_SUB
```

This problem can be avoided by utilizing the AIJ log server process or by setting the database number of cluster nodes greater than one.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.3.2 RMU/LOAD/DEFER_INDEX_UPDATES ACCVIO with a Hashed Partitioned Index

Bug 3262298

If an RMU/LOAD/DEFER_INDEX_UPDATES command was used and a partitioned hashed index was the only or primary index defined for the table being loaded, an access violation would occur and an RMUBUGCHK.DMP and RDSBUGCHK.DMP would be generated. The exception in the RDSBUGCHK.DMP would occur in the routine PSII\$INSERT_T. This was because the test for a hashed index did not work correctly so an attempt was made to update a hashed index using a sorted index routine. This problem has been fixed.

The following example shows that an access violation occurred if the /DEFER_INDEX_UPDATES qualifier was used for an RMU/LOAD of a table for which the only or the primary index was a partitioned hashed index.

```
SQL> att 'fi [.to]mf_personnel';
SQL> show table job_history
Information for table JOB_HISTORY

Columns for table JOB_HISTORY:
Column Name          Data Type          Domain
-----
EMPLOYEE_ID          CHAR(5)            ID_NUMBER
JOB_CODE              CHAR(4)            JOB_CODE
JOB_START             DATE VMS           STANDARD_DATE
JOB_END               DATE VMS           STANDARD_DATE
DEPARTMENT_CODE      CHAR(4)            DEPARTMENT_CODE
SUPERVISOR_ID        CHAR(5)            ID_NUMBER

Table constraints for JOB_HISTORY:
No constraints found

Constraints referencing table JOB_HISTORY:
No constraints found

Indexes on table JOB_HISTORY:
JOB_HISTORY_HASH          with column EMPLOYEE_ID
  Duplicates are allowed
  Type is Hashed Scattered
  Compression is DISABLED
Store clause:            STORE
                          using (EMPLOYEE_ID)
                          in EMPIDS_LOW
                          with limit of ('00200')
                          in EMPIDS_MID
                          with limit of ('00400')
                          otherwise in EMPIDS_OVER

Comment:                hash index for job_history

Storage Map for table JOB_HISTORY:
  JOB_HISTORY_MAP

Triggers on table JOB_HISTORY:
No triggers found

SQL> exit
```

Oracle® Rdb for OpenVMS

```
$rmu/load/defer/log [.to]mf_personnel job_history job_history.unl
%RDMS-I-BUGCHKDMP, generating bugcheck dump file
  DEVICE:[LOGIN]RDSBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
  virtual address=00000050, PC=00186C66, PSL=03C00000
%RMU-I-BUGCHKDMP, generating bugcheck dump file
  DEVICE:[LOGIN]RMUBUGCHK.DMP;
%RMU-I-DATRECREAD, 274 data records read from input file.
%RMU-I-DATRECSTO, 0 data records stored.
%RMU-F-FTL_LOAD, Fatal error for LOAD operation at 13-JAN-2004 13:32:24.76
$type [login]RDSBUGCHK.DMP

***** Exception at 0093EBB0 : PSII$INSERT_T + 000000B7
```

The workaround for this problem is to not use the qualifier /DEFER_INDEX_UPDATES when doing the RMU/LOAD of a table for which the only or primary index is a partitioned hashed index.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.3.3 RMU/RECOVER/ONLINE Without /JUST_CORRUPT or /AREA Qualifiers Bugchecks

As described in the documentation, when issuing an RMU/RECOVER command, the /ONLINE qualifier can only be use in conjunction with the /JUST_PAGE or /AREA qualifier.

In previous releases, if the /ONLINE qualifier was used alone, it resulted in an RMU bugcheck as shown below.

```
$ rmu/recover/log/online AIJ_BCK.AIJ
%RMU-E-RECF FAILED, fatal, unexpected roll-forward error detected at AIJ record 1
%COSI-F-BUGCHECK, internal consistency failure
%RMU-F-FATALOSI, Fatal error from the Operating System Interface.
%RMU-I-BUGCHKDMP, generating bugcheck dump file disk:[dir]RMUBUGCHK.DMP;
%RMU-F-FTL_RCV, Fatal error for RECOVER operation at <timestamp>
```

The exception and the call frame in the dump are as follows.

```
***** Exception at 007212B8 : KUTREC$ABORT + 000005D8
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 0071F5E4 : KUTREC$SETUP + 00000214
Saved PC = 00474198 : RMUREC$RECOVER_ACTION + 00000948
Saved PC = 004737D8 : RMUCLI$RECOVER + 000004C8
Saved PC = 003A56F4 : RMU_DISPATCH + 00000D44
Saved PC = 003A4508 : RMU_STARTUP + 000004A8
Saved PC = 001E002C : RMU$MAIN + 0000002C
```

Now RMU will detect missing /JUST_PAGE or /AREA qualifiers when /ONLINE is used and will return an error message, as in the following example.

```
$ rmu/recover/log/online AIJ_BCK.AIJ
%RMU-F-CONFLSWIT, conflicting qualifiers /ONLINE and /AREAS or /JUST_CORRUPT
missing
%RMU-F-FTL_RCV, Fatal error for RECOVER operation at <timestamp>
```

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.3.4 Invalid RMU-E-INASPAREA Message in RMU/VERIFY/ROOT

Bug 3424732

While doing an RMU/VERIFY of the database root, each live storage area entry in the database root was correctly reloaded to get its current state on all cluster nodes accessing the database but the corresponding snapshot area entry was not reloaded by RMU/VERIFY to bring it up to date. If a new database storage area had just been created on one node in the cluster, this could lead to the following incorrect RMU/VERIFY error message being output on another cluster node which was doing an RMU/VERIFY of the database root.

```
RMU-E-INASPAREA, Live area AREA_NAME points to a snapshot area that is
inactive.
```

Repeating the RMU/VERIFY would not show this error since this would reload the snapshot entry and bring it up to date. This problem has been fixed.

The following example shows the sequence of adding a storage area entry to a database on one node and then verifying the database on another node that gave the RMU-E-INASPAREA error.

On one cluster node...

```
$ @RDM$DEMO:PERSONNEL SQL M NOCDD
$ SQL
SQL> ALTER DATABASE FILENAME MF_PERSONNEL OPEN IS MANUAL;
SQL> ALTER DATABASE FILENAME MF_PERSONNEL RESERVE 1 STORAGE AREA;
SQL> EXIT;
$ RMU/OPEN MF_PERSONNEL
```

On another node in the cluster...

```
$ RMU/OPEN MF_PERSONNEL
```

Back to first cluster node...

```
$ SQL
SQL> ALTER DATABASE FILENAME MF_PERSONNEL ADD STORAGE AREA TEST;
SQL> EXIT;
```

Back to the second cluster node...

```
$ RMU/VERIFY MF_PERSONNEL
%RMU-E-INASPAREA, Live area TEST points to a snapshot area that is inactive.
```

The workaround for this problem is do the RMU/VERIFY on the same cluster node where the storage area was created or to repeat the RMU/VERIFY on the second cluster node if the RMU-E-INASPAREA error is returned on the first RMU/VERIFY.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.3.5 RMU Support For /DENSITY = SDLT320

Oracle Rdb RMU commands that support the /DENSITY qualifier (ie, RMU/BACKUP, RMU/BACKUP/AFTER_JOURNAL and RMU/OPTIMIZE_AIJ) now support the keyword "SDLT320" for use with SuperDLT320 tape drives.

2.3.6 Incorrect Verification of Invalid Reference DBKEYS in Ranked Indices

Bugs 3009262 and 3537202

If an index of *TYPE IS SORTED RANKED* had invalid reference dbkeys in duplicate overflow nodes, RMU/VERIFY/INDEX could fail to report the error. If the error was reported, it was reported as an informational message.

This type of corruption was possibly introduced by the problem referenced in [Section 2.1.5](#).

In the following example, the index corruption is correctly reported.

```
$ RMU/VERIFY/INDEX=MY_IDX MY_DB
%RMU-E-BADREFDBK, Invalid reference pointer 358:263034:0 for duplicate B-tree
node 299:88538:0
%RMU-I-DUPRECDBK, the last duplicate record dbkey was 358:294862:1
%RMU-I-DUPOWNDBK, Dbkey of owner of this duplicate node is 299:89267:1
%RMU-I-BTRERPATH, parent B-tree node of 299:89267:1 is at 299:89245:0
%RMU-I-BTRERPATH, parent B-tree node of 299:89245:0 is at 299:89138:1
%RMU-I-BTRERPATH, parent B-tree node of 299:89138:1 is at 299:88639:1
%RMU-I-BTRERPATH, parent B-tree node of 299:88639:1 is at 299:88437:1
%RMU-I-BTRROODBK, root dbkey of B-tree is 299:88437:1
```

This problem represents a real corruption in the index and can lead to bugchecks as described in [Section 2.1.5](#).

As a workaround for this problem, the offending index must be dropped and recreated.

This problem has been corrected in Oracle Rdb Release 7.0.7.2.

2.4 LogMiner Errors Fixed

2.4.1 RMU /UNLOAD /AFTER_JOURNAL Incorrect NULL Bit Setting When VARCHAR is Last Column

Bug 3309002

In prior versions of Oracle Rdb, the RMU /UNLOAD /AFTER_JOURNAL command could incorrectly process the null bit vector for tables with the last column being a VARCHAR data type. The LogMiner was not correctly calculating the position of the null bit vector within the data record and could pick up stray bit patterns as the null bit vector content. The effect of not using the correct null bit vector content could be NULL column values being incorrectly returned as not NULL or not NULL column values being incorrectly returned as NULL.

The following example script demonstrates one possible effect of this problem. The column Z\$I6 should be returned as NULL but is being extracted as a zero value:

```
$ SQL$
CREATE DATA FILE FOO;
CREATE TABLE T1 (
  Z$I1 INT, Z$I2 INT, Z$I3 INT, Z$I4 INT,
  Z$I5 INT, Z$I6 INT, X$I7 INT, Z$I8 INT,
  Z$I9 INT, Z$I10 INT, Z$V1 VARCHAR(10));
COMMIT;
DISCONNECT ALL;
ALTER DATA FILE FOO ADD JOURNAL J1 FILE J1 JOURNAL ENABLE;
EXIT;
$ RMU/SET LOGMINER/ENABLE FOO.RDB
$ RMU/BACKUP/NOLOG FOO.RDB NLA0:FOO
$ RMU/BACKUP/AFTER/NOLOG FOO.RDB NLA0:FOO
$ SQL$
ATTACH 'FILE FOO';
INSERT INTO T1 VALUES (1,2,3,4,5,NULL,7,8,9,10,'TEST');
1 ROW INSERTED
COMMIT;
EXIT;
$ RMU/BACKUP/AFTER/NOLOG FOO.RDB B1.AIJ
$ RMU/UNLOAD/AFTER_JOURNAL/NOLOG FOO.RDB B1.AIJ -
  /TABLE=(NAME=T1,OUTPUT=T1.TXT) /FORMAT=DUMP
$ SEARCH T1.TXT Z$
Z$I1          : 1
Z$I2          : 2
Z$I3          : 3
Z$I4          : 4
Z$I5          : 5
Z$I6          : 0
Z$I8          : 8
Z$I9          : 9
Z$I10         : 10
Z$V1         : (4) TEST
$ EXIT
```

This problem has been corrected in Oracle Rdb Release 7.0.7.2. The RMU /UNLOAD /AFTER_JOURNAL command now correctly determines the location of the null bit vector within the record when the final column

of the record is a VARCHAR.

2.5 Oracle Trace Errors Fixed

2.5.1 Oracle TRACE COLLECT FORMAT Command Could Not Create a Database After Rdb Upgrade

There was a problem with Oracle Trace Release 2.4.1 where, after upgrading to Rdb Release 7.0.7.1, the COLLECT FORMAT command was no longer able to successfully create the database where the formatted data is stored. This problem has been fixed in Oracle TRACE Release 2.4.1.1 Update 01. Note that this problem has been fixed in Oracle TRACE, not Oracle Rdb.

The following example shows the problem that the Oracle TRACE COLLECT FORMAT command had creating a database when Rdb was upgraded to Release 7.0.7.1.

```
$ COLLECT FORMAT SAMPLE_DATA.DAT SAMPLE_DATA.RDB
%EPC-E-FMT_FAILURE, Formatting failed
-RDMS-F-NOEUACCESS, unable to acquire exclusive access to database
%EPC-E-OPFAIL, Operation failed
```

This problem has been corrected in Oracle TRACE Release 2.4.1.1 Update 01.

Chapter 3

Software Errors Fixed in Oracle Rdb Release 7.0.7.1

This chapter describes software errors that are fixed by Oracle Rdb Release 7.0.7.1.

3.1 Software Errors Fixed That Apply to All Interfaces

3.1.1 Recovery of Empty Optimized AIJ Does Not Update the Sequence Number

Bug 2581948

The recovery of an empty optimized AIJ does not update the next AIJ sequence number. This will prevent the recovery of the next AIJ as shown below. An optimized AIJ file can be empty if the corresponding AIJ file contains only rolled-back transactions. See the following example.

```
$ sql$
SQL> attach 'filename opt_aij';
SQL> insert into t1 values (1);
1 row inserted
SQL> rollback;
SQL> exit
$ RMU /BACKUP /AFTER OPT_AIJ OPT_AIJ_BCK1
```

- Note that opt_aij_bck1 contains a single rolled-back transaction

```
$ sql$
SQL> attach 'filename opt_aij';
SQL> insert into t1 values (2);
1 row inserted
SQL> commit;
SQL> exit
$ RMU /BACKUP /AFTER OPT_AIJ OPT_AIJ_BCK2
```

! opt_aij_bck2 contains a single committed transaction

```
$ RMU /OPTIMIZE /AFTER OPT_AIJ_BCK1 OPT_AIJ_OPT1
$ RMU /OPTIMIZE /AFTER OPT_AIJ_BCK2 OPT_AIJ_OPT2
$ RMU /RESTORE /NOCDD OPT_AIJ
```

! Recovering the empty Optimized AIJ file opt_aij_opt1

```
$ RMU /RECOVER /LOG OPT_AIJ_OPT1.OAIJ
%RMU-I-LOGRECDB, recovering database file $111$DUA4:[VIGIER.OPTAIJ.DB]OPT_AIJ.RDB;1
%RMU-I-LOGOPNAIJ, opened journal file RAID1:[VIGIER.OPTAIJ.DB]OPT_AIJ_OPT1.OAIJ;1 at 8-JAN-200
%RMU-I-AIJONEDONE, AIJ file sequence 0 roll-forward operations completed
%RMU-I-LOGRECOVR, 0 transactions committed
%RMU-I-LOGRECOVR, 0 transactions rolled back
%RMU-I-LOGRECOVR, 0 transactions ignored
%RMU-I-AIJNOACTIVE, there are no active transactions
%RMU-I-AIJSUCCEs, database recovery completed successfully
%RMU-I-AIJNXTSEQ, to continue this AIJ file recovery, the sequence number needed will be 0
%RMU-I-AIJALLDONE, after-image journal roll-forward operations completed
%RMU-W-NOTRANAPP, no transactions in this journal were applied
%RMU-I-AIJSUCCEs, database recovery completed successfully
%RMU-I-AIJFNLSEQ, to start another AIJ file recovery, the sequence number needed will be 0
%RMU-I-AIJNOENABLED, after-image journaling has not yet been enabled
```

- Note that the next AIJ Sequence number has been left as 0

- Recovering the second Optimized AIJ file opt_aij_opt2

```
$ RMU /RECOVER /LOG OPT_AIJ_OPT2.OAIJ
%RMU-I-LOGRECD, recovering database file $111$DUA4:[VIGIER.OPTAIJ.DB]OPT_AIJ.RDB;1
%RMU-F-AIJNORCVR, recovery of this journal must start with sequence 0
%RMU-F-FTL_RCV, Fatal error for RECOVER operation at 8-JAN-2003 10:23:38.32
```

The recovery fails since it does not have the expected AIJ sequence number.

A workaround is to put all the optimized AIJ files in the same command line.

```
$ RMU /RECOVER /LOG OPT_AIJ_OPT1.OAIJ,OPT_AIJ_OPT2.OAIJ
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.2 Left Outer Join Query With OR Predicate Returns Wrong Results

Bugs 2836144 and 1837522

The following left outer join query with an OR predicate, having an equality predicate of constant values on the left side, and another equality predicate of a column and a constant value on the right side, returns the wrong result. It should find 274 rows but it finds only one row.

```
set flags 'strategy,detail';
select count(*) from
  job_history as c1
  left outer join
  employees as c2 on (c1.employee_id = c2.employee_id)
  where
    1=1 OR c1.job_code = 'JNTR';
Tables:
  0 = JOB_HISTORY
  1 = EMPLOYEES
Aggregate: 0:COUNT (*)
Conjunct: 0.JOB_CODE = 'JNTR'
Conjunct: 0.JOB_CODE = 'JNTR'
Match (Left Outer Join)
  Outer loop
    Conjunct: (1 = 1) OR (0.JOB_CODE = 'JNTR')
    Get Retrieval by index of relation 0:JOB_HISTORY
      Index name JH_EMPLOYEE_ID [0:0]
  Inner loop (zig-zag)
    Index only retrieval of relation 1:EMPLOYEES
      Index name EMP_EMPLOYEE_ID [0:0]

  1
1 row selected
```

This problem was partially fixed in Bug 1837522 where the equality contains one column and one constant.

As a workaround, the query works if the left and right side of the OR predicate is swapped as in the following example.

```
select count(*) from
  job_history as c1
  left outer join
```

```

employees as c2 on (c1.employee_id = c2.employee_id)
where
    c1.job_code = 'JNTR' OR 1=1;
Tables:
    0 = JOB_HISTORY
    1 = EMPLOYEES
Aggregate: 0:COUNT (*)
Conjunct: 0.JOB_CODE = 'JNTR'
Conjunct: 0.JOB_CODE = 'JNTR'
Match      (Left Outer Join)
Outer loop
    Conjunct: (0.JOB_CODE = 'JNTR') OR (1 = 1)
    Get      Retrieval by index of relation 0:JOB_HISTORY
             Index name  JH_EMPLOYEE_ID [0:0]
Inner loop      (zig-zag)
    Index only retrieval of relation 1:EMPLOYEES
    Index name  EMP_EMPLOYEE_ID [0:0]

        274
1 row selected

```

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.3 Left Outer Join Query With OR Predicate Returns Wrong Results

Bugs 2845172, 2836144 and 1837522

The following left outer join query with an OR predicate, having an equality predicate of constant values on the left side, and another equality predicate of a column and a constant value on the right side, returns the wrong result. It should find 274 rows, but it finds only 1 row.

```

set flags 'strategy,detail';
select
    t1.home_area_cd,t1.home_ph_num,
    t1.bus_area_cd,t1.bus_ph_num
from
    t1 left outer join t2  on t1.acct_num = t2.acct_num
where
    (t1.home_area_cd=999) OR
    (t1.bus_area_cd=310 and t1.bus_ph_num=5355400) ;
Tables:
    0 = T1
    1 = T2
Conjunct: 0.HOME_AREA_CD = 999                <== Notel : MISSING OR predicate
Conjunct: (0.BUS_AREA_CD = 310) AND (0.BUS_PH_NUM = 5355400)
Cross block of 2 entries      (Left Outer Join)
Cross block entry 1
OR index retrieval
    Conjunct: 0.HOME_AREA_CD = 999
    Get      Retrieval by index of relation 0:T1
             Index name  T1_IDX1 [1:1]
             Keys: 0.HOME_AREA_CD = 999
    Conjunct: NOT (0.HOME_AREA_CD = 999) AND (0.BUS_AREA_CD = 310) AND (
                0.BUS_PH_NUM = 5355400)
    Get      Retrieval by index of relation 0:T1
             Index name  T1_IDX2 [2:2]
             Keys: (0.BUS_AREA_CD = 310) AND (0.BUS_PH_NUM = 5355400)

```

```
Cross block entry 2
  Conjunct: 0.ACCT_NUM = 1.ACCT_NUM
  Get      Retrieval sequentially of relation 1:T2
0 rows selected
```

Note1: Notice that the OR predicate is missing between the two conjuncts.

This problem is similar to the query in Bug 2836144, but the only difference is that this query uses static OR retrieval strategy.

As a workaround, the query works if the left and right side of the OR predicate is swapped in this example but it won't do the trick with the customer's original query where the OR predicate has an additional equality predicate. For example:

```
select
  t1.home_area_cd,t1.home_ph_num,
  t1.bus_area_cd,t1.bus_ph_num
from
  t1 left outer join t2  on t1.acct_num = t2.acct_num
where
  (t1.home_area_cd=999 and t1.home_ph_num=999) or
  (t1.bus_area_cd=310 and t1.bus_ph_num=5355400);
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.4 Query Bugchecks When IN Clause Contains More Than Two DBKEYS

Bug 2742592

The following query bugchecks when the IN clause contains more than 2 dbkeys.

```
create table prova (i integer, c char(32));
insert into prova values (1,'one');
insert into prova values (2,'two');
insert into prova values (3,'three');

select * from prova where dbkey in (:dbk1,:dbk2,:dbk3);
declare :dbk1 char(8);
declare :dbk2 char(8);
declare :dbk3 char(8);
declare c1 table cursor for select dbkey from prova;
open c1;
fetch c1 into :dbk1;
fetch c1 into :dbk2;
fetch c1 into :dbk3;
close c1;
```

The following query works.

```
select * from prova where dbkey in (:dbk1,:dbk2);
OR index retrieval
  Conjunct      Firstn  Get      Retrieval by DBK of relation PROVA
  Conjunct      Firstn  Get      Retrieval by DBK of relation PROVA
      I      C
```

```
      1  one
      2  two
2 rows selected
```

However, with more than 2 dbkeys in the IN clause, it bugchecks.

```
select * from prova where dbkey in (:dbk1,:dbk2,:dbk3);
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USD04:[ONG]RDSBUGCHK.DMP;
```

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.5 Processes Loop at IPL2 When VLM Feature Used

Bug 2859466

It was possible for processes to become stuck in a loop in the Oracle Rdb image RDMPRV.EXE at VMS interrupt priority level (IPL) 2. Because the process was looping at an elevated IPL, it was not possible to delete the process. A system reboot was necessary to get rid of the looping processes. This problem only occurred when the Very Large Memory (VLM) feature was utilized for database shared memory or row caches.

Often the initial symptom of this problem was unresolved database stalls. Because the looping processes often held database locks, and since the looping prevented the processes from responding to blocking AST requests, other database processes would stall waiting for the looping processes to release their locks. Analysis of the processes not responding to blocking AST requests revealed that those processes were looping in the RDMPRV.EXE image.

To avoid this problem, disable the VLM feature for database shared memory or row caches.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.6 DBR Bugchecks at DBR\$DDTM_RESOLVE + 000003F4

The database recovery process (DBR) would sometimes bugcheck with the following failure:

```
***** Exception at 0005EC94 : DBR$DDTM_RESOLVE + 000003F4
%COSI-F-BUGCHECK, internal consistency failure
```

This bugcheck would occur when the OpenVMS system service \$GETDTIW would return an undocumented transaction state value (11, or "IN_DOUBT") for an unresolved transaction. The DBR would fail since it was not expecting that transaction state.

Subsequent attempts to open or attach to the database would usually succeed. The second query using the \$GETDTIW would typically return an expected state value.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.7 Replication Option and LogMiner Features Active at the Same Time

Bug 2903092

When both the Replication Option and LogMiner(TM) features are enabled, it is possible for the "pre-delete" record contents stored in the after-image journal file for the LogMiner to contain incorrect contents. This problem may be indicated by errors from the RMU /UNLOAD /AFTER_JOURNAL command such as the following example:

```
%RMU-W-RECVERDIF, Record at DBKEY 819:20615:8 in table
"FOOBAR" version 262 does not match current version.
```

This problem was caused by an incorrect buffer being used when writing the "pre-delete" record contents for the LogMiner. This buffer was also used by the Replication Option code path and the existing saved content was lost.

This problem has been corrected in Oracle Rdb Release 7.0.7.1. A different buffer is now used to save "pre-delete" record content data.

3.1.8 RMU /UNLOAD /AFTER_JOURNAL Created .RRD Content Clarification

Bug 2916639

An optional "RECORD_DEFINITION" keyword can be used with the RMU /UNLOAD /AFTER_JOURNAL command to create a template .RRD (record definition) file that can be used to load a transaction table. The TABLE_DEFINITION keyword can also be used to create a template SQL procedure to create such a transaction table.

The behaviour of the RMU /UNLOAD /AFTER_JOURNAL command when used with these keywords is to append the string "RDB_LM_" to the table name to create the record name in the .RRD or .SQL file.

However, when the existing table name exceeds 24 characters in length, the resultant name for the transaction table in the .SQL or .RRD file exceeds 31 characters and is no longer a valid table name in an Rdb database. In these cases, a decision about the table name to be used must be made and the .RRD or .SQL file must be manually modified.

Oracle Rdb engineering is considering alternatives for future releases to help reduce the impact of this behaviour.

3.1.9 Page Locks Not Released When LOCKING IS PAGE LEVEL

Bug 2959599

When the LOCKING IS PAGE LEVEL storage area attribute was enabled, it was possible for a process to obtain a page lock and not release it when a blocking AST was delivered by another process. Other processes would stall waiting for the page until the owning process removed the page from its buffer pool. The process could continue to hold the lock even after committing its transaction.

When this problem occurred, the output from the *RMU /SHOW LOCKS /MODE=CULPRIT* would be similar to the following:

```
=====
SHOW LOCKS/LOCK/MODE=CULPRIT Information
=====

-----
Resource: page 4443

      ProcessID Process Name      Lock ID   System ID Requested  Granted
-----
Blocker: 2541851A DKOM_TSG_004... 5300BA60  0001002A
Waiting: 25418513 DKOM_BRZ_PAS... 48001A32  0001002A  PW      NL
```

The OpenVMS system analyzer (SDA) utility SHOW LOCK command would show output similar to the following:

```
Lock id: 5300BA60          PID: 0061011A   Flags: VALBLK  CONVERT NOQUEUE
Par. id: 21010D65          SUBLCKs: 0      SYNCSTS SYSTEM PROTECT
LKB: FFFFFFFF.7E352E50    BLKAST: 00000000
Priority: 0000

Granted at PW 00000000-FFFFFFF

Resource: 0000115B 00000050 P...[...] Status: PROTCT
Length 08 00000000 00000000 .....
Exec. mode 00000000 00000000 .....
System 00000000 00000000 .....
```

Local copy

Note that in the above output, the blocking AST address (BLKAST) is zero and the lock is granted in PW mode.

To avoid this problem, set the storage area to LOCKING IS ROW LEVEL.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.10 Looping or Bugchecks in DIO\$FETCH_DBKEY for SORTED RANKED Indexes

Bugs 2936413 and 3017913

If a query specified a range of key values that fell between any existing key values in a SORTED RANKED index and the dynamic optimizer was used, it was possible to experience looping or bugchecks.

If a query specified a range of key values such that the range appeared in upper level index nodes but not in lower level index nodes, it was possible that Oracle Rdb could enter an infinite loop in the routines PSII2FINDSEPINTERVAL2 and PSII2EXPANDNODESEPS2.

If a query specified an empty range on an index, it was possible that bugchecks could occur such as:

```
SYSTEM-F-ACCVIO, access violation
```

```
Exception occurred at DIO$FETCH_DBKEY + 00000200  
Called from PSII2FINDSEPINTERVAL2 + 0000015C  
Called from PSII2ESTIMATECARD2 + 0000008C  
Called from RDMS$$NDX_ESTIM + 000003B0
```

An empty range is one such as that specified by the condition *WHERE MY_FIELD > 'A' AND MY_FIELD < 'B'*. If *MY_FIELD* were a *CHAR(1)* field, it would be impossible to find a key value to match the specified condition.

In some cases, Rdb would fail to detect that the range was empty and would try and pursue index estimation for the dynamic optimizer. This could result in Rdb attempting to read an erroneous dbkey, or to repeatedly re-read the same dbkey as the next index node.

Oracle Rdb now correctly detects an empty range and returns an estimate of zero rows for the index estimation.

The bugcheck problem can be avoided by specifying ranges that include at least one possible key value. This key value need not exist in the database.

The looping problem can be avoided by rebuilding the offending index or by disabling the dynamic optimizer.

This problem only occurs on indexes of *TYPE IS SORTED RANKED*. So another workaround is to use indexes of *TYPE IS HASHED* or *TYPE IS SORTED* if appropriate.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.11 RMU/CLOSE/WAIT Hangs Waiting for ALS to Terminate

It was possible for the *RMU/CLOSE/WAIT* command to hang if the database had *LOG SERVER IS AUTOMATIC* specified and a user attempted to attach to the database immediately after the *RMU/CLOSE* command was issued. This problem was introduced in Oracle Rdb Release 7.0.7.

The *RMU/CLOSE/WAIT* command will ensure that database recovery processes (DBRs) are executed before closing the database. In Release 7.0.7, if a user attempted to attach to the database while the DBR process was executing, when the DBR process completed, the database monitor would incorrectly start a new AIJ Log Server (ALS) process. The monitor would then wait for the ALS process to terminate before responding to the *RMU/CLOSE/WAIT* command. Since the ALS process had never been instructed to terminate, the database shutdown would never complete and no response was ever delivered to the RMU process that issued the close request.

The only way to avoid this problem is to prevent users from attempting to attach to the database immediately after it has been closed. When the problem occurs, the ALS process must be manually deleted by using the *DCL STOP/IDENTIFICATION* command or by issuing another *RMU/CLOSE* command with the */ABORT=DELPRC* qualifier.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.12 Query With EXISTS Clause and COMPUTED BY Column Returns Wrong Results

Bug 2452636

The following query, checking for EXISTS clause with an equality predicate involving a column of "COMPUTED BY", returns wrong results when applying a match strategy.

Columns for table T2:

Column Name	Data Type	Domain
A_DATE		DATE VMS
A_SEC_NAME		CHAR(12)
A_DET_METHOD		INTEGER
A_CALC_PRICE		INTEGER
A_BAD_CMPBY		INTEGER

```

Computed:  by
          case
            when (A_CALC_PRICE > 0 and
                  A_CALC_PRICE =
                    (select T0.A_PRICE from T0
                     where T0.A_SEC_NAME = T0.A_SEC_NAME
                       and T0.A_DATE = T0.A_DATE
                       and T0.A_PRICE_SOURCE = 'GIC'
                       and T0.A_PRICE_TYPE = 'STL' limit to 1 rows))
            then A_DET_METHOD
            else 5
            end

```

Indexes on table T2:

```

T2_NDX
          with column A_SEC_NAME
          and column A_DATE

```

```

set flags 'strategy,detail';
sel a_data from T1 s
where exists (select * from T2 p
             where p.A_SEC_NAME=s.A_SEC_NAME and
                   p.a_date=s.a_date and
                   p.a_bad_cmpby = 2);

```

Tables:

```

0 = T1
1 = T2
2 = T0

```

Conjunct: <agg0> <> 0

Match

Outer loop

Sort: 0.A_SEC_NAME(a), 0.A_DATE(a)

Cross block of 2 entries

Cross block entry 1

Aggregate: 1:VIA (2.A_PRICE)

Firstn: 1

Leaf#01 Ffirst 2:T0 Card=10

Bool: (2.A_SEC_NAME = 2.A_SEC_NAME) AND (2.A_DATE = 2.A_DATE) AND (2.A_PRICE_SOURCE = 'GIC') AND (2.A_PRICE_TYPE = 'STL')

BgrNdx1 T0_NDX [2:2] Fan=9

Keys: (2.A_PRICE_SOURCE = 'GIC') AND (2.A_PRICE_TYPE = 'STL')

Bool: (2.A_SEC_NAME = 2.A_SEC_NAME) AND (2.A_DATE = 2.A_DATE)

Cross block entry 2

Leaf#02 BgrOnly 0:T1 Card=2

BgrNdx1 T1_NDX [0:0] Fan=10

Inner loop (zig-zag)

Aggregate-F1: 0:COUNT-ANY (<subselect>)

Get Retrieval by index of relation 1:T2

Oracle® Rdb for OpenVMS

```
Index name  T2_NDX [0:0]
A_DATA
  14
  10
2 rows selected
```

Notice that the following equality predicate, "p.a_bad_cmpby = 2", involving the "Computed by" column A_BAD_CMPBY is missing in the above strategy and thus, the query returns the wrong result of 2 rows. P.a_bad_cmpby represents the following CASE statement:

```
case
  when (A_CALC_PRICE > 0 and
        A_CALC_PRICE =
          (select T0.A_PRICE from T0
           where T0.A_SEC_NAME = T0.A_SEC_NAME
             and T0.A_DATE = T0.A_DATE
             and T0.A_PRICE_SOURCE = 'GIC'
             and T0.A_PRICE_TYPE = 'STL' limit to 1 rows))
  then A_DET_METHOD
  else 5
end
```

Notice that the select statement references only a single table T0.

The missing conjunct is supposed to be generated in the following format with the aggregate value represented by <agg0>.

```
Conjunct: CASE (WHEN ((1.A_CALC_PRICE > 0) AND (1.A_CALC_PRICE = <agg0>))
                THEN 1.A_DET_METHOD ELSE 5) = 2
```

The reason that the conjunct is not created in the query is that the inner match leg contains only the context "1:T2" while the conjunct involves an aggregate value with external context "0:T0" from the outer match leg.

Here is one workaround for the problem. This query works if an outline is used to change the strategy from match to cross, since the inner cross leg contains both the contexts from the outer and inner cross legs.

```
sel a_data from T1 s
where exists (select * from T2 p
             where p.A_SEC_NAME=s.A_SEC_NAME and
                   p.a_date=s.a_date and
                   p.a_bad_cmpby = 2);
~S: Outline "TEST_BUG_OUTLINE" used
Tables:
  0 = T1
  1 = T2
  2 = T0
Cross block of 3 entries
Cross block entry 1
  Aggregate: 0:VIA (2.A_PRICE)
  Firstn: 1
  Leaf#01 FFirst 2:T0 Card=10
    Bool: (2.A_SEC_NAME = 2.A_SEC_NAME) AND (2.A_DATE = 2.A_DATE) AND (
          2.A_PRICE_SOURCE = 'GIC') AND (2.A_PRICE_TYPE = 'STL')
    BgrNdx1 T0_NDX [2:2] Fan=9
      Keys: (2.A_PRICE_SOURCE = 'GIC') AND (2.A_PRICE_TYPE = 'STL')
      Bool: (2.A_SEC_NAME = 2.A_SEC_NAME) AND (2.A_DATE = 2.A_DATE)
Cross block entry 2
  Leaf#02 FFirst 0:T1 Card=2
```

Oracle® Rdb for OpenVMS

```

    BgrNdx1 T1_NDX [0:0] Fan=10
Cross block entry 3
Conjunct: <agg1> <> 0
Aggregate-F1: 1:COUNT-ANY (<subselect>)
Leaf#03 FFirst 1:T2 Card=2
    Bool: (1.A_SEC_NAME = 0.A_SEC_NAME) AND (1.A_DATE = 0.A_DATE) AND (CASE (
        WHEN ((1.A_CALC_PRICE > 0) AND (1.A_CALC_PRICE = <agg0>
            1.A_DET_METHOD ELSE 5) = 2)
    BgrNdx1 T2_NDX [2:2] Fan=10
        Keys: (1.A_SEC_NAME = 0.A_SEC_NAME) AND (1.A_DATE = 0.A_DATE)
-- Rdb Generated Outline : 25-JUN-2003 11:30
create outline QO_F5E5D311487F5E17_00000000
id 'F5E5D311487F5E1776847CFB5A9C308B'
mode 0
as (
    query (
-- For loop
        subquery (
            subquery (
                T0 2    access path index      T0_NDX
            )
            join by cross to
            T1 0    access path index      T1_NDX
            join by cross to
            subquery (
                T2 1    access path index      T2_NDX
            )
        )
    )
)
compliance optional      ;
0 rows selected

```

Here is another possible workaround. The query also works if the COMPUTED BY column is changed to join the tables T0 and T2 instead of single table T0.

```

alter table T2 add column A_GOOD_CMPBY
computed by
case
when (A_CALC_PRICE > 0 and
    A_CALC_PRICE =
        (select T0.A_PRICE from T0
            where T0.A_SEC_NAME = T2.A_SEC_NAME
            and T0.A_DATE = T2.A_DATE
            and T0.A_PRICE_SOURCE = 'GIC'
            and T0.A_PRICE_TYPE = 'STL' limit to 1 rows))
then A_DET_METHOD
else 5
end;

sel a_data from T1 s
where exists (select * from T2 p
            where p.A_SEC_NAME=s.A_SEC_NAME and
            p.a_date=s.a_date and
            p.a_good_cmpby = 2);

Tables:
    0 = T1
    1 = T2
    2 = T0
Conjunct: <agg0> <> 0

```

```

Match
  Outer loop      (zig-zag)
    Get          Retrieval by index of relation 0:T1
      Index name T1_NDX [0:0]
    Inner loop
      Aggregate: 0:COUNT-ANY (<subselect>)
      Cross block of 2 entries
        Cross block entry 1
          Get      Retrieval by index of relation 1:T2
            Index name T2_NDX [0:0]
          Cross block entry 2
            Conjunct: CASE (WHEN ((1.A_CALC_PRICE > 0) AND (1.A_CALC_PRICE = <aggl>))
                          ) THEN 1.A_DET_METHOD ELSE 5) = 2
            Aggregate: 1:VIA (2.A_PRICE)
            Firstn: 1
            Leaf#01 FFirst 2:T0 Card=10
              Bool: (2.A_SEC_NAME = 1.A_SEC_NAME) AND (2.A_DATE = 1.A_DATE) AND (
                    2.A_PRICE_SOURCE = 'GIC') AND (2.A_PRICE_TYPE = 'STL')
              BgrNdx1 T0_NDX [4:4] Fan=9
                Keys: (2.A_PRICE_SOURCE = 'GIC') AND (2.A_PRICE_TYPE = 'STL') AND (
                      2.A_SEC_NAME = 1.A_SEC_NAME) AND (2.A_DATE = 1.A_DATE)
0 rows selected

```

Notice that the conjunct with aggregate value now appears in the Cross block entry 2 under the Inner loop of the match strategy.

This query works since the context "1:T2" is joined by cross strategy with the context "2:T0" and the conjunct with aggregate value is properly resolved with all the contexts available.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.13 Bugcheck in DIO\$FREE_CURRENT_LOCK for Sorted Ranked Indexes

Bug 2874671

If a cursor was used to fetch rows, and the transaction was committed between fetches, and the retrieval strategy involved a backwards scan of an index of *TYPE IS SORTED RANKED*, then Rdb could generate a bugcheck dump.

The problem occurred because Oracle Rdb mistakenly released a lock prematurely. Later, when Oracle Rdb was truly finished with the resource, a bugcheck occurred because a second attempt was made to release the same lock.

The following example uses the MF_PERSONNEL database to demonstrate the problem.

```

SQL> at 'f mf_personnel';
SQL> drop index EMP_EMPLOYEE_ID;
SQL> commit;
SQL> create unique index EMP_EMPLOYEE_ID
cont>      on EMPLOYEES (EMPLOYEE_ID asc)
cont>      type is SORTED ranked
cont>      node size 430
cont>      disable compression;
SQL> commit work;

```

Oracle® Rdb for OpenVMS

```
SQL> declare transaction read write isolation level read committed;
SQL> declare t2 table cursor with hold preserve all for
cont> select      *
cont> from        employees
cont> where       employee_id > '00300'
cont> and        employee_id < '00400'
cont> order by   employee_id desc;
SQL> open t2;
SQL> fetch t2;
EMPLOYEE_ID  LAST_NAME          FIRST_NAME  MIDDLE_INITIAL
ADDRESS_DATA_1  ADDRESS_DATA_2    CITY
STATE  POSTAL_CODE  SEX  BIRTHDAY    STATUS_CODE
00374      Andriola      Leslie     Q
111 Boston Post Rd.                Salisbury
NH      03268        M      19-Mar-1955  1

SQL> commit;
SQL> fetch t2;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file
MBRADLEY_USR:[BRADLEY]RDSBUGCHK.DMP;
```

The problem can be avoided by disabling the backward scan feature using the *RDMS\$DISABLE_REVERSE_SCAN* logical name.

The problem does not occur if all rows are fetched in the same transaction.

The problem does not occur if the index being scanned is not of *TYPE IS SORTED RANKED*.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.14 Ranked Index Overflow Node Corruption on Insert

Bug 3009262

A problem in the way Rdb chooses the insertion point for the dbkey in a Ranked Index duplicate node may, in rare circumstances, cause a corruption of the index entry overflow node. This node corruption may manifest itself as a bugcheck when trying to access the records associated with the overflow node, as in the following example.

```
***** Exception at 00C72D78 : PSII2SCANGETNEXTBBCDUPLICATE + 00000128
%COSI-F-BUGCHECK, internal consistency failure
```

Or, this node corruption may be seen as a corrupt index warning when *RMU/VERIFY/INDEX* is used to verify the ranked index. For example:

```
%RMU-I-BTRDUPCAR, Inconsistent duplicate cardinality (C1) of 221 specified
                    for entry 1 at dbkey 85:807:1.
                    Actual count of duplicates is 251.
%RMU-I-BTRERPATH, parent B-tree node of 85:807:1 is at 85:806:0
%RMU-I-BTRROODBK, root dbkey of B-tree is 85:806:0
%RMU-W-DATNOTIDX, Row in table TT11 is not in any indexes.
                    Logical dbkey is 83:82:2.
%RMU-W-BADIDXREL, Index INDTT either points to a non-existent record or
                    has multiple pointers to a record in table TT11.
                    The logical dbkey in the index is 83:127:8.
```

A dump of the offending overflow node will show that even though the reference dbkey is correct, all the dbkeys in the overflow node may be incorrect and possibly not valid record dbkeys.

The problem may only occur on Sorted Ranked index entries that are volatile enough so that more than one overflow node is required to hold the duplicates and that, due to removal of dbkeys from the entry, at least one non-final overflow node has subsequently been removed from the entry.

A subsequent insertion of a new duplicate in that entry may, if the insertion dbkey happens to be greater than the last dbkey in an overflow node but less than the reference dbkey of the next overflow node, incorrectly update the overflow node causing the subsequent corruption.

This problem only occurs with Sorted Ranked indexes.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.15 Illegal Page Count Error in the Dynamic Optimizer

Bug 3045841

When a very large table was queried, and the access strategy used the dynamic optimizer, and the first two indexes used both returned more than 1024 dbkeys, an illegal page count error could be generated.

To encounter this error, the table had to have close to one billion (1,000,000,000) rows.

In addition, the retrieval strategy for the query must use the dynamic optimizer where at least two indices return more than 1024 dbkeys.

In this case, the dynamic optimizer will allocate memory for a dbkey bitmap. In calculating the size of that bitmap, integer (signed longword) arithmetic was used, and an integer overflow could cause the calculation to result in a negative number being used as the requested amount of memory.

The following example shows a query on a table containing one billion (1,000,000,000) rows.

```
SQL> select * from t1
cont>         where f1 >0 and f2 > 0
cont>         optimize for total time;
%COSI-F-UNEXPERR, unexpected system error
-SYSTEM-F-ILLPAGCNT, illegal page count parameter
```

A bugcheck dump may also be generated with the following exception and call sequence:

```
***** Exception at 00FE3664 : COSI_MEM_GET_VM + 000007B4
%COSI-F-UNEXPERR, unexpected system error
-SYSTEM-F-ILLPAGCNT, illegal page count parameter
Saved PC = 00FE2E68 : COSI_MEM_GET_POOL + 00000048
Saved PC = 00E3CCE8 : RDMS$$EXE_LEAF + 00001A38
Saved PC = 00E290B4 : RDMS$$EXE_OPEN + 00000764
```

The problem would not occur if any of the conditions described above were not true.

The problem can be avoided by disabling the dynamic optimizer for the query by using a query outline with the clause *EXECUTION OPTIONS NONE*.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.16 Bugcheck During Create Index with Mapping Values

Bug 2992315

SQL Create Index on a table with multiple storage areas and mapping values fails with a bugcheck.

The following example shows a typical stack trace.

```
***** Exception at 007CFA67 : PSIIBUILD$BUILD_FROM_BOTTOM + 00000693
Saved PC = 80000014 : symbol not found
Saved PC = 007CC32E : PSII$CREATE_TREE + 000000B6
Saved PC = 006AE447 : RDMS$$KOD_CREATE_TREE + 00000168
Saved PC = 006AE1AF : RDMS$$KOD_CREATE_INDEX + 00000320
Saved PC = 0061C416 : RDMS$$CREATE_INDEX_INFO + 0000234F
```

Possible workarounds include: create the index on an Alpha node or do not use mapping values.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.17 Error %RDMS-E-NOSOL_FOUND in Full Outer Join Query

Bug 2669656

The following full outer natural join query between columns of different data types fails using the cross strategy (it should succeed applying the match strategy):

```
create table x1 (f1 char(10));
create table x2 (f1 integer);
create table x3 (f1 char(10));

insert into x1 value ('1');
insert into x2 value (1);
insert into x2 value (-1);
insert into x3 value ('1');
insert into x3 value ('-1');

select * From x1 natural full outer join x2;
~S: Full OJ query with cross strategy was not possible
%RDMS-E-NOSOL_FOUND, No possible solution has been found by Rdb optimizer
```

As a workaround, the following full outer join between columns of the same data type works applying a match strategy.

```
select * From x1 natural full outer join x3;
Tables:
  0 = X1
  1 = X3
Match (Full Outer Join)
  Outer loop
    Sort: 0.F1(a)
    Get Retrieval sequentially of relation 0:X1
```

```

Inner loop
  Temporary relation
  Sort: 1.F1(a)
  Get Retrieval sequentially of relation 1:X3
F1
-1
1
2 rows selected

```

The following query should apply a match strategy as suggested by the outline bug_outline.

```

create outline bug_outline
id '6CBC3F110B75FD48C256CE94DCEB8A1F'
mode 0
as (
  query (
-- For loop
    subquery (
      X1 0    access path sequential
      join by match to
      X2 1    access path sequential
    )
  )
)
compliance optional      ;

select * from x1 , x2 where x1.f1 = x2.f1;
~S: Outline "BUG_OUTLINE" used
~S: Full compliance with the outline was not possible
Tables:
  0 = X1
  1 = X2
Cross block of 2 entries
Cross block entry 1
  Get      Retrieval sequentially of relation 0:X1
Cross block entry 2
  Conjunct: 0.F1 = 1.F1
  Get      Retrieval sequentially of relation 1:X2
X1.F1      X2.F1
1          1
1 row selected

```

But the query works if one of the join predicates is cast as the same data type as the other as in the following example.

```

create outline bug_good_outline
id '0EEB4BC11CF012EDB10AEA1C16EC0E70'
mode 0
as (
  query (
-- For loop
    subquery (
      X1 0    access path sequential
      join by match to
      X2 1    access path sequential
    )
  )
)
compliance optional      ;

```



```

select * from x1 , x2 where cast(x1.f1 as integer) = x2.f1;
~S: Outline "BUG_GOOD_OUTLINE" used
Tables:
  0 = X1
  1 = X2
Conjunct: CAST (0.F1 AS INT) = 1.F1
Match
  Outer loop
    Sort: CAST (0.F1 AS INT)(a)
    Get Retrieval sequentially of relation 0:X1
  Inner loop
    Temporary relation
    Sort: 1.F1(a)
    Get Retrieval sequentially of relation 1:X2
X1.F1          X2.F1
1              1
1 row selected

```

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.18 TRUNCATE TABLE and RMU /REPAIR Corruption Corrected

Previously, when using the TRUNCATE TABLE statement followed by a RMU /REPAIR /ABM command executed one or more times, various corruptions of the SPAM and ABM structures could occur. This could sometimes lead to table records that had been truncated "reappearing" in the table unexpectedly.

The following example demonstrates one possible symptom of this problem detected by RMU /VERIFY:

```

$ SQL$
  CREATE DATABASE FILENAME 'TEST.RDB'
    CREATE STORAGE AREA RDB$SYSTEM FILENAME 'RDB$SYSTEM.RDA'
    CREATE STORAGE AREA AREA FILENAME 'AREA.RDA';
  CREATE TABLE TAB (ID INTEGER);
  CREATE INDEX IND ON TAB (ID) STORE IN AREA;
  CREATE STORAGE MAP TM FOR TAB STORE IN AREA;
  COMMIT;
  INSERT INTO TAB VALUES (1);
1 row inserted
  COMMIT;
  TRUNCATE TABLE TAB;
  COMMIT;
  EXIT;
$!
$ RMU/VER/ALL/NOLOG TEST
%RMU-I-NODATANDX, no data records in index IND
$!
$ RMU/REPAIR/ABM TEST
%RMU-I-FULBACREQ, A full backup of this database should be
performed after RMU REPAIR
$ RMU/VER/ALL/NOLOG TEST
%RMU-W-AIPLAREID, area inventory page 152 entry #9 contains a
reference to logical area 58 that is nonexistent
%RMU-W-BADABMPTR, invalid larea for ABM page 5 in storage area 2.
The SPAM page entry for this page is for a different larea.
SPAM larea_dbid : 0 page larea_dbid: 58.
%RMU-W-BADABMPTR, invalid larea for ABM page 6 in storage area 2.
The SPAM page entry for this page is for a different larea.

```

Oracle® Rdb for OpenVMS

```
SPAM larea_dbid : 0 page larea_dbid: 58.
%RMU-W-BADABMPTR, invalid larea for ABM page 7 in storage area 2.
The SPAM page entry for this page is for a different larea.
SPAM larea_dbid : 0 page larea_dbid: 58.
%RMU-E-BADABMPAG, error verifying ABM pages
%RMU-I-NODATANDX, no data records in index IND
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1. The RMU /REPAIR /ABM command no longer incorrectly updates the SPAM and ABM structures after a TRUNCATE TABLE operation.

3.1.19 UNION Query With Two Left Outer Joins in First Leg Returns Wrong Results

Bugs 3076004 and 2529598

The following UNION query with left outer join should return 1 row.

```
set flags 'strategy,detail';
select routing_id from
  (select
    C4.ROUTING_ID,
    C2.FY,
    C2.PAY_PERIOD
  from
    ROSTER as C2
  left outer join
    WORK_AUTH as C4
    on (C2.AUTH_NBR = C4.AUTH_NBR)
  left outer join
    TAX_BEN as C5
    ON (C2.AUTH_NBR = C5.TAX_BEN_NBR)

  union

  select
    C6.ROUTING_ID,
    C7.FY,
    C7.PAY_PERIOD
  from
    WORK_AUTH as C6, PAY_PERIOD as C7)
  AS DT (ROUTING_ID, FY, PAY_PERIOD)
where
  fy='2004' and pay_period='01' and routing_id = 'R91297';
```

Tables:

```
0 = ROSTER
1 = WORK_AUTH
2 = TAX_BEN
3 = WORK_AUTH
4 = PAY_PERIOD
```

Merge of 1 entries

Merge block entry 1

Reduce: <mapped field>, <mapped field>, <mapped field>

Sort: <mapped field>(a), <mapped field>(a), <mapped field>(a)

Merge of 2 entries

Merge block entry 1

Cross block of 2 entries (Left Outer Join)

Cross block entry 1

Cross block of 2 entries (Left Outer Join)

Cross block entry 1

Leaf#01 BgrOnly 0:ROSTER Card=2

Oracle® Rdb for OpenVMS

```
BgrNdx1 ROSTER_NDX [2:2] Fan=15
  Keys: (<mapped field> = '2004') AND (<mapped field> = '01')
Cross block entry 2
  Leaf#02 BgrOnly 1:WORK_AUTH Card=1
  Bool: 0.AUTH_NBR = 1.AUTH_NBR
  BgrNdx1 WORK_AUTH_NDX [1:1] Fan=16
  Keys: <mapped field> = 'R91297'
Cross block entry 2
  Conjunct: 0.AUTH_NBR = 2.TAX_BEN_NBR
  Index only retrieval of relation 2:TAX_BEN
  Index name TAX_BEN_NDX [1:1]
  Keys: 0.AUTH_NBR = 2.TAX_BEN_NBR
Merge block entry 2
Cross block of 2 entries
  Cross block entry 1
  Conjunct: 3.ROUTING_ID = 'R91297'
  Index only retrieval of relation 3:WORK_AUTH
  Index name WORK_AUTH_NDX [1:1]
  Keys: <mapped field> = 'R91297'
  Cross block entry 2
  Conjunct: (4.FY = '2004') AND (4.PAY_PERIOD = '01')
  Index only retrieval of relation 4:PAY_PERIOD
  Index name PAY_PERIOD_NDX [2:2]
  Keys: (<mapped field> = '2004') AND (<mapped field> = '01')
ROUTING_ID
R91297
NULL
2 rows selected
```

This problem is similar to the previous Bug 2529598 where the fix did not cover the current case with two left outer joins in the first UNION leg. The conjunct "routing_id = 'R91297'" is generated at the top of the second UNION (merge) leg but not at the top of the first leg and thus returns the wrong result.

The query works if the legs of the UNION clause are swapped, as in the following example:

```
select routing_id from
(
  select
    C6.ROUTING_ID,
    C7.FY,
    C7.PAY_PERIOD
  from
    WORK_AUTH as C6, PAY_PERIOD as C7
union
  select
    C4.ROUTING_ID,
    C2.FY,
    C2.PAY_PERIOD
  from
    ROSTER as C2
    left outer join
      WORK_AUTH as C4
        on (C2.AUTH_NBR = C4.AUTH_NBR)
    left outer join
      TAX_BEN as C5
        ON (C2.AUTH_NBR = C5.TAX_BEN_NBR)
)
AS DT (ROUTING_ID, FY, PAY_PERIOD)
where
  fy='2004' and pay_period='01' and routing_id = 'R91297';
```

Oracle® Rdb for OpenVMS

Tables:

```
0 = WORK_AUTH
1 = PAY_PERIOD
2 = ROSTER
3 = WORK_AUTH
4 = TAX_BEN
```

Conjunct: <mapped field> = 'R91297'

Merge of 1 entries

Merge block entry 1

Reduce: <mapped field>, <mapped field>, <mapped field>

Sort: <mapped field>(a), <mapped field>(a), <mapped field>(a)

Conjunct: 1.FY = '2004'

Conjunct: 1.PAY_PERIOD = '01'

Merge of 2 entries

Merge block entry 1

Cross block of 2 entries

Cross block entry 1

Conjunct: 0.ROUTING_ID = 'R91297'

Index only retrieval of relation 0:WORK_AUTH

Index name WORK_AUTH_NDX [1:1]

Keys: <mapped field> = 'R91297'

Cross block entry 2

Conjunct: (1.FY = '2004') AND (1.PAY_PERIOD = '01')

Index only retrieval of relation 1:PAY_PERIOD

Index name PAY_PERIOD_NDX [2:2]

Keys: (<mapped field> = '2004') AND (<mapped field> = '01')

Merge block entry 2

Cross block of 2 entries (Left Outer Join)

Cross block entry 1

Cross block of 2 entries (Left Outer Join)

Cross block entry 1

Leaf#01 BgrOnly 2:ROSTER Card=2

BgrNdx1 ROSTER_NDX [2:2] Fan=15

Keys: (<mapped field> = '2004') AND (<mapped field> = '01')

Cross block entry 2

Conjunct: 2.AUTH_NBR = 3.AUTH_NBR

Get Retrieval by index of relation 3:WORK_AUTH

Index name WORK_AUTH_NDX [0:0]

Cross block entry 2

Leaf#02 BgrOnly 4:TAX_BEN Card=1

Bool: 2.AUTH_NBR = 4.TAX_BEN_NBR

BgrNdx1 TAX_BEN_NDX [1:1] Fan=17

Keys: 2.AUTH_NBR = 4.TAX_BEN_NBR

ROUTING_ID

R91297

1 row selected

Notice that the "Conjunct: <mapped field> = 'R91297'" is now generated on top of both UNION legs. Even though the conjunct is NOT efficiently optimized by being pushed down to the second leg, at least the query works correctly.

There is no known workaround other than the above-mentioned query with UNION legs swapped.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.20 Logical Area Record Erasure Count Not Updated for Cached Rows

Bug 3099718

Previously, logical area statistics for record erase operations were not correctly counted when erasing rows in a row cache.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.21 Query With Sum Function of Two Select Counts Bugchecks

Bug 2649215

A query with a sum function of two select counts could produce a bugcheck.

```

select sum ((select count (*) - (select count (*) from T1
                                where T1.F1 = T2.F1
                                and T1.F2 = T2.F2)
            from T2
            where F3 = 'B'
            group by T2.F1, T2.F2)
)
from rdb$database;
%DEBUG-I-DYNMODSET, setting module RDMS$PREEXEASN
%SYSTEM-F-BREAK, breakpoint fault at PC=003995E9, PSL=03C00004
break on exception at RDMS$PREEXEASN\RDMS$$FIND_VALID_SEG_CRTV\%LINE 8443 in
EAD 1

```

The query works if one of the equality predicates is removed, as in the following example.

```

select sum ((select count (*)
            - (select count (*) from T1 where
              T1.F1 = T2.F1
              and T1.F2 = T2.F2
            )
            from T2 where F3 = 'B'
            group by T2.F1, T2.F2
            ))
from rdb$database;
Tables:
  0 = RDB$DATABASE
  1 = T2
  2 = T1
Aggregate: 0:SUM (<agg1>)
Cross block of 2 entries
Cross block entry 1
Aggregate: 1:VIA (<mapped field> - <agg2>)
Cross block of 2 entries
Cross block entry 1
Aggregate: 2:COUNT (*)
Index only retrieval of relation 2:T1
Index name U1_T1 [1:1]
Keys: 2.F1 = 1.F1

```

```

Cross block entry 2
  Aggregate: 3:COUNT (*)
  Conjunct: 1.F3 = 'B'
  Get      Retrieval by index of relation 1:T2
           Index name U1_T2 [0:0]
Cross block entry 2
  Retrieval sequentially of relation 0:RDB$DATABASE

```

1
1 row selected

The query also works if the SUM function is removed.

```

select count (*) - (select count (*) from T1
                    where T1.F1 = T2.F1
                    and T1.F2 = T2.F2)
                from T2
                where F3 = 'B'
                group by T2.F1, T2.F2
                ;

```

Tables:

```

0 = T2
1 = T1

```

Cross block of 2 entries

```

Cross block entry 1
  Aggregate: 0:COUNT (*)
  Conjunct: 0.F3 = 'B'
  Get      Retrieval by index of relation 0:T2
           Index name U1_T2 [0:0]
Cross block entry 2
  Aggregate: 1:COUNT (*)
  Index only retrieval of relation 1:T1
  Index name U1_T1 [2:2]      Direct lookup
  Keys: (1.F1 = 0.F1) AND (1.F2 = 0.F2)

```

0
1 row selected

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.22 Left Outer Join Query With CONCAT Function Returns Wrong Results

Bugs 3139961, 2836144, 1837522

The following query with CONCAT function returns wrong results (should return 4 rows).

```

set flags 'strategy,detail';
select T1.YEAR,T1.WEEK,T1.KODE
FROM T1 LEFT OUTER JOIN T2
      ON T1.YEAR = T2.YEAR AND T1.KODE = T2.KODE,
T3 where
(T1.YEAR||T1.WEEK <='200337') and
T3.IPROC = T1.IPROC AND
T3.ITeam = T1.ITeam      ;

```

Tables:

0 = T1
 1 = T2
 2 = T3

Cross block of 2 entries

Cross block entry 1

Index only retrieval of relation 2:T3

Index name T3_IND1 [0:0]

Cross block entry 2

Conjunct: ((0.YEAR < SUBSTRING ('200337' FROM 0 FOR 4)) AND
 <error: missing expression>) OR ((0.YEAR = SUBSTRING ('200337'
 FROM 0 FOR 4)) AND (0.WEEK <= SUBSTRING ('200337' FROM 4)))

Cross block of 2 entries (Left Outer Join)

Cross block entry 1

Leaf#01 FFirst 0:T1 Card=4

Bool: (((0.YEAR < SUBSTRING ('200337' FROM 0 FOR 4)) AND NOT
 MISSING (0.WEEK)) OR ((0.YEAR = SUBSTRING ('200337'
 FROM 0 FOR 4)) AND (0.WEEK <= SUBSTRING ('200337' FROM 4)))
) AND (2.IPROC = 0.IPROC) AND (2 ITEAM = 0.ITEAM)

BgrNdx1 T1_IND2 [0:0] Fan=15

BgrNdx2 T1_IND1 [0:0] Fan=14

Bool: ((0.YEAR < SUBSTRING ('200337' FROM 0 FOR 4)) AND NOT
 MISSING (0.WEEK)) OR ((0.YEAR = SUBSTRING ('200337'
 FROM 0 FOR 4)) AND (0.WEEK <= SUBSTRING ('200337' FROM 4)
))

Cross block entry 2

Conjunct: (0.YEAR = 1.YEAR) AND (0.KODE = 1.KODE)

Index only retrieval of relation 1:T2

Index name T2_IND1 [2:2] Direct lookup

Keys: (0.YEAR = 1.YEAR) AND (0.KODE = 1.KODE)

T1.YEAR	T1.WEEK	T1.KODE
2003	26	270
2003	34	270

2 rows selected

The problem is caused by the fix made for Bug 1837522 in Oracle Rdb Release 7.0.6.2 where a left outer join query with OR predicate returns wrong results.

Even though this query apparently does not contain an OR predicate, the Oracle Rdb optimizer transforms the CONCAT function into an OR expression with two SUBSTRING functions, as seen in the following detail strategy.

```

Bool: (((0.YEAR < SUBSTRING ('200337' FROM 0 FOR 4)) AND
NOT MISSING (0.WEEK)) OR
((0.YEAR = SUBSTRING ('200337' FROM 0 FOR 4)) AND
(0.WEEK <= SUBSTRING ('200337' FROM 4)))) AND
(2.IPROC = 0.IPROC) AND (2.ITEAM = 0.ITEAM)
    
```

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.23 RCS Bugchecks at DIOCCH\$UNMARK_GRIC_ENT

Bug 2502144

In prior releases of Oracle Rdb, it was possible for the Record Cache Server (RCS) process to bugcheck in DIOCCH\$UNMARK_GRIC_ENT. This problem was due to an incorrect check in the RCS process while evaluating the amount of locked space on a database page that is owned by the RCS.

When the RCS process writes an erased or resized record back to the database page, it must return the resultant locked space on the page back to free space. This is because the RCS is not allowed to "own" locked space. As part of this action, the RCS was checking to make sure that the amount of locked space being returned exactly matched the length of the space being returned by the erased or resized record. This check was not taking into account the possibility that there was existing locked space on the page marked with the TID of the RCS. In such a case, the RCS could find that it "owned" more locked space than it expected and would bugcheck.

This problem has been corrected in Oracle Rdb Release 7.0.7.1. The RCS process now correctly evaluates and releases all locked space on the page that is owned by the RCS after writing an erased or resized record back to the database.

3.1.24 Wrong Sort Order for Query with Aggregate, Group By, Order By

Bug 3077857

Before the problem was fixed, the query in the example below returned the results in ascending order of the NTS_FLAG column, in contradiction to the explicit DESCENDING attribute in the ORDER BY clause. The problem appeared when RDM\$BIND_BUFFERS was set to 200 but not when it was set to 20.

```
set flags 'detail,strategy';

select  Max(SERVICE_SUM.NTS_FLAG) NTS_FLAG
  from  SERVICE_SUM, INFO_PROVIDER_CODE_PARENT
  where SERVICE_SUM.CALL_END_DATE between
        DATE ANSI '2003-07-16' and DATE ANSI '2003-07-16'
        and
        SERVICE_SUM.DIALED_NO = INFO_PROVIDER_CODE_PARENT.DIALED_NO
  group by SERVICE_SUM.NTS_FLAG
  order by SERVICE_SUM.NTS_FLAG DESC;
```

The change in the value for RDM\$BIND_BUFFERS had the effect of changing the optimizer query strategy from a CROSS join to a MATCH join. Using interactive SQL, for example, this could be seen by first enabling output of the optimizer strategy (see the SET FLAGS statement in the preceding example). The combination of MATCH join, aggregate query, ascending GROUP BY sort, and descending ORDER BY sort, plus the optimizer logic to try to eliminate unnecessary sorts were the conditions needed for the problem to appear.

As a workaround, the problem could be made to go away by reducing the value in RDM\$BIND_BUFFERS from 200 to 20 or by creating a query outline forcing a CROSS join strategy.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.25 Left Outer Join Query with SUBSTRING and CHAR_LENGTH Bugchecks

Bug 2851595

The following left outer join query with SUBSTRING and CHAR_LENGTH bugchecks during the process of creating a match retrieval strategy.

```
select e.employee_id
  from employees e left outer join job_history jh
    on (e.employee_id =
        substring (jh.employee_id from 1
                   for char_length (e.employee_id));
%DEBUG-I-DYNMODSET, setting module RDMS$PREEXEMSC
%SYSTEM-F-BREAK, breakpoint fault at PC=003A4CD6, PSL=03C00004
break on exception at RDMS$PREEXEMSC\RDMS$$FIND_MEMBER_EQV_CLASS\%LINE 4486
DBG> SH CA
  module name      routine name      line          rel PC      ab
*RDMS$PREEXEMSC  RDMS$$FIND_MEMBER_EQV_CLASS
*RDMS$PREEXE     RDMS$$SETUP_SORT32
*RDMS$PREEXE_CREATE
                  RDMS$$CREATE_RSS$MTCH
```

The bugcheck is caused by a match order using equi-join keys of the outer join query, created from the following predicate:

```
(e.employee_id = substring (jh.employee_id from 1
                           for char_length (e.employee_id))
```

where the match key on the right side depends on its own context "employees e" from its left side.

As a workaround, the query works if the function CHAR_LENGTH is applied with JH table:

```
select e.employee_id
  from employees e left outer join job_history jh
    on (e.employee_id =
        substring (jh.employee_id from 1
                   for char_length (jh.employee_id));
```

Tables:

```
0 = EMPLOYEES
1 = JOB_HISTORY
```

Match (Left Outer Join)

Outer loop

```
Index only retrieval of relation 0:EMPLOYEES
  Index name EMP_EMPLOYEE_ID [0:0]
```

Inner loop

```
Temporary relation
Sort: SUBSTRING (1.EMPLOYEE_ID FROM (1 - 1) FOR CHAR_LENGTH (1.EMPLOYEE_ID))
(a)
```

```
Index only retrieval of relation 1:JOB_HISTORY
  Index name JH_EMPLOYEE_ID [0:0]
```

```
E.EMPLOYEE_ID
00164
00164
...etc...
```

274 rows selected

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.26 Join Query with GROUP_BY/ORDER_BY Returns Wrong Order

Bug 2851708

The following query with GROUP_BY/ORDER_BY, joining two derived tables with similar GROUP_BY clauses, returns results in the wrong order (should return in descending order).

```

set flags 'strategy, detail';
select  t1.name, t2.name, t1.datum, t2.datum
  from (select name, datum from a
        group by name, datum) t1
      join
        (select name, datum from b
        group by name, datum) t2
      on (t1.name = t2.name and t1.datum = t2.datum)
  group by t1.name, t2.name, t1.datum, t2.datum
  order by t1.name desc, t1.datum desc ;
Tables:
  0 = A
  1 = B
Reduce: 0.NAME, 0.DATUM, 1.NAME, 1.DATUM
Cross block of 2 entries
  Cross block entry 1
    Merge of 1 entries
      Merge block entry 1
        Reduce: 0.NAME, 0.DATUM
        Sort: 0.NAME(a), 0.DATUM(a)          <== See Note:
        Get Retrieval sequentially of relation 0:A
  Cross block entry 2
    Merge of 1 entries
      Merge block entry 1
        Reduce: 1.NAME, 1.DATUM
        Sort: 1.NAME(d), 1.DATUM(d)
        Conjunct: (0.NAME = 1.NAME) AND (0.DATUM = 1.DATUM)
        Get Retrieval sequentially of relation 1:B
T1.NAME   T1.DATUM
AAAA     1-JAN-2000 00:00:00.00
BBBB     1-JAN-2000 00:00:00.00
2 rows selected

```

NOTE: The sort keys should be descending instead of ascending.

This problem occurs when the main query contains a GROUP BY clause on the columns of the two joined derived tables with GROUP BY with the ORDER BY clause referencing the columns of the first table using descending order.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.27 Query Joining Two Derived Tables of a View with UNION Overflows the Stack

Bug 3194445

The following query, joining two derived tables of a view with UNION, overflows the stack.

```
select
  (select period_start_date from gl_period_vw
   where
     company_no = glpv.company_no and
     fiscal_yr = glpv.fiscal_yr
   )
  as p2_period_start_date,
  (select period_end_date from gl_period_vw
   where
     company_no = glpv.company_no and
     fiscal_yr = glpv.fiscal_yr
   )
  as p2_period_end_date
from gl_period_vw glpv, gl_period glp
;
%RDB-F-IMP_EXC, facility-specific limit exceeded
-RDB-I-TEXT, internal error -- query solution not found
```

The view is defined as:

```
create view GL_PERIOD_VW
  (COMPANY_NO,
   FISCAL_YR,
   PERIOD_START_DATE,
   PERIOD_END_DATE) as
select
  C2.COMPANY_NO,
  C2.FISCAL_YR,
  C2.PERIOD_START_DATE,
  C2.PERIOD_END_DATE
from GL_PERIOD C2
  where (C2.PERIOD_NO = 0)
union all
select
  C3.COMPANY_NO,
  C3.FISCAL_YR,
  C3.PERIOD_START_DATE,
  C3.PERIOD_END_DATE
from GL_PERIOD C3, INTEGER_LIST C4
  where (C3.PERIOD_NO = 0)
;
```

This problem is caused by the fix made for Bug 2649215 where a query before bugchecks.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.28 Query with Shared Expression in Two Predicates Returns Wrong Results

Bug 3201864

The following query with shared expression in two predicates should return 1 row but instead returns 2 rows.

```

set flags 'strategy,detail';

SELECT  T1.CODE,T2.ALUM, T1.ALUM, T2.CODE
FROM    T1, T2, T3 WHERE
        T2.TUNE      = T1.TUNE      AND
        T2.SLAM      = T1.SLAM      AND
        T2.STAR      = T1.STAR      AND
        (T1.CODE     >= T2.ALUM      AND
         (T1.ALUM    <= T2.CODE OR T2.CODE IS NULL)) AND
        T3.TUNE      = T1.TUNE      AND
        T3.SLAM      = T1.SLAM      AND
        T3.STAR      = T1.STAR      AND
        T2.TUNE      = '240167-1447' AND
        T2.SLAM      <= '31-dec-2002' AND
        (T2.CODE     >= '01-jan-2002' OR T2.CODE IS NULL) AND
        NOT EXISTS
        (SELECT T4.TUNE FROM T2 T4 WHERE
         T4.TUNE = T2.TUNE      AND
         T4.MAKE = '0'          AND
         T4.SLAM = T2.SLAM      AND
         T4.STAR = T2.STAR AND
         T4.SITE <> T2.SITE AND
         T1.ALUM >= T4.ALUM AND
         T1.CODE <= T4.CODE );

```

Tables:

```

0 = T1
1 = T2
2 = T3
3 = T2

```

Cross block of 4 entries

Cross block entry 1

Conjunct: 1.TUNE = '240167-1447'

Conjunct: 1.SLAM <= '31-DEC-2002'

Leaf#01 FFirst 1:T2 Card=2

Bool: (1.CODE >= '1-JAN-2002') OR MISSING (1.CODE)

BgrNdx1 T2_IND_1 [1:2] Fan=7

Keys: (1.TUNE = '240167-1447') AND (1.SLAM <= '31-DEC-2002')

Cross block entry 2

Leaf#02 FFirst 0:T1 Card=1

Bool: (1.TUNE = 0.TUNE) AND (1.SLAM = 0.SLAM) AND (1.STAR = 0.STAR) AND
(0.CODE >= 1.ALUM)

BgrNdx1 T1_IND [1:1] Fan=11

Keys: 1.TUNE = 0.TUNE

BgrNdx2 T1_IND_1 [3:3] Fan=8

Keys: (1.TUNE = 0.TUNE) AND (1.SLAM = 0.SLAM) AND (1.STAR = 0.STAR)

Cross block entry 3

Conjunct: <agg0> = 0

Aggregate-F1: 0:COUNT-ANY (<subselect>)

Leaf#03 FFirst 3:T2 Card=2

Bool: (3.TUNE = 1.TUNE) AND (3.MAKE = '0') AND (3.SLAM = 1.SLAM) AND
(3.STAR = 1.STAR) AND (3.SITE <> 1.SITE) AND (0.ALUM >= 3.ALUM)
AND (0.CODE <= 3.CODE)

Oracle® Rdb for OpenVMS

```

BgrNdx1 T2_IND_1 [3:4] Fan=7
  Keys: (3.TUNE = 1.TUNE) AND (3.SLAM = 1.SLAM) AND
        (3.STAR = 1.STAR) AND (0.ALUM >= 3.ALUM)
  Bool: (3.SLAM <= '31-DEC-2002') AND (3.TUNE = '240167-1447')
Cross block entry 4
  Index only retrieval of relation 2:T3
  Index name T3_IND [3:3]
  Keys: (2.TUNE = 0.TUNE) AND (2.SLAM = 0.SLAM) AND (2.STAR = 0.STAR)
T1.CODE          T2.ALUM          T1.ALUM
T2.CODE
22-FEB-2002 00:00:00.00    1-OCT-2001 00:00:00.00    22-FEB-2002 00:00:00.00
31-JAN-2002 00:00:00.00

22-FEB-2002 00:00:00.00    1-FEB-2002 00:00:00.00    22-FEB-2002 00:00:00.00
31-DEC-2003 00:00:00.00
2 rows selected

```

The predicate "T2.CODE IS NULL" is shared by two OR clauses referencing other different tables as in the following example.

```

the first predicate (T1.ALUM <= T2.CODE OR T2.CODE IS NULL)) references
table T1 and T2, and
the second predicate (T2.CODE >= '01-jan-2002' OR T2.CODE IS NULL)
references table T2.

```

However, in the detailed strategy display, the first predicate is missing under the cross block entry 2 for table "0:T1".

As a workaround, the query works if the SQL flag 'MAX_STABILITY' is defined OR the logical name RDMS\$MAX_STABILITY is defined.

```
set flags 'max_stability';
```

Tables:

```

0 = T1
1 = T2
2 = T3
3 = T2

```

Cross block of 4 entries

Cross block entry 1

```
Conjunct: (1.CODE >= '1-JAN-2002') OR MISSING (1.CODE)
```

```
Get Retrieval by index of relation 1:T2
```

```
Index name T2_IND_1 [1:2]
```

```
Keys: (1.TUNE = '240167-1447') AND (1.SLAM <= '31-DEC-2002')
```

Cross block entry 2

```
Conjunct: (1.SLAM = 0.ALUM) AND (1.STAR = 0.STAR) AND (0.CODE >= 1.ALUM)
```

```
AND ((0.ALUM <= 1.CODE) OR MISSING (1.CODE))
```

```
Get Retrieval by index of relation 0:T1
```

```
Index name T1_IND [1:1]
```

```
Keys: 1.TUNE = 0.TUNE
```

Cross block entry 3

```
Conjunct: <agg0> = 0
```

```
Aggregate-F1: 0:COUNT-ANY (<subselect>)
```

```
Conjunct: (3.MAKER = '0') AND (3.SITE <> 1.SITE) AND (0.CODE <= 3.CODE)
```

```
Get Retrieval by index of relation 3:T2
```

```
Index name T2_IND_1 [3:4]
```

```
Keys: (3.TUNE = 1.TUNE) AND (3.SLAM = 1.SLAM) AND
```

```
(3.STAR = 1.STAR) AND (0.ALUM >= 3.ALUM)
```

Oracle® Rdb for OpenVMS

```
      Bool: (3.SLAM <= '31-DEC-2002') AND (3.TUNE = '240167-1447')
Cross block entry 4
  Index only retrieval of relation 2:T3
  Index name  T3_IND [3:3]
  Keys: (2.TUNE = 0.TUNE) AND (2.ALUM = 0.ALUM) AND (2.STAR = 0.STAR)
T1.CODE          T2.ALUM          T1.ALUM
T2.CODE
22-FEB-2002 00:00:00.00    1-FEB-2002 00:00:00.00    22-FEB-2002 00:00:00.00
31-DEC-2003 00:00:00.00
```

1 row selected

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.29 Wrong Index Retrieval is Selected in Query with GTR Predicate

Bug 3144382

The wrong index retrieval is selected in a query with a GTR predicate.

```
set flags 'strategy,detail';

create index t_i on t(v1,v2,v3,v4,v5);

select * from t where
  v1='D' and v2 > 18 and
  v3 > 7 and v4 > 17 and v5 > 4
  order by v1,v2,v3,v4,v5;
Tables:
  0 = T
Conjunct: (0.V1 = 'D') AND (0.V2 > 18) AND (0.V3 > 7) AND (0.V4 > 17) AND (0.V5
  > 4)
Get      Retrieval by index of relation 0:T
  Index name  T_I [1:1]
  Keys: 0.V1 = 'D'
  Bool: (0.V2 > 18) AND (0.V3 > 7) AND (0.V4 > 17) AND (0.V5 > 4)
...etc...
180 rows selected
```

The query takes about 10.3 seconds (elapsed time) on a VAX machine.

In Oracle Rdb Release 7.0–62, the query runs fast (approximately 1 second) with the following strategy:

```
Tables:
  0 = T
Conjunct: ((0.V1 = <cvar>) AND (0.V2 >
<cvar>) AND (0.V3 > <cvar>)
  AND (0.V4 > <cvar>) AND (0.V5 > <cvar>))
Get      Retrieval by index of relation 0:T
  Index name  T_I [2:1] Bool
  Key: (0.V1 = <cvar>) AND (0.V2 > <cvar>)
  Bool: (0.V3 > <cvar>) AND (0.V4 > <cvar>) AND (0.V5 > <cvar>)
```

As a workaround, the query works if the SQL flag 'NOMAX_SOLUTION' is defined OR the logical name RDMS\$DISABLE_MAX_SOLUTION is defined.

```

set flags 'NOMAX_SOLUTION';

select * from t where
  v1='D' and v2 > 18 and
  v3 > 7 and v4 > 17 and v5 > 4
order by v1,v2,v3,v4,v5;
Tables:
  0 = T
Conjunct: (0.V1 = 'D') AND (0.V2 > 18) AND (0.V3 > 7) AND (0.V4 > 17) AND (0.V5
  > 4)
Get      Retrieval by index of relation 0:T
Index name  T_I [2:1]
Keys: (0.V1 = 'D') AND (0.V2 > 18)
Bool: (0.V3 > 7) AND (0.V4 > 17) AND (0.V5 > 4)

```

The query takes about 2.22 seconds (elapsed time) on a VAX machine.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.30 Memory Corruption When Using Explicit 2PC

Bug 3230612

It was possible for an application to experience memory corruption when the following conditions were true:

- Explicit two-phase commit (2PC) transactions were being utilized. That is, transactions that were started and ended by using the system services SYS\$START_TRANS and SYS\$END_TRANS.
- The application would:
 1. Detach from the database involved in the 2PC transaction
 2. Attach to a database
 3. Start a new 2PC transaction
- The application would allocate memory after detaching from the database and before attaching to a database.

This problem can be avoided by not using explicit two-phase commit or by remaining attached to all databases.

For more information regarding explicit two-phase commit transactions, see the Guide to Distributed Transactions.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.1.31 Query Applying Zero Shortcut Returns Wrong Results

Bug 3216607

The following query applying zero shortcut optimization should return 7 rows.

```

set flags 'strategy,detail';
SQL> select wip_number, plan_sequence_code
cont> from pv_errors
cont> where wip_number = 'PML_257224' and plan_sequence_code = 1;

```

Oracle® Rdb for OpenVMS

```
~S#0004
Leaf#01 FFirst PV_ERRORS Card=151726
  BgrNdx1 PV_ERRORS_HASH [1:1] Fan=1
  BgrNdx2 PV_ERRORS_SORTED_001 [2:2] Fan=13
~E#0004.01(1) Estim   Ndx:Lev/Seps/DBKeys 1:_6 2:1/0/0 ZeroShortcut
0 rows selected
```

A number of optimization changes have been made to the latest releases of Rdb to improve how the optimizer chooses which indexes should be used to help retrieve data.

One such improvement has been new index estimation procedures that give much more precise estimation of the effectiveness of the use of a particular index for retrieval.

This problem may occur when the optimizer incorrectly determines that one of the indices contains no records that could possibly satisfy the selection criteria and therefore immediately discontinues the processing of the query, delivering no records.

This is called 'zero shortcut' optimization which can be seen in the execution trace of the query above.

```
~E#0004.01(1) Estim   Ndx:Lev/Seps/DBKeys 1:_6 2:1/0/0 ZeroShortcut
```

As a workaround for the problem, the query returns the correct result by adding an ORDER BY clause.

```
SQL> select wip_number, plan_sequence_code
cont> from pv_errors
cont> where wip_number = 'PML_257224' and plan_sequence_code = 1
cont> order by wip_number;
```

WIP_NUMBER	PLAN_SEQUENCE_CODE
PML_257224	1
PML_257224	1
PML_257224	1
PML_257224	1
PML_257224	1
PML_257224	1
PML_257224	1

```
~S#0003
```

```
Leaf#01 Sorted PV_ERRORS Card=151726
  FgrNdx PV_ERRORS_HASH [1:1] Fan=1
  BgrNdx1 PV_ERRORS_SORTED_001 [2:2] Fan=13
WIP_NUMBER      PLAN_SEQUENCE_CODE
PML_257224      1
PML_257224      1
PML_257224      1
PML_257224      1
PML_257224      1
PML_257224      1
PML_257224      1
~E#0003.01(1) FgrNdx  Sorted  DBKeys=8  Fetches=0+0  RecsOut=7
~E#0003.01(1) FgrNdx  Sorted  DBKeys=8  Fetches=0+0  RecsOut=7
PML_257224      1
7 rows selected
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2 SQL Errors Fixed

3.2.1 Unexpected DATEEQLILL Error During IMPORT With CREATE INDEX or CREATE STORAGE MAP

Bug 1094071

When the SQL IMPORT statement includes CREATE STORAGE MAP or CREATE INDEX statements which use TIMESTAMP or DATE ANSI literals in the WITH LIMIT OF clauses, it fails with the following error:

```
%SQL-F-UNSDATXPR, Unsupported date expression  
-SQL-F-DATEEQLILL, Operands of date/time comparison are incorrect
```

The same CREATE STORAGE MAP or CREATE INDEX statements work correctly when used outside of the IMPORT statement.

This error is generated because the SQL IMPORT statement tries to validate the data type of the column against that of the literal value. However, during this phase of the IMPORT the table does not yet exist.

A workaround for this problem is to use DATE VMS literals in the WITH LIMIT OF clause and allow the Rdb Server to perform the data type conversion at runtime.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.2 Incorrect Unit for the DETECTED ASYNC PREFETCH THRESHOLD Option

Bug 2838771

In prior versions of Oracle Rdb V7.0, the THRESHOLD option of the DETECTED ASYNC PREFETCH clause required the units to be PAGES. However, this value is really specified in BUFFERS.

To avoid confusion, the SQL syntax has now been changed to use the BUFFERS unit for this clause. The older syntax is now deprecated as shown in the following example:

```
SQL> alter database  
cont>     filename 'DB$:PERSONNEL'  
cont>     detected async prefetch is ENABLED  
cont>     (depth is 4 buffers, threshold is 4 pages);  
%SQL-I-DEPR_FEATURE, Deprecated Feature:  
PAGES is replaced with BUFFERS
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1. In addition, the RMU Extract command will output the new corrected syntax.

```
SQL> alter database
```

```
cont> filename 'DB$:PERSONNEL'
cont> detected async prefetch is ENABLED
cont> (depth is 4 buffers, threshold is 4 buffers);
```

3.2.3 DECLARE LOCAL TEMPORARY TABLE Limited to 10 Tables Per Session

Bug 2911428

In prior releases of Oracle Rdb, access to tables declared using the DECLARE LOCAL TEMPORARY TABLE statement in interactive and dynamic SQL would fail. This problem does not occur when DECLARE LOCAL TEMPORARY TABLE is used in a CREATE MODULE statement.

The following example shows the reported error.

```
SQL> select * from demo.MODULE.prodn_new_constraints_table;
%RDB-E-OBSOLETE_METADA, request references metadata objects that no longer exist
-RDMS-F-TABIDNOTDEF, relation ID, 11, is not defined in database
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1. This limitation has been removed from interactive and dynamic SQL. A workaround would be to declare tables #10 and #11 and not use those temporary tables.

3.2.4 IVP or Other Failure with Dynamic SQL if SQL\$INT is Installed /RESIDENT

Bug 2950983

In SYS\$STARTUP:SQL\$STARTUP.COM, if the line RESIDENT = "/RESIDENT" was present, (for example, not commented out), the IVP failed while running the dynamic SQL test.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.5 Bugcheck at PSIINDEX\$FIND_ENTS_EXACT + 54

Bug 2448304

When a combination of metadata operations was being performed in the same transaction, it was possible that an OpenVMS Alpha bugcheck could be generated with an exception similar to the following example.

```
***** Exception at 0106BC24 : PSIINDEX$FIND_ENTS_EXACT + 00000054
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
  virtual address=0000000050000038, PC=000000000106BC24, PS=0000000B
```

This problem was caused by an incorrect use of internal memory management optimizations. This incorrect use can be identified by the unusual virtual address of 0000000050000038.

The problem can be avoided by reducing the number of metadata operations performed in a single transaction.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.6 Multistatement Procedures Used with Connections Resulted in %RDB-E-OBSOLETE_METADA Error Message

Bugs 1879521 and 1808821

In prior releases of Oracle Rdb, there was a problem with multiple connections and the use of multistatement procedures. Specifically, Oracle Rdb requires a special internal module to be set up for multistatement procedures. In the case of two or more connections calling the same multistatement procedure, the module setup was not done for the second connection. This was incorrect behavior and resulted in the following error message:

```
%RDB-E-OBSOLETE_METADA, request references metadata objects that no longer exist
```

The correct behaviour is to insure that the module setup is performed when a database switch occurs for the first time.

Note: this problem was originally reported as fixed in Release 7.0.6.4 but the fix was inadvertently left out of that release.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.7 IMPORT May Generate ACCVIO Exception During Import of a Module

In previous releases of Oracle Rdb, the IMPORT command would fail with an ACCVIO exception during import of a module that contained an external function or procedure.

The following example shows the exception.

```
SQL> import database
cont>   from SQL_CREATE_MODULE_4G_EXP
cont>   filename SQL_CREATE_MODULE_4G_DB
cont> ;
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
virtual address=00000000, PC=0029908E, PSL=03C00001
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1. SQL IMPORT now correctly handles these types of modules.

3.2.8 Unexpected ACCVIO When Reporting Incompatible Character Set Assignments

Bug 3098432

In prior releases of Rdb, an incompatible character set assignment might cause Rdb to generate a malformed message vector which can lead to an ACCVIO while trying to display the message.

The following example shows the problem.

```
SQL> select count(*) from rdb$database
cont> where _dec_mcs'a'=translate('a' using rdb$dec_kanji);
%RDB-E, invalid or unsupported data conversion
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual address=
00000000004F2DE4, PC=FFFFFFFF80207624, PS=0000001B
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.9 SQL-F-NODBFILe When SQL Modules are Compiled With /CONNECT

Bug 3002999

When multiple SQL modules were called from the same main program and a disconnect was done, under some circumstances calls to SQL modules after the disconnect would fail with the following message:

```
%SQL-F-NODBFILe, Alias <ALIAS_NAME> is missing a declaration
```

where <ALIAS_NAME> is one of the aliases declared by one of the SQL modules. The behaviour will occur whenever the following is true:

- ◆ The SQL modules are compiled with /CONNECT (which is the default).
- ◆ Some aliases have a run time resolution for the filename or pathname.
- ◆ One or more modules do not declare one of the aliases which has a run-time resolution.
- ◆ The first call after the disconnect is to a module which does not have a declaration of an alias with a run-time resolution.

As a workaround, ensure that the first SQL module called after the disconnect declares all aliases which have a run-time resolution.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.10 GET DIAGNOSTICS Keyword CONNECTION_NAME Returned Incorrect Value

Bug 880810

In previous versions of Oracle Rdb, the GET DIAGNOSTICS keyword CONNECTION_NAME was not always returning the correct value. The following example shows the problem. The result 'RDB\$DEFAULT_CONNECTION' was not expected but rather 'MF'.

```
SQL> declare :c char(31);
SQL>
SQL> attach 'alias aa filename sql$database';
SQL> connect to 'alias aa filename db$:mf_personnel' as 'MF';
SQL>
```

```
SQL> set connect 'MF';
SQL> begin
cont> get diagnostic :c = connection_name;
cont> end;
SQL> print :c;
      C
      RDB$DEFAULT_CONNECTION
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1. CONNECT now correctly passes the connection name to the Rdb server so that it can be returned by GET DIAGNOSTICS.

3.2.11 Errors Not Reported by SQL

Bug 3083552

In some cases, when a bugcheck was produced, the SQLCODE returned was 0 indicating that no error had occurred.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.12 Variable Updated Though No Rows Found

Bug 2863676

Variables in SQL statements were updated even though the statement did not result in any data (for example, no rows selected). This occurred with declared variables in interactive SQL as well as host variables in application programs. The problem is illustrated by the following interactive SQL example.

```
SQL> at 'f personnel';
SQL> declare :my_int integer;
SQL> begin set :my_int = 0; end;
SQL> select 1 from rdb$database where exists
cont> (select * from employees where employee_id = '99999');
0 rows selected
SQL> select 1 into :my_int from rdb$database where exists
cont> (select * from employees where employee_id = '99999');
SQL> print :my_int;
      MY_INT
      1
1 row selected
```

In the above example, the variable ":my_int" should have remained zero since no rows were selected.

As a workaround, the application must examine the SQLSTATE or SQLCODE after each such statement and restore the original value of the variable as needed.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.2.13 Partitioning Clause Not Working in Embedded SQL

Bug 3112810

In some cases, a sequential scan on partitioned data where the partitioning uses a VARCHAR type column can fail (or possibly produce wrong data results).

The following example shows a typical error message:

```
char a_alert[6] = "XYZ05";

EXEC SQL
  SET TRANSACTION READ WRITE WAIT RESERVING
    T_LIMIT PARTITION (3) FOR EXCLUSIVE READ;

EXEC SQL
  SELECT count(*)
  into   :i
  FROM   T_LIMIT
  WHERE  A_ALERT = :a_alert
        and (
          EXTRACT (YEAR FROM A_DATE_CREATED) * 10000 +
          EXTRACT (MONTH FROM A_DATE_CREATED) * 100 +
          EXTRACT (DAY FROM A_DATE_CREATED)
        ) = 20031017;

%RDB-E-UNRES_REL, relation T_LIMIT in specified request is not a relation
reserved in specified transaction
-RDMS-E-UNRES_AREA, area 67 within relation "T_LIMIT" is not reserved
```

Rdb is expecting the type of the parameter in the WHERE clause to agree with the type of the column it is comparing against (which is CHAR(5) in this case).

Interactive SQL does the data conversion and comparison correctly but SQL\$PRE/CC sometimes gets confused.

Possible workarounds include:

1. Use a compound statement to convert the type from VARCHAR to CHAR.
Here we recode the example query (above) as a compound statement which assigns the input parameter into an internally declared CHAR(5) variable as follows:

```
EXEC SQL
  BEGIN declare :my_char char(5);
  set :my_char = :a_alert;
  SELECT count(*)
  into   :i
  FROM   T_LIMIT
  WHERE  A_ALERT = :my_char
        and (
          EXTRACT (YEAR FROM A_DATE_CREATED) * 10000 +
          EXTRACT (MONTH FROM A_DATE_CREATED) * 100 +
          EXTRACT (DAY FROM A_DATE_CREATED)
        ) = 20031017; END;
```

The above statement uses the "MY_CHAR" variable to convert the type of the input from VARCHAR (which is the only way that SQL\$PRE/CC will send it) into CHAR.

2. Use Dynamic SQL to avoid host variable inputs.

For example, in the reproducer program declare SQLDA areas as follows:

```
EXEC SQL INCLUDE SQLDA;
...
struct SQLDA_STRUCT *OUT_AREA;
OUT_AREA = calloc(1, sizeof(struct SQLDA_STRUCT) +
                 (10*sizeof(struct SQLVAR_STRUCT)) );
OUT_AREA->SQLN = 10;
```

Then recode the query as a dynamic statement with no input parameters like this:

```
sprintf(sql_stmt,
        "SELECT count(*) \
        into ? \
        FROM T_LIMIT \
        WHERE A_ALERT = '%s' \
        and ( \
            EXTRACT (YEAR FROM A_DATE_CREATED) * 10000 + \
            EXTRACT (MONTH FROM A_DATE_CREATED) * 100 + \
            EXTRACT (DAY FROM A_DATE_CREATED) \
            ) = 20031017", a_alert);
EXEC SQL PREPARE stmt1 from :sql_stmt;
EXEC SQL DESCRIBE stmt1 OUTPUT into :OUT_AREA;
OUT_AREA->SQLVAR[0].SQLDATA = (char *)&i;
EXEC SQL EXECUTE stmt1 into descriptor :OUT_AREA;
```

3. Use Dynamic SQL to force host variable parameter inputs to CHAR in an input SQLDA. This method has the advantage that the request can be reused. However, it is imperative that the input strings be space padded to the length of the data (five in our examples).

For example, declare the SQLDA items as in the Workaround 2 example above plus the following:

```
struct SQLDA_STRUCT *IN_AREA;
IN_AREA = calloc(1, sizeof(struct SQLDA_STRUCT) +
                 (10*sizeof(struct SQLVAR_STRUCT)) );
IN_AREA->SQLN = 10;
```

Then recode the query to use an input SQLDA as follows:

```
strcpy(sql_stmt, "SELECT count(*) \
into ? \
FROM T_LIMIT \
WHERE A_ALERT = ? \
and ( \
    EXTRACT (YEAR FROM A_DATE_CREATED) * 10000 + \
    EXTRACT (MONTH FROM A_DATE_CREATED) * 100 + \
    EXTRACT (DAY FROM A_DATE_CREATED) \
    ) = 20031017");

EXEC SQL PREPARE stmt2 from :sql_stmt;
EXEC SQL DESCRIBE stmt2 OUTPUT into :OUT_AREA;
EXEC SQL DESCRIBE stmt2 INPUT into :IN_AREA;
```

Oracle® Rdb for OpenVMS

```
IN_AREA->SQLVAR[0].SQLTYPE = SQLDA_CHAR;
IN_AREA->SQLVAR[0].SQLDATA = a_alert;
OUT_AREA->SQLVAR[0].SQLDATA = (char *)&i;

EXEC SQL EXECUTE stmt2 into descriptor :OUT_AREA
      using descriptor :IN_AREA;
```

In this example, the second DESCRIBE sets *IN_AREA->SQLVAR[0].SQLTYPE* to *SQLDA_VARCHAR*. But overriding this value to *SQLDA_CHAR* causes dynamic SQL to build a request that works. Again please note that when using *SQLDA_CHAR*, the data string (*a_alert* in the example) must be space padded to the length specified in *IN_AREA->SQLVAR[0].SQLLEN*.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.3 RDO and RDML Errors Fixed

3.3.1 RDML /DATE_TYPE Qualifier Default is Now NOEMPTY_RECORDS

RDML /PASCAL generates a record type for date items. This type is either an empty record or a record composed of two longword elements depending on the value of the /DATE_TYPE qualifier (EMPTY_RECORDS or NOEMPTY_RECORDS respectively). Earlier versions of the Pascal compiler behaved as desired when empty records were used in assignments and fetches. More current versions of the Pascal compiler give the following warning message when empty records are used for dates:

```
%PASCAL-W-EMPTYVAR, Fetching an empty record with an explicit size  
attribute may not yield expected results
```

Furthermore, statements flagged with the above warning frequently do not yield the desired result but instead treat the record as having a zero length. Using the /DATE_TYPE=NOEMPTY_RECORD qualifier avoids this problem.

In order to make RDML /PASCAL more usable, NOEMPTY_RECORDS has been made the default value for the DATE_TYPE. Additionally, using an explicit /DATE_TYPE=EMPTY_RECORDS on the RDML /PASCAL command line will now result in the following warning message being issued by RDML:

```
%RDML-W-DATE_NOEMPTY, /DATE_TYPE=EMPTY_RECORDS specified;  
use /DATE_TYPE=NOEMPTY_RECORDS
```

As a workaround, use /DATE_TYPE=NOEMPTY_RECORDS.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.4 Oracle RMU Errors Fixed

3.4.1 RMU /COLLECT May Default to a READ WRITE Transaction

Bug 2898115

When no transaction type is specified for the *RMU/COLLECT OPTIMIZER_STATISTICS* command, RMU examines the database storage areas and if any storage area has snapshots disabled, a read write transaction is used. If no storage areas have snapshots disabled, then a read only transaction is used.

Currently Rdb does not allow different areas on the same database to have snapshots enabled or disabled. However, the attribute is recorded in the database root file against each storage area.

In previous versions, if a database had reserved and unused storage area slots, this check could erroneously examine the unused storage area slots and conclude that areas on the database had snapshots disabled. This could cause RMU to adopt a read write transaction as the default transaction mode for *RMU/COLLECT*.

The following example shows *RMU/COLLECT* executing against a database with reserved storage areas erroneously selecting a read write transaction. The example also shows the use of debug flags to display the transaction modes used by the command.

```
$ DEFINE RDMS$SET_FLAGS TRANSACTION
$ RMU /COLLECT OPTIMIZER MF_PERSONNEL
~T Compile transaction (1) on db: 1
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_WRITE (read write)
~T Start_transaction (1) on db: 1, db count=1
~T Commit_transaction on db: 1
~T Prepare_transaction on db: 1
~T Compile transaction (1) on db: 2
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_READ (read only)
~T Start_transaction (1) on db: 2, db count=1
~T Commit_transaction on db: 2
~T Prepare_transaction on db: 2
~T Compile transaction (1) on db: 3
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_WRITE (read write)
~T Start_transaction (1) on db: 3, db count=1
~T Commit_transaction on db: 3
~T Prepare_transaction on db: 3
~T Compile transaction (1) on db: 4
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_WRITE (read write)
~T Start_transaction (1) on db: 4, db count=1
~T Commit_transaction on db: 4
~T Prepare_transaction on db: 4
```

The following example shows the corrected behaviour.

```

$ RMU /COLLECT OPTIMIZER MF_PERSONNEL /TRANS=READ_ONLY
~T Compile transaction (1) on db: 1
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_READ (read only)
~T Start_transaction (1) on db: 1, db count=1
~T Commit_transaction on db: 1
~T Prepare_transaction on db: 1
~T Compile transaction (1) on db: 2
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_READ (read only)
~T Start_transaction (1) on db: 2, db count=1
~T Commit_transaction on db: 2
~T Prepare_transaction on db: 2
~T Compile transaction (1) on db: 3
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_READ (read only)
~T Start_transaction (1) on db: 3, db count=1
~T Commit_transaction on db: 3
~T Prepare_transaction on db: 3
~T Compile transaction (1) on db: 4
~T Transaction Parameter Block: (len=2)
0000 (00000) TPB$K_VERSION = 1
0001 (00001) TPB$K_WRITE (read write)
~T Start_transaction (1) on db: 4, db count=1
~T Commit_transaction on db: 4
~T Prepare_transaction on db: 4

```

Note that the read write transaction at the end of the example is used to update information in the metadata and is normal and required.

Specifying the transaction type on the command line by using the */TRANSACTION_TYPE=READ_ONLY* qualifier will avoid this problem.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.4.2 RMU/VERIFY/CONSTRAINTS Problems With Named Tables And Constraints

Bug 3016789

Oracle Rdb V7.0 RMU/VERIFY/CONSTRAINTS verified all constraints even if a list of constraints was specified. Oracle Rdb V7.1 RMU/VERIFY/CONSTRAINTS verified all constraints for all tables even if a list of tables was specified. This behaviour has been corrected so that only constraints for a specified list of tables or only the constraints specified will be verified.

The following example shows that for Oracle Rdb V7.0 RMU/VERIFY/CONSTRAINTS, all constraints were verified even if a list of constraints was specified and that for Oracle Rdb V7.1 RMU/VERIFY/CONSTRAINTS, all constraints for all tables were verified even if a list of tables was specified.

Oracle® Rdb for OpenVMS

```
$DEFINE RDMS$DEBUG_FLAGS "H"
$ rmu/show version
Executing RMU for Oracle Rdb V7.0-7
$ rmu/verify/constraint=(constraint=degrees_foreign2)/log mf_personnel
%RMU-I-BGNROOVER, beginning root verification
%RMU-I-ENDROOVER, completed root verification
%RMU-I-BGNVCONST, beginning verification of constraints for database
device:[directory]MF_PERSONNEL.RDB;1
~H Extension (VERIFY CONSTRAINTS) Item List: (len=0)
~H: ..verify constraint "WORK_STATUS_PRIMARY_STATUS_CODE"
~H: ..verify constraint "STATUS_NAME_VALUES"
~H: ..verify constraint "STATUS_TYPE_VALUES"
~H: ..verify constraint "EMPLOYEES_PRIMARY_EMPLOYEE_ID"
~H: ..verify constraint "EMP_SEX_VALUES"
~H: ..verify constraint "EMP_STATUS_CODE_VALUES"
~H: ..verify constraint "JOBS_PRIMARY_JOB_CODE"
~H: ..verify constraint "WAGE_CLASS_VALUES"
~H: ..verify constraint "DEPARTMENTS_PRIMARY1"
~H: ..verify constraint "JOB_HISTORY_FOREIGN1"
~H: ..verify constraint "JOB_HISTORY_FOREIGN2"
~H: ..verify constraint "JOB_HISTORY_FOREIGN3"
~H: ..verify constraint "SALARY_HISTORY_FOREIGN1"
~H: ..verify constraint "COLLEGES_PRIMARY_COLLEGE_CODE"
~H: ..verify constraint "DEGREES_FOREIGN1"
~H: ..verify constraint "DEGREES_FOREIGN2"
~H: ..verify constraint "DEG_DEGREE_VALUES"
~H: ..verify constraint "CANDIDATES_LAST_NAME_NOT_NULL"
~H: ..verify constraint "RESUMES_UNIQUE_EMPLOYEE_ID"
~H: ..verify constraint "RESUMES_FOREIGN1"
%RMU-I-ENDVCONST, completed verification of constraints for database
device:[directory]MF_PERSONNEL.RDB;1

$ rmu/show version
Executing RMU for Oracle Rdb V7.1-10
$ rmu/verify/constraint=(table=colleges)/log mf_personnel
%RMU-I-BGNROOVER, beginning root verification
%RMU-I-ENDROOVER, completed root verification
%RMU-I-BGNVCONST, beginning verification of constraints for database
device:[directory]MF_PERSONNEL.RDB;1
~H Extension (VERIFY CONSTRAINTS) Item List: (len=0)
~H: ..verify constraint "WORK_STATUS_PRIMARY_STATUS_CODE"
~H: ..verify constraint "STATUS_NAME_VALUES"
~H: ..verify constraint "STATUS_TYPE_VALUES"
~H: ..verify constraint "EMPLOYEES_PRIMARY_EMPLOYEE_ID"
~H: ..verify constraint "EMP_SEX_VALUES"
~H: ..verify constraint "EMP_STATUS_CODE_VALUES"
~H: ..verify constraint "JOBS_PRIMARY_JOB_CODE"
~H: ..verify constraint "WAGE_CLASS_VALUES"
~H: ..verify constraint "DEPARTMENTS_PRIMARY1"
~H: ..verify constraint "JOB_HISTORY_FOREIGN3"
~H: ..verify constraint "JOB_HISTORY_FOREIGN1"
~H: ..verify constraint "JOB_HISTORY_FOREIGN2"
~H: ..verify constraint "COLLEGES_PRIMARY_COLLEGE_CODE"
~H: ..verify constraint "DEGREES_FOREIGN1"
~H: ..verify constraint "DEGREES_FOREIGN2"
~H: ..verify constraint "DEG_DEGREE_VALUES"
~H: ..verify constraint "CANDIDATES_LAST_NAME_NOT_NULL"
~H: ..verify constraint "RESUMES_UNIQUE_EMPLOYEE_ID"
~H: ..verify constraint "RESUMES_FOREIGN1"
%RMU-I-ENDVCONST, completed verification of constraints for database
device:[directory]MF_PERSONNEL.RDB;1
```

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.4.3 RMU Unload Incorrectly Using DBKEY SCOPE IS ATTACH

Bug 2623790

In prior releases of Oracle Rdb V7.0, it was possible that RMU Unload would default to DBKEY SCOPE IS ATTACH. This problem has been corrected in Oracle Rdb Release 7.0.7.1.

This problem occurs when applications attach to the Rdb database without a database parameter block (DPB). This can be determined by defining RDMS\$SET_FLAGS "DATABASE" prior to running the application or RMU command.

A typical attach to the database will display lines describing the explicitly set "Database Parameter Buffer" including the DBKEY SCOPE.

```
$ SQL$
SQL> attach 'filename SCRATCH';
ATTACH #1, Database DISK3:[DATABASES.V71]SCRATCH.RDB;1
~P Database Parameter Buffer (version=2, len=78)
0000 (00000) RDB$K_DPB_VERSION2
0001 (00001) RDB$K_FACILITY_ALL
0002 (00002) RDB$K_DPB2_IMAGE_NAME "NODE::DISK1:[DIR]SQL$71.EXE;1"
003F (00063) RDB$K_FACILITY_ALL
0040 (00064) RDB$K_DPB2_DBKEY_SCOPE (Transaction)
0044 (00068) RDB$K_FACILITY_ALL
0045 (00069) RDB$K_DPB2_REQUEST_SCOPE (Attach)
0049 (00073) RDB$K_FACILITY_RDB_VMS
004A (00074) RDB$K_DPB2_CDD_MAINTAINED (No)
RDMS$BIND_WORK_FILE = "DISK2:[TEST]RDMSTTBL$KMCMLSDFBXX.TMP;" (Visible = 0)
```

In the case of RMU Unload, this buffer is absent and defaults should be used. Unfortunately, the DBKEY SCOPE is undefined and is being incorrectly set to DBKEY SCOPE IS ATTACH.

```
$ RMU /Unload SCRATCH
ATTACH #1, Database DISK3:[DATABASES.V71]SCRATCH.RDB;1
RDMS$BIND_WORK_FILE = "DISK2:[TEST]RDMSTTBL$KMY10MNHXX.TMP;" (Visible = 0)
```

It is possible that other applications or 4GL tools suffer from this problem also. It can be confirmed using RDMS\$SET_FLAGS as shown here, or using the RMU/SHOW STATISTICS command to see if any process is waiting for the lock "waiting for database key scope" in the *Active User Stall Messages screen*.

3.5 LogMiner Errors Fixed

3.5.1 LogMiner Elimination of Processing Unneeded AIJ Files

By default, after-image journal files are processed in the order that they are presented to the RMU `UNLOAD AFTER_JOURNAL` command. The `ORDER_AIJ_FILES` qualifier specifies that the input after-image journal files are to be processed in ascending order by sequence number. This can be of benefit when you use wildcard (* or %) processing of a number of input files. The .AIJ files are each opened, the first block is read (to determine the sequence number), and the files are closed prior to the sorting operation.

The `/RESTART=restart-point` qualifier can be used to specify an AIJ Extract Restart Control Point (AERCP) that indicates the location to begin the extraction. The AERCP indicates the transaction sequence number (TSN) of the last extracted transaction along with a location in the .AIJ file where a known "Micro-quiet point" exists.

When the Restart qualifier is not specified and no input after-image journal files are specified on the command line, the Continuous LogMiner process starts extracting at the beginning of the earliest modified online after-image journal file.

Previously, all specified input after-image journal files were always processed. This behaviour has been altered when both the `RESTART` and `ORDER_AIJ_FILES` qualifiers are specified. In this situation, any after-image journal files containing only sequence numbers prior to the "Micro-quiet point" sequence number (indicated in the AIJ Extract Restart Control Point (AERCP)) can be eliminated from processing after the order of sequence numbers has been determined by the sort operation. Eliminating the unneeded after-image journal files can speed the restart operation.

3.5.2 `/TRANSACTION_TYPE` Qualifier for RMU `/UNLOAD /AFTER_JOURNAL`

The RMU `/UNLOAD /AFTER_JOURNAL` command would start either a read-write or a read-only transaction to read the database metadata. A read-only transaction would be started if the database was set to "SNAPSHOTS ARE IMMEDIATE"; otherwise a read-write transaction would be started.

The qualifier `"/TRANSACTION_TYPE="` has been added to the RMU `/UNLOAD /AFTER_JOURNAL` command to allow explicit control over the transaction type used when reading the database metadata.

The following keywords may be specified to the `"/TRANSACTION_TYPE="` qualifier.

- ◆ `AUTOMATIC` – The transaction type will depend upon the current database settings for snapshots (enabled, deferred, or disabled), transaction modes available to this user, and the standby status of this database.
- ◆ `READ_ONLY` – Starts a `READ ONLY` transaction.
- ◆ `WRITE` – Starts a `READ WRITE` transaction.
- ◆ `ISOLATION_LEVEL` – The transaction isolation level. It accepts the following keywords:

- ◇ READ_COMMITTED
- ◇ REPEATABLE_READ
- ◇ SERIALIZABLE

Please refer to the Oracle Rdb7 SQL Reference Manual under the SET TRANSACTION statement for a complete description of these isolation levels.

- ◆ WAIT – Will wait indefinitely on a locked resource.
- ◆ WAIT=n – This instructs Rdb to wait 'n' seconds before aborting the wait and the RMU session. Specifying a wait timeout interval of zero (0) is equivalent to specifying NOWAIT.
- ◆ NOWAIT – Will not wait on locked resources.

If the qualifier /TRANSACTION_TYPE is omitted or specified with no options, then the default is /TRANSACTION_TYPE=(AUTOMATIC, WAIT).

Although a WRITE transaction is started on the database, RMU /UNLOAD /AFTER_JOURNAL does not attempt to write to the database tables.

3.6 RMU Show Statistics Errors Fixed

3.6.1 RMU /SHOW STATISTICS Writes Invalid Configuration File

Bug 3108571

Starting in Oracle Rdb Release 7.0.6.3, the RMU /SHOW STATISTICS utility could write an invalid line to a save configuration file. In particular, the "CHECKPOINT_TX" line would be omitted and the line containing "CHECKPOINT_BLOCK_COUNT" would be corrupted.

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

3.7 Hot Standby Errors Fixed

3.7.1 Starting LRS on Master Database Caused Shutdown

Bug 1775983

If a database administrator attempted to start the Hot Standby feature and mistakenly specified the master database as the standby database then the master database was shutdown.

For example, note that in the following command the /MASTER qualifier is being used against the master database. This causes Hot Standby to attempt to configure the master database as a standby database.

```
$ RMU /REPLICATE AFTER_JOURNAL START [.MASTER]mf_personnel -  
  /MASTER_ROOT=[.STANDBY]standby_personnel
```

When the above command was issued, the AIJ Log Roll-forward Server (LRS) failed with the following error:

```
RDMS-F-STBYDBINUSE, standby database cannot be  
exclusively accessed for replication
```

Whenever the LRS failed, the database was automatically shutdown. A running application could be accidentally shutdown due to this type of command error.

This problem has been corrected in Oracle Rdb Release 7.0.7.1. Now, if the LRS fails with a STBYDBINUSE error, the database is not shutdown.

Chapter 4 Enhancements

4.1 Enhancements Provided in Oracle Rdb Release 7.0.7.2

4.1.1 Rdb Optional Site-Specific Startup Procedure

The Oracle Rdb startup procedure RMONSTART(xx).COM now supports an optional site-specific startup procedure to be executed after the Rdb Monitor (RDMMON) process has been started. If the file SYS\$STARTUP:RDB\$SYSTARTUP(xx).COM (where xx indicates the version number for multi-version Rdb kits) is found, it is executed as a DCL command procedure by the RMONSTART(xx).COM procedure.

SYS\$STARTUP:RDB\$SYSTARTUP(xx).COM is intended to contain site-specific tasks to be executed after the Rdb monitor procedure has completed. Such tasks might include opening databases or starting layered products that depend on the Rdb monitor process having been started.

If a site wishes to use this capability, the RDB\$SYSTARTUP(xx).COM procedure must be created in SYS\$STARTUP (either in the common SYS\$COMMON:[SYS\$STARTUP] directory or a node-specific SYS\$SPECIFIC:[SYS\$STARTUP] directory). The Rdb installation procedure does not provide or replace this file.

4.1.2 Oracle Rdb SGA API

Oracle Rdb maintains an extensive set of online performance statistics that provide valuable dynamic information regarding the status of an active database. The system global area (SGA) application programming interface (API) described in this document provides a way to retrieve these database performance statistics.

The SGA API automates retrieving database statistics available only through the RMU Show Statistics command. The SGA API provides the only way to retrieve statistics for Oracle Rdb databases from an application. Using the SGA API provides fast access to the data without effecting the execution of the server.

Previously, the Oracle Rdb SGA API was available as a separate software option to be downloaded, installed and maintained independently of the Oracle Rdb kit. Each time a new version of Oracle Rdb was installed, the SGA API would have to be updated. If the SGA API was not updated, it would, in many cases, fail to work correctly.

This problem has been partly resolved. Most of the contents of the SGA API separate software option are now automatically provided in the RDM\$DEMO directory during the Oracle Rdb kit installation procedure. Please refer to the SGA API documentation available in RDM\$DEMO in various formats as SGAAPI.PS, SGAAPI.HTML and SGAAPI.TXT for additional information.

Existing Users of the SGA API May Have to Modify Procedures

Existing users of the Oracle Rdb SGA API should refer to the documentation as there will be some minor changes required. In particular, the KUSRMUSHRxx.EXE sharable image is now provided in SYS\$LIBRARY and the

KUSRMUSHRxx.OPT linker options file has been updated to reference this sharable image in its new location.

4.1.3 CHRONO_FLAG Replaces Older CRONO_FLAG Keyword

The SET FLAGS statement and the RDMS\$SET_FLAGS logical now accept the keyword CHRONO_FLAG as an alternate spelling of the existing keyword CRONO_FLAG. The latter keyword is deprecated and will be removed from a future release of Rdb V7.1.

The new keyword will also be displayed by the SHOW FLAGS statement.

4.2 Enhancements Provided in Oracle Rdb Release 7.0.7.1

4.2.1 RDM\$BIND_SNAP_QUIET_POINT Logical No Longer Used

Bug 2656534

If the logical RDM\$BIND_SNAP_QUIET_POINT was defined to be "0" on a system that was used for the standby database in a Hot Standby configuration, it was not possible to start database replication. Attempts to start replication would fail with:

```
RDMS-F-HOTSNAPQUIET, quiet points must be enabled
for snapshot transactions during hot standby replication
```

However, defining the logical to "1" can cause processes with long running READ ONLY transactions to prevent database backups from proceeding.

This logical was introduced in Oracle Rdb Release 6.0 to allow database administrators to override the 6.0 requirement that READ ONLY transactions hold the quiet-point lock. Defining the logical to "1" (the default) would provide better performance when the Fast Commit feature was enabled and processes frequently switched between READ ONLY and READ WRITE transactions. Since that time, improvements have been made to the quiet-point lock algorithms that make this logical no longer necessary. Since releases 7.0.6.3 and 7.1.0.1, READ ONLY transactions would continue to hold the quiet-point lock until a backup process requested the lock. When the lock was requested, READ ONLY processes would release the quiet point lock as soon as the currently executing database request finished, if the RDM\$BIND_SNAP_QUIET_POINT logical was defined as "0". This made it no longer necessary to have the logical defined as "1".

Since the quiet-point lock behavior now behaves optimally, even with the Fast Commit feature enabled, the RDM\$BIND_SNAP_QUIET_POINT logical is no longer needed and thus has been removed. Oracle Rdb will now behave as if the logical is always defined to be "0".

This problem has been corrected in Oracle Rdb Release 7.0.7.1.

4.2.2 Determining Which Oracle Rdb Options Are Installed

When installing Oracle Rdb Server on OpenVMS you can choose from five components to install:

1. Oracle Rdb
2. Programmer for Rdb (Rdb Compilers)
3. Hot Standby
4. Power Utilities
5. Common Components

Starting with Rdb 7.0, you can determine what Rdb options were selected during the installation of Rdb by running the program SYS\$SYSTEM:RDBINS<RdbVersionVariant>.EXE. For example:

```
$ RUN SYS$SYSTEM:RDBINS70
```

Oracle® Rdb for OpenVMS

Installed: Oracle Rdb,Rdb Compilers,Hot Standby,Power Utilities

Previously, however, the output of the RDBINS program could not easily be redirected. Attempts to redefine SYSS\$OUTPUT would not allow the program output to be captured.

This problem has been resolved. The RDBINS program now allows redirection of the output to SYSS\$OUTPUT. The RDBINS program also creates a DCL symbol RDB\$INSTALLED_SELECTIONS containing the same output string as is displayed to SYSS\$OUTPUT.

4.2.3 New Procedure RDB\$IMAGE_VERSIONS.COM

The command procedure RDB\$IMAGE_VERSIONS.COM is supplied in SYSS\$LIBRARY by the Rdb installation procedure. The RDB\$IMAGE_VERSIONS command procedure can be used to display the image identification string and image link date/time from various Oracle Rdb or potentially related images in SYSS\$SYSTEM, SYSS\$LIBRARY and SYSS\$MESSAGE. This procedure can be used to determine exactly what images are installed on the system.

RDB\$IMAGE_VERSIONS.COM accepts an optional parameter. If passed, this parameter specifies a specific file or wildcard to lookup and display information for. By default, filenames starting with RD*, SQL*, RM*, and COSI* and ending with .EXE are searched for and displayed.

The following example shows how to use the RDB\$IMAGE_VERSIONS command procedure.

```
Decrdb RTA1:> @RDB$IMAGE_VERSIONS
SYS$SYSROOT:[SYSEXE]RDB$NATCONN71.EXE;1 SQL*NET V7.1-55 8-MAY-2002 15:56
SYS$COMMON:[SYSEXE]RDBINS.EXE;4 ORACLE RDB V7.0 14-NOV-2002 17:20
SYS$COMMON:[SYSEXE]RDBINS70.EXE;33 ORACLE RDB V7.0 7-MAR-2003 15:30
SYS$COMMON:[SYSEXE]RDBINS71.EXE;5 ORACLE RDB V7.1 9-APR-2003 10:58
SYS$COMMON:[SYSEXE]RDBPRE.EXE;5 V7.0-65 10-SEP-2002 16:02
SYS$COMMON:[SYSEXE]RDBPRE70.EXE;37 V7.0-7 28-FEB-2003 23:24
SYS$COMMON:[SYSEXE]RDBPRE71.EXE;5 V7.1-101 8-APR-2003 16:49
SYS$COMMON:[SYSEXE]RDBSERVER.EXE;9 RDB/RSV V7.0-65 5-SEP-2002 21:01
SYS$COMMON:[SYSEXE]RDBSERVER70.EXE;41 RDB/RSV V7.0-7 27-FEB-2003 17:29
SYS$COMMON:[SYSEXE]RDBSERVER71.EXE;5 RDB/RSVV7.1-101 7-APR-2003 17:43
SYS$COMMON:[SYSEXE]RDMABS.EXE;5 RDB V7.0-65 10-SEP-2002 16:01
.
.
.
```

Chapter 5

Documentation Corrections

This chapter provides information not currently available in the Oracle Rdb documentation set.

5.1 Documentation Corrections

5.1.1 Database Server Process Priority Clarification

By default, the database servers (ABS, ALS, DBR, LCS, LRS, RCS) created by the Rdb monitor inherit their VMS process scheduling base priority from the Rdb monitor process. The default priority for the Rdb monitor process is 15.

Individual server priorities can be explicitly controlled via system-wide logical names as described in [Table 5-1](#).

Table 5-1 Server Process Priority Logical Names

Logical Name	Use
RDM\$BIND_ABS_PRIORITY	Base Priority for the ABS Server process
RDM\$BIND_ALS_PRIORITY	Base Priority for the ALS Server process
RDM\$BIND_DBR_PRIORITY	Base Priority for the DBR Server process
RDM\$BIND_LCS_PRIORITY	Base Priority for the LCS Server process
RDM\$BIND_LRS_PRIORITY	Base Priority for the LRS Server process
RDM\$BIND_RCS_PRIORITY	Base Priority for the RCS Server process

When the Hot Standby feature is installed, the RDMAIJSERVER account is created specifying an account priority of 15. The priority of AIJ server processes on your system can be restricted with the system-wide logical name RDM\$BIND_AIJSRV_PRIORITY. If this logical name is defined to a value less than 15, an AIJ server process will adjust its base priority to the value specified when the AIJ server process starts. Values from 0 to 31 are allowed for RDM\$BIND_AIJSRV_PRIORITY, but the process is not able to raise its priority above the RDMAIJSERVER account value.

For most applications and systems, Oracle discourages changing the server process priorities.

5.1.2 Waiting for Client Lock Message

The Oracle Rdb7 Guide to Database Performance and Tuning contains a section in Chapter 3 that describes the Performance Monitor Stall Messages screen. The section contains a list describing the "Waiting for" messages. The description of the "waiting for client lock" message was missing from the list.

A client lock indicates that an Rdb metadata lock is in use. The term client indicates that Rdb is a client of the Rdb locking services. The metadata locks are used to guarantee memory copies of the metadata (table, index and column definitions) are consistent with the on-disk versions.

The "waiting for client lock" message means the database user is requesting an incompatible locking mode. For example, when trying to drop a table which is in use, the drop operation requests a PROTECTED WRITE lock on the metadata object (such as a table) which is incompatible with the existing PROTECTED READ lock currently used by others of the table.

These metadata locks consist of three longwords. The lock is displayed in text format first, followed by its hexadecimal representation. The text version masks out non–printable characters with a dot (.).

The leftmost value seen in the hexadecimal output contains the id of the object. The id is described below for tables and views, functions, procedures, and modules.

- ◆ For tables and views, the id represents the unique value found in the RDB\$RELATION_ID column of the RDB\$RELATIONS system relation for the given table.
- ◆ For routines, the id represents the unique value found in the RDB\$ROUTINE_ID column of the RDB\$ROUTINES system relation for the given routine.
- ◆ For modules, the id represents the unique value found in the RDB\$MODULE_ID column of the RDB\$MODULES system relation for the given module.

The next value displayed signifies the object type. The following table describes objects and their hexadecimal type values.

Table 5–2 Objects and Their Hexadecimal Type Value

Object	Hexadecimal Value
Tables or views	00000004
Routines	00000016
Modules	00000015

The last value in the hexadecimal output represents the lock type. The value 55 indicates this is a client lock.

The following example shows a "waiting for client lock" message from a Stall Messages screen:

```
Process.ID Since..... Stall.reason..... Lock.ID.
46001105:2 10:40:46.38 - waiting for client '.....' 000000190000000400000055
```

To determine the name of the referenced object given the lock ID, the following queries can be used based on the object type:

```
SQL> select RDB$RELATION_NAME from RDB$RELATIONS where RDB$RELATION_ID = 25;
SQL> select RDB$MODULE_NAME from RDB$MODULES where RDB$MODULE_ID = 12;
SQL> select RDB$ROUTINE_NAME from RDB$ROUTINES where RDB$ROUTINE_ID = 7;
```

Because the full client lock output is long, it may require more space than is allotted for the Stall.reason column and therefore can be overwritten by the Lock.ID. column output.

For more detailed lock information, perform the following steps:

- ◆ Press the L option from the horizontal menu to display a menu of lock IDs.
- ◆ Select the desired lock ID.

5.1.3 Clarification of PREPARE Statement Behavior

Bug 2581863

According to the Oracle Rdb7 SQL Reference Manual, Volume 3 page 7–227, when using a statement–id parameter for PREPARE "if that parameter is an integer, then you must explicitly initialize that integer to zero before executing the PREPARE statement".

This description is not correct and should be replaced with this information:

1. If the statement–id is non–zero and does not match any prepared statement (the id was stale or contained a random value), then an error is raised:
%SQL–F–BADPREPARE, Cannot use DESCRIBE or EXECUTE on a statement that is not prepared
2. If the statement–id is non–zero, or the statement name is one that has previously been used and matches an existing prepared statement, then that statement is automatically released prior to the prepare of the new statement. Please refer to the RELEASE statement for further details.
3. If the statement–id is zero or was automatically released, then a new statement–id is allocated and the statement prepared.

Please note that if you use statement–name instead of a statement–id–parameter then SQL will implicitly declare an id for use by the application. Therefore, the semantics described apply similarly when using the statement–name. See the RELEASE statement for details.

5.1.4 SQL EXPORT Does Not Save Some Database Attributes

Bug 2574640

The SQL EXPORT and IMPORT commands do not support several database attributes.

The following attributes are not saved by EXPORT for this release of Oracle Rdb.

- ◆ CHECKPOINT TIMED EVERY n SECONDS
- ◆ CHECKPOINT { ALL | UPDATED } ROWS TO { BACKING FILE | DATABASE }
- ◆ RMU/SET LOGMINER

This problem has been corrected in Oracle Rdb Release 7.1. The EXPORT protocol has been extended to support these and other database attributes. If you have databases with these attributes enabled then after the IMPORT completes, use a SQL script to re–establish the settings for the new database.

5.1.5 RDM\$BIND_LOCK_TIMEOUT_INTERVAL Overrides the Database Parameter

Bug 2203700

When starting a transaction, there are three different values that are used to determine the lock timeout interval for that transaction. Those values are:

1. The value specified in the SET TRANSACTION statement
2. The value stored in the database as specified in CREATE or ALTER DATABASE
3. The value of the logical name RDM\$BIND_LOCK_TIMEOUT_INTERVAL

The timeout interval for a transaction is the smaller of the value specified in the SET TRANSACTION statement and the value specified in CREATE DATABASE. However, if the logical name RDM\$BIND_LOCK_TIMEOUT_INTERVAL is defined, the value of this logical name overrides the value specified in CREATE DATABASE.

The description of how these three values interact, found in several different parts of the Rdb documentation set, is incorrect and will be replaced by the description above.

The lock timeout value in the database can be dynamically modified from the Locking Dashboard in RMU/SHOW STATISTICS. The Per-Process Locking Dashboard can be used to dynamically override the logical name RDM\$BIND_LOCK_TIMEOUT_INTERVAL for one or more processes.

5.1.6 New Request Options for RDO, RDBPRE and RDB\$INTERPRET


This release note was included in the V70A Release Notes but had gotten dropped somewhere along the line.

For this release of Rdb, two new keywords have been added to the handle-options for the DECLARE_STREAM, the START_STREAM (undeclared format) and FOR loop statements. These changes have been made to RDBPRE, RDO and RDB\$INTERPRET at the request of several RDO customers.

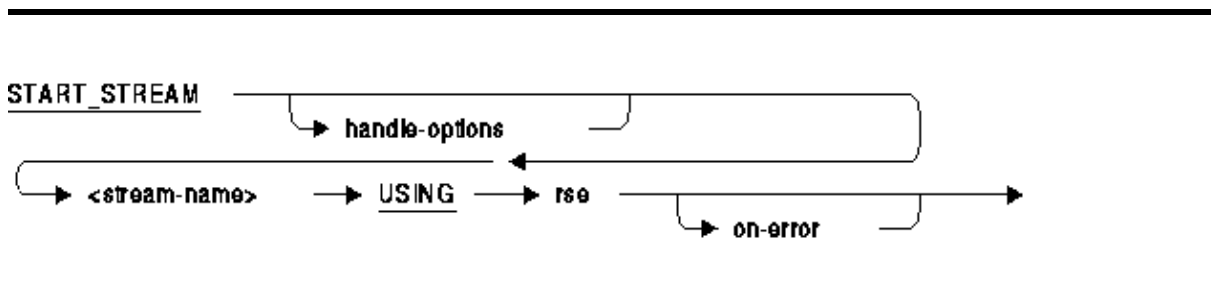
In prior releases, the handle-options could not be specified in interactive RDO or RDB\$INTERPRET. This has changed in Rdb but these allowed options will be limited to MODIFY and PROTECTED keywords. For RDBPRE, all options listed will be supported. These option names were chosen to be existing keywords to avoid adding any new keywords to the RDO language.

The altered statements are shown in Example 5-1, Example 5-2 and Example 5-3.


Example 5-1 DECLARE_STREAM Format



Example 5-2 START_STREAM Format




Example 5-3 FOR Format



Each of these statements references the syntax for the HANDLE-OPTIONS which has been revised and is shown below.

handle-options =



The following options are available for HANDLE-OPTIONS:

- ◆ **REQUEST_HANDLE** specifies the request handle for this request. This option is only valid for RDBPRE and RDML applications. It cannot be used with RDB\$INTERPRET, nor interactive RDO.
- ◆ **TRANSACTION_HANDLE** specifies the transaction handle under which this request executes. This option is only valid for RDBPRE and RDML applications. It cannot be used with RDB\$INTERPRET, nor interactive RDO.
- ◆ **MODIFY** specifies that the application will modify all (or most) records fetched from the stream or for loop. This option can be used to improve application performance by avoiding lock promotion from SHARED READ for the FETCH to PROTECTED WRITE access for the nested MODIFY or ERASE statement. It can also reduce DEADLOCK occurrence because lock promotions are avoided.
This option is valid for RDBPRE, RDB\$INTERPRET, and interactive RDO. This option is not currently available for RDML.

For example:

```
RDO> FOR (MODIFY) E IN EMPLOYEES WITH E.EMPLOYEE_ID = "00164"
cont>     MODIFY E USING E.MIDDLE_INITIAL = "M"
cont>     END_MODIFY
cont>     END_FOR
```

This FOR loop uses the MODIFY option to indicate that the nested MODIFY is an unconditional statement and so aggressive locking can be undertaken during the fetch of the record in the FOR loop.

- ◆ **PROTECTED** specifies that the application may modify records fetched by this stream by a separate and independent MODIFY statement. Therefore, this stream should be protected from interference (aka Halloween affect). The optimizer will select a snapshot of the rows

and store them in a temporary relation for processing, rather than traversing indexes at the time of the FETCH statement. In some cases this may result in poorer performance when the temporary relation is large and overflows from virtual memory to a temporary disk file, but the record stream will be protected from interference. The programmer is directed to the documentation for the Oracle Rdb logical names RDMS\$BIND_WORK_VM and RDMS\$BIND_WORK_FILE.

This option is valid for RDBPRE, RDB\$INTERPRET, and interactive RDO. This option is not currently available for RDML.

The following example creates a record stream in a BASIC program using Callable RDO:

```
RDMS_STATUS = RDB$INTERPRET ( 'INVOKE DATABASE PATHNAME "PERSONNEL"' )

RDMS_STATUS = RDB$INTERPRET ( 'START_STREAM (PROTECTED) EMP USING ' + &
                               'E IN EMPLOYEES' )

RDMS_STATUS = RDB$INTERPRET ( 'FETCH EMP' )

DML_STRING = 'GET ' +
              '!VAL = E.EMPLOYEE_ID;' +
              '!VAL = E.LAST_NAME;' +
              '!VAL = E.FIRST_NAME' +
              'END_GET'

RDMS_STATUS = RDB$INTERPRET (DML_STRING, EMP_ID, LAST_NAME, FIRST_NAME)
```

In this case the FETCH needs to be protected against MODIFY statements which execute in other parts of the application.

5.1.7 Missing Descriptions of RDB\$FLAGS from HELP File

The HELP file for Oracle Rdb Release 7.0 describes the system tables for Oracle Rdb and was missing these updated descriptions of the RDB\$FLAGS column for several tables.

Table 5-3 Changed Columns for RDB\$INDICES Table

Column Name	Data Type	Domain Name	Comments
RDB\$FLAGS	Integer	RDB\$FLAGS	A bit mask where the bits have the following meaning when set:
			Bit 0: This index is of type HASHED.
			Bit 1: This index uses the MAPPING VALUES clause to compress integer value ranges.
			Bit 2: If this is a HASHED index then it is of type ORDERED. If clear this indicates the index is of type SCATTERED.
			Bit 3: Reserved for future use.
			Bit 4: This index has run length compression enabled (ENABLE COMPRESSION).

Bit 5: This index is no longer used (MAINTENANCE IS DISABLED).
Bit 6 through 10: Reserved for future use.
Bit 11: This index has duplicates compressed (DUPLICATES ARE COMPRESSED).
Bit 12: This index is of type SORTED RANKED.
Bits 13 through 31: Reserved for future use.

Table 5–4 Changed Columns for RDB\$RELATIONS Table

Column Name	Data Type	Domain Name	Comments
RDB\$FLAGS	Integer	RDB\$FLAGS	A bit mask where the bits have the following meaning when set:
			Bit 0: This relation is a view.
			Bit 1: This relation is not compressed.
			Bit 2: The SQL clause, WITH CHECK OPTION, is used in this view definition.
			Bit 3: Indicates a special internal system relation.
			Bit 4: This view is not an ANSI updatable view.
			Bit 5: This is an imported table in the Distributed Option for Rdb catalog.
			Bit 6: This is a passthru table in the Distributed Option for Rdb catalog.
			Bit 7: This is a partitioned view in the Distributed Option for Rdb catalog.
			Bit 8: This table has compression defined by the storage map. When set Bit 1 in this bit mask is ignored.
			Bit 9: This is a temporary table.
			Bit 10: When bit 9 is set this is a global temporary table, when clear it indicates a local temporary table.
			Bit 11: When bit 9 is set this indicates that the rows in the temporary table should be deleted upon COMMIT.
			Bit 12: Reserved for future use.
			Bit 13: A table (via a computed by column) or view references a local temporary table.
			Bit 14: Reserved for future use.
			Bit 15: This is a system table with a special storage map.
			Bits 16 through 31: Reserved for future use.

Table 5–5 Changed Columns for RDB\$STORAGE_MAPS Table

Column Name	Data	Domain Name	Comments
-------------	------	-------------	----------

	Type		
RDB\$FLAGS	Integer	RDB\$FLAGS	A bit mask where the bits have the following meaning when set:
			Bit 0: This table or index is mapped to page format MIXED areas.
			Bit 1: This partition is not compressed.
			Bit 2: This is a strictly partitioned storage map, the partitioning columns become read only for UPDATE.
			Bit 3 through 31: Reserved for future use.

5.1.8 A Way to Find the Transaction Type of a Particular Transaction Within the Trace Database

The table EPC\$1_221_TRANSACTION in the formatted Oracle Trace database has a column LOCK_MODE_START of longword datatype. The values of this column indicate the type of transaction a particular transaction was.

Value	Transaction type
-----	-----
8	Read only
9	Read write
14	Batch update

5.1.9 Clarification of SET FLAGS Option DATABASE_PARAMETERS

Bug 1668049

The Oracle Rdb7 SQL Reference Manual describes the option DATABASE_PARAMETERS in Table 7-6 in the SET FLAGS section. However, this keyword generates output only during ATTACH to the database which happens prior to the SET FLAGS statement executing.

This option is therefore only useful when used with the RDMS\$SET_FLAGS logical name which provides similar functionality.

```

$ define RDMS$SET_FLAGS "database_parameters"
$ sql$
SQL> Attach 'File db$:scratch';
ATTACH #1, Database BLUGUM$DKA300:[SMITHI.DATABASES.V70]SCRATCH.RDB;1
~P Database Parameter Buffer (version=2, len=79)
0000 (00000) RDB$K_DPB_VERSION2
0001 (00001) RDB$K_FACILITY_ALL
0002 (00002) RDB$K_DPB2_IMAGE_NAME "NODE::DISK:[DIR]SQL$70.EXE;1"
0040 (00064) RDB$K_FACILITY_ALL
0041 (00065) RDB$K_DPB2_DBKEY_SCOPE (Transaction)
0045 (00069) RDB$K_FACILITY_ALL
0046 (00070) RDB$K_DPB2_REQUEST_SCOPE (Attach)
004A (00074) RDB$K_FACILITY_RDB_VMS
004B (00075) RDB$K_DPB2_CDD_MAINTAINED (No)

```

```
RDMS$BIND_WORK_FILE = "DISK:[DIR]RDMSTTBL$UEOU3LQ0RV2.TMP;" (Visible = 0)
SQL> Exit
DETACH #1
```

5.1.10 Additional Information About Detached Processes

Oracle Rdb documentation omits necessary detail on running Oracle Rdb from a detached process.

Applications run from a detached process must ensure that the OpenVMS environment is established correctly before running Oracle Rdb. Otherwise, Oracle Rdb will not execute.

Attempts to attach to a database and execute an Oracle Rdb query from applications running as detached processes will result in an error similar to the following:

```
%RDB-F-SYS_REQUEST, error from system services request
-SORT-E-OPENOUT, error opening {file} as output
-RMS-F-DEV, error in device name or inappropriate device type for
operation
```

The problem occurs because a detached process does not normally have the logical names `SYSS$LOGIN` or `SYSS$SCRATCH` defined.

There are two methods that can be used to correct this:

1. Use the DCL command procedure `RUN_PROCEDURE` to run the `ACCOUNTS` application: `RUN_PROCEDURE.COM` includes the single line:

```
$ RUN ACCOUNTS_REPORT
```

Then execute this procedure using this command:

```
$ RUN/DETACH/AUTHORIZE SYS$SYSTEM:LOGINOUT/INPUT=RUN_PROCEDURE
```

This solution executes `SYSS$SYSTEM:LOGINOUT` so that the command language interface (DCL) is activated. This causes the logical names `SYSS$LOGIN` and `SYSS$SCRATCH` to be defined for the detached process. The `/AUTHORIZE` qualifier also ensures that the users' process quota limits (PQLs) are used from the system authorization file rather than relying on the default PQL system parameters, which are often insufficient to run Oracle Rdb.

2. If DCL is not desired, and `SYSS$LOGIN` and `SYSS$SCRATCH` are not defined, then prior to executing any Oracle Rdb statement, you should define the following logical names:

```
◇ RDMS$BIND_WORK_FILE
```

Define this logical name to allow you to reduce the overhead of disk I/O operations for matching operations when used in conjunction with the `RDMS$BIND_WORK_VM` logical name. If the virtual memory file is too small, then overflow to disk will occur at the disk and directory location specified by `RDMS$BIND_WORK_FILE`.

For more information on `RDMS$BIND_WORK_FILE` and

`RDMS$BIND_WORK_VM`, see the Oracle Rdb Guide to Database Performance and Tuning.

◇ SORTWORK0, SORTWORK1, and so on

The OpenVMS sort/merge utility (SORT/MERGE) attempts to create sort work files in SYS\$SCRATCH. If the SORTWORK logical names exist, the utility will not require the SYS\$SCRATCH logical. However, note that not all queries will require sorting, and that some sorts will be completed in memory and so will not necessarily require disk space.

If you use the logical RDMS\$BIND_SORT_WORKFILES, you will need to define further SORTWORK logical names as described in the Oracle Rdb Guide to Database Performance and Tuning.

You should also verify that sufficient process quotas are specified on the RUN/DETACH command line, or defined as system PQL parameters to allow Oracle Rdb to execute.

5.1.11 The Halloween Problem

When a cursor is processing rows selected from a table, it is possible that another separate query can interfere with the retrieval of the cursor by modifying the index column's key values used by the cursor.

For instance, if a cursor selects all EMPLOYEES with LAST_NAME >= 'M', it is likely that the query will use the sorted index on LAST_NAME to retrieve the rows for the cursor. If an update occurs during the processing of the cursor which changes the LAST_NAME of an employee from "Mason" to "Rickard", then it is possible that that employee row will be processed twice. First, when it is fetched with name "Mason", and then later when it is accessed by the new name "Rickard".

The Halloween problem is a well known problem in relational databases. Access strategies which optimize the I/O requirements, such as Index Retrieval, can be subject to this problem. Interference from queries by other sessions are avoided by locking and are controlled by the ISOLATION LEVEL options in SQL, or the CONCURRENCY/CONSISTENCY options in RDO/RDML.

Oracle Rdb avoids this problem if it knows that the cursors subject table will be updated. For example, if the SQL syntax UPDATE ... WHERE CURRENT OF is used to perform updates of target rows, or if the RDO/RDML MODIFY statement uses the context variable for the stream. Then the optimizer will choose an alternate access strategy if an update can occur which may cause the Halloween problem. This can be seen in the access strategy in the example below as a "Temporary relation" being created to hold the result of the cursor query.

When you use interactive or dynamic SQL, the UPDATE ... WHERE CURRENT OF or DELETE ... WHERE CURRENT OF statements will not be seen until after the cursor is declared and opened. In these environments, you must use the FOR UPDATE clause to specify that columns selected by the cursor will be updated during cursor processing. This is an indication to the Rdb optimizer so that it protects against the Halloween problem in this case. This is shown in the two examples below [Example 5-1](#) and [Example 5-2](#).

[Example 5-1](#) shows that the EMP_LAST_NAME index is used for retrieval. Any update performed will possibly be subject to the Halloween problem.

Example 5-1 Interactive Cursor with no Halloween Protection

```
SQL> set flags 'strategy';
SQL> declare emp cursor for
cont> select * from employees where last_name >= 'M'
cont> order by last_name;
```

```
SQL> open emp;
Conjunct      Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_LAST_NAME [1:0]
SQL> close emp;
```

Example 5–2 shows that the query specifies that the column LAST_NAME will be updated by some later query. Now the optimizer protects the EMP_LAST_NAME index used for retrieval by using a "Temporary Relation" to hold the query result set. Any update performed on LAST_NAME will now avoid the Halloween problem.

Example 5–2 Interactive Cursor with Halloween Protection

```
SQL> set flags 'strategy';
SQL> declare emp2 cursor for
cont> select * from employees where last_name >= 'M'
cont> order by last_name
cont> for update of last_name;
SQL> open emp2;
Temporary relation      Conjunct      Get
Retrieval by index of relation EMPLOYEES
  Index name  EMP_LAST_NAME [1:0]
SQL> close emp2;
```

When you use the SQL precompiler or the SQL module language compiler, it can be determined from usage that the cursor context will possibly be updated during the processing of the cursor because all cursor related statements are present within the module. This is also true for the RDML/RDBPRE precompilers when you use the DECLARE_STREAM and START_STREAM statements and use the same stream context to perform all MODIFY and ERASE statements.

The point to note here is that the protection takes place during the open of the SQL cursor (or RDO stream), not during the subsequent UPDATE or DELETE.

If you execute a separate UPDATE query which modifies rows being fetched from the cursor, then the actual rows fetched will depend upon the access strategy chosen by the Rdb optimizer. As the query is separate from the cursors query (i.e. doesn't reference the cursor context) then the optimizer does not know that the cursor selected rows are potentially updated and so can not perform the normal protection against the Halloween problem.

5.1.12 RDM\$BIND_MAX_DBR_COUNT Documentation Clarification

Bug 1495227

The Rdb7 Guide to Database Performance and Tuning Manual, Volume 2, Page A–18, incorrectly describes the use of the RDM\$BIND_MAX_DBR_COUNT logical.

Following is an updated description. Note that the difference in actual behavior between what is in the existing documentation and software is that the logical name only controls the number of database recovery processes created at once during "node failure" recovery (that is, after a system or monitor crash or other abnormal shutdown).

When an entire database is abnormally shut down (due, for example, to a system failure), the database will have to be recovered in a "node failure" recovery mode. This recovery will be performed by another monitor in the cluster if the database is opened on another node or will be performed the next time the database is opened.

The RDM\$BIND_MAX_DBR_COUNT logical name and the RDB_BIND_MAX_DBR_COUNT configuration parameter define the maximum number of database recovery (DBR) processes to be simultaneously invoked by the database monitor during a "node failure" recovery.

This logical name and configuration parameter apply only to databases that do not have global buffers enabled. Databases that utilize global buffers have only one recovery process started at a time during a "node failure" recovery.

In a "node failure" recovery situation with the Row Cache feature enabled (regardless of the global buffer state), the database monitor will start a single database recovery (DBR) process to recover the Row Cache Server (RCS) process and all user processes from the oldest active checkpoint in the database.

5.1.13 RMU /UNLOAD /AFTER_JOURNAL NULL Bit Vector Clarification

Each output record from the RMU /UNLOAD /AFTER_JOURNAL command includes a vector (array) of bits. There is one bit for each field in the data record. If a null bit value is 1, the corresponding field is NULL; if a null bit value is 0, the corresponding field is not NULL and contains an actual data value. The contents of a data field that is NULL are not initialized and are not predictable.

The null bit vector begins on a byte boundary. The field RDB\$LM_NBV_LEN indicates the number of valid bits (and thus, the number of columns in the table). Any extra bits in the final byte of the vector after the final null bit are unused and the contents are unpredictable.

The following example C program demonstrates one possible way of reading and parsing a binary output file (including the null bit vector) from the RMU /UNLOAD /AFTER_JOURNAL command. This sample program has been tested using Oracle Rdb V7.0.5 and HP C V6.2-009 on OpenVMS Alpha V7.2-1. It is meant to be used as a template for writing your own program.

```

/* DATATYPES.C */

#include <stdio.h>
#include <descrip.h>
#include <starlet.h>
#include <string.h>

#pragma member_alignment __save
#pragma nomember_alignment

struct { /* Database key structure */
    unsigned short    lno;    /* line number */
    unsigned int      pno;    /* page number */
    unsigned short    dbid;   /* area number */
} dbkey;

typedef struct { /* Null bit vector with one bit for each column */
    unsigned          n_tinyint    :1;
    unsigned          n_smallint   :1;

```

Oracle® Rdb for OpenVMS

```

unsigned          n_integer   :1;
unsigned          n_bigint    :1;
unsigned          n_double    :1;
unsigned          n_real      :1;
unsigned          n_fixstr    :1;
unsigned          n_varstr    :1;
} nbv_t;

struct { /* LogMiner output record structure for table DATATYPES */
char          rdb$lm_action;
char          rdb$lm_relation_name [31];
int           rdb$lm_record_type;
short        rdb$lm_data_len;
short        rdb$lm_nbv_len;
__int64      rdb$lm_dbk;
__int64      rdb$lm_start_tad;
__int64      rdb$lm_commit_tad;
__int64      rdb$lm_tsn;
short        rdb$lm_record_version;
char         f_tinyint;
short        f_smallint;
int          f_integer;
__int64      f_bigint;
double       f_double;
float        f_real;
char         f_fixstr[10];
short        f_varstr_len; /* length of varchar */
char         f_varstr[10]; /* data of varchar */
nbv_t        nbv;
} lm;

#pragma member_alignment __restore

main ()
{  char timbuf [24];
   struct dsc$descriptor_s dsc = {
       23, DSC$K_DTYPE_T, DSC$K_CLASS_S, timbuf};
   FILE *fp = fopen ("datatypes.dat", "r", "ctx=bin");

   memset (&timbuf, 0, sizeof(timbuf));

   while (fread (&lm, sizeof(lm), 1, fp) != 0)
   {
       printf ("Action      = %c\n",      lm.rdb$lm_action);
       printf ("Table        = %.*s\n",      sizeof(lm.rdb$lm_relation_name),
           lm.rdb$lm_relation_name);

       printf ("Type          = %d\n",      lm.rdb$lm_record_type);
       printf ("Data Len     = %d\n",      lm.rdb$lm_data_len);
       printf ("Null Bits    = %d\n",      lm.rdb$lm_nbv_len);

       memcpy (&dbkey, &lm.rdb$lm_dbk, sizeof(lm.rdb$lm_dbk));
       printf ("DBKEY        = %d:%d:%d\n", dbkey.dbid,
           dbkey.pno,
           dbkey.lno);

       sys$asctim (0, &dsc, &lm.rdb$lm_start_tad, 0);
       printf ("Start TAD    = %s\n", timbuf);

       sys$asctim (0, &dsc, &lm.rdb$lm_commit_tad, 0);
       printf ("Commit TAD   = %s\n", timbuf);
   }
}

```

Oracle® Rdb for OpenVMS

```

printf ("TSN          = %Ld\n",      lm.rdb$lm_tsn);
printf ("Version      = %d\n",      lm.rdb$lm_record_version);

if (lm.nbv.n_tinyint == 0)
    printf ("f_tinyint  = %d\n",    lm.f_tinyint);
else     printf ("f_tinyint  = NULL\n");

if (lm.nbv.n_smallint == 0)
    printf ("f_smallint = %d\n",    lm.f_smallint);
else     printf ("f_smallint = NULL\n");

if (lm.nbv.n_integer == 0)
    printf ("f_integer  = %d\n",    lm.f_integer);
else     printf ("f_integer  = NULL\n");

if (lm.nbv.n_bigint == 0)
    printf ("f_bigint   = %Ld\n",   lm.f_bigint);
else     printf ("f_bigint   = NULL\n");

if (lm.nbv.n_double == 0)
    printf ("f_double   = %f\n",    lm.f_double);
else     printf ("f_double   = NULL\n");

if (lm.nbv.n_real == 0)
    printf ("f_real     = %f\n",    lm.f_real);
else     printf ("f_real     = NULL\n");

if (lm.nbv.n_fixstr == 0)
    printf ("f_fixstr   = %.*s\n",    sizeof (lm.f_fixstr),
                                                lm.f_fixstr);
else     printf ("f_fixstr   = NULL\n");

if (lm.nbv.n_varstr == 0)
    printf ("f_varstr   = %.*s\n",    lm.f_varstr_len, lm.f_varstr);
else     printf ("f_varstr   = NULL\n");

printf ("\n");
}
}

```

Example sequence of commands to create a table, unload the data and display the contents with this program:

```

SQL> ATTACH 'FILE MF_PERSONNEL';
SQL> CREATE TABLE DATATYPES (
    F_TINYINT TINYINT
  ,F_SMALLINT SMALLINT
  ,F_INTEGER INTEGER
  ,F_BIGINT BIGINT
  ,F_DOUBLE DOUBLE PRECISION
  ,F_REAL REAL
  ,F_FIXSTR CHAR (10)
  ,F_VARSTR VARCHAR (10));
SQL> COMMIT;
SQL> INSERT INTO DATATYPES VALUES (1, NULL, 2, NULL, 3, NULL, 'THIS', NULL);
SQL> INSERT INTO DATATYPES VALUES (NULL, 4, NULL, 5, NULL, 6, NULL, 'THAT');
SQL> COMMIT;
SQL> EXIT;
$ RMU /BACKUP /AFTER_JOURNAL MF_PERSONNEL AIJBCK.AIJ
$ RMU /UNLOAD /AFTER_JOURNAL MF_PERSONNEL AIJBCK.AIJ -
    /TABLE = (NAME=DATATYPES, OUTPUT=DATATYPES.DAT)

```

```
$ CC DATATYPES.C
$ LINK DATATYPES.OBJ
$ RUN DATATYPES.EXE
```

5.1.14 Location of Host Source File Generated by the SQL Precompilers

Bug 478898

When the SQL precompiler generates host source files (like .c, .pas, .for) from the precompiler source files, it locates these files based on the /obj qualifier located on the command line given to the SQL precompiler.

The following examples show the location where the host source file is generated. When /obj is not specified on the command line, the object and the host source file take the name of the SQL precompiler source files with the extensions of .obj and .c respectively.

```
LUND> sqlpre/cc scc_try_mli_successful.sc
LUND> dir scc_try_mli_successful.*
```

Directory MYDISK:[LUND]

```
SCC_TRY_MLI_SUCCESSFUL.C;1          SCC_TRY_MLI_SUCCESSFUL.OBJ;2
SCC_TRY_MLI_SUCCESSFUL.SC;2
```

Total of 3 files.

When /obj is specified on the command line, the object and the host source take the name given on the qualifier switch. It uses the default of the SQL precompiler source if a filespec is not specified. It uses the defaults of .obj and .c if the extension is not specified. If the host language is other than C, then it uses the appropriate host source extension (like .pas, .for, etc). The files also default to the current directory if a directory specification is not specified.

```
LUND> sqlpre/cc/obj=myobj scc_try_mli_successful.sc
LUND> dir scc_try_mli_successful.*
```

Directory MYDISK:[LUND]

```
SCC_TRY_MLI_SUCCESSFUL.SC;2
```

Total of 1 file.

```
LUND> dir myobj.*
```

Directory MYDISK:[LUND]

```
MYOBJ.C;1          MYOBJ.OBJ;2
```

Total of 2 files.

```
LUND> sqlpre/cc/obj=MYDISK:[lund.tmp] scc_try_mli_successful.sc
LUND> dir scc_try_mli_successful.*
```

Directory MYDISK:[LUND]

```
SCC_TRY_MLI_SUCCESSFUL.SC;2
```

```
Total of 1 file.
LUND> dir MYDISK:[lund.tmp]scc_try_mli_successful.*

Directory MYDISK:[LUND.TMP]

SCC_TRY_MLI_SUCCESSFUL.C;1                SCC_TRY_MLI_SUCCESSFUL.OBJ;2

Total of 2 files.
```

This problem has been corrected in Oracle Rdb Release 7.0.6.

5.1.15 Suggestion to Increase GH_RSRVPGCNT Removed

The Oracle Rdb7 for OpenVMS Installation and Configuration Guide contains a section titled "Installing Oracle Rdb Images as Resident on OpenVMS Alpha" that includes information about increasing the value of the OpenVMS system parameter GH_RSRVPGCNT when you modify the RMONSTART.COM or SQL\$STARTUP.COM procedures to install Rdb images with the /RESIDENT qualifier.

Note that modifying the parameter GH_RSRVPGCNT is only ever required if the RMONSTART.COM or SQL\$STARTUP.COM procedures have been manually modified to install Rdb images with the /RESIDENT qualifier. Furthermore, if the RMONSTART.COM and SQL\$STARTUP.COM procedures are executed during the system startup procedure (directly from SYSTARTUP_VMS.COM, for example), then there is no need to modify the GH_RSRVPGCNT parameter.

Oracle and HP suggest that you do not modify the value of the GH_RSRVPGCNT system parameter unless it is absolutely required. Some versions of OpenVMS on some hardware platforms require that GH_RSRVPGCNT be zero in order to ensure the highest level of system performance.

5.1.16 Clarification of the DDLDONOTMIX Error Message

Bug 454080

The ALTER DATABASE statement performs two classes of functions: changing the database root structures in the .RDB file and modifying the system metadata in the RDB\$SYSTEM storage area. The first class of changes do not require a transaction to be active. However, the second class requires that a transaction be active. Oracle Rdb does not currently support the mixing of these two classes of ALTER DATABASE clauses.

When you mix clauses that fall into both classes, the error message DDLDONOTMIX "the {SQL-syntax} clause can not be used with some ALTER DATABASE clauses" is displayed, and the ALTER DATABASE statement fails.

```
SQL> alter database filename MF_PERSONNEL
cont> dictionary is not used
cont> add storage area JOB_EXTRA filename JOB_EXTRA;
%RDB-F-BAD_DPB_CONTENT, invalid database parameters in the database parameter
block (DPB)
-RDMS-E-DDLDONOTMIX, the "DICTIONARY IS NOT USED" clause can not be used with
some ALTER DATABASE clauses
```

The following clauses may be mixed with each other but may not appear with other clauses such as ADD

STORAGE AREA, or ADD CACHE.

- DICTIONARY IS [NOT] REQUIRED
- DICTIONARY IS NOT USED
- MULTISCHEMA IS { ON | OFF }
- CARDINALITY COLLECTION IS { ENABLED | DISABLED }
- METADATA CHANGES ARE { ENABLED | DISABLED }
- WORKLOAD COLLECTION IS { ENABLED | DISABLED }

If the DDLDONOTMIX error is displayed, then restructure the ALTER DATABASE into two statements, one for each class of actions.

```
SQL> alter database filename MF_PERSONNEL
cont> dictionary is not used;
SQL> alter database filename MF_PERSONNEL
cont> add storage area JOB_EXTRA filename JOB_EXTRA;
```

5.1.17 Compressed Sorted Index Entry Stored in Incorrect Storage Area

This note was originally included in the Oracle Rdb Release 7.0.1.3 and 7.0.2 Release Notes. The logical name documented in the note for those releases was documented incorrectly. Below is a corrected note.

In specific cases, in versions V6.1 and V7.0 of Oracle Rdb, when a partitioned, compressed sorted index was created after the data was inserted into the table, b–tree entries may have been inserted into the wrong storage area.

All of the following criteria must be met in order for the possibility of this problem to occur:

- CREATE INDEX is issued after there are records already in the table on which the index is being created
- index must be partitioned over a single column
- index must have compression enabled
- scale factor must be zero on the columns of the index
- no collating sequences specified on the columns of the index
- no descending indexes
- MAPPING VALUES must not be specified

RMU/DUMP/AREA=xx will show that the b–tree entry was not stored in the expected storage area. However, in versions V6.1 and V7.0 of Oracle Rdb, the rows of the table can still be successfully retrieved.

The following example shows the problem:

```
create database
  filename foo
create storage area Area_1
  filename Area_1
create storage area Area_2
  filename Area2;

create table T1
```



```

(C1 integer);

! insert data into table prior to index creation
insert into T1 values (0);
commit;

! create index with COMPRESSION ENABLED
create index Index_1
  on T1 (C1)
  enable compression
  store using (C1)
  in Area_1 with limit of (0)
  otherwise in Area_2;
COMMIT;
!
! Dump out the page for b-tree in AREA_1, there are 0 bytes stored.
! There should be 5 bytes stored for the b-tree entry.
!
RMU/DUMP/AREA=AREA_1
.
.
.
.          .... total B-tree node size: 430
          0030 2003 0240 line 0 (2:5:0) index: set 48
          002F FFFFFFFF FFFF 0244 owner 47:-1:-1
                   0000 024C 0 bytes of entries <---***** no entry
                   8200 024E level 1, full suffix
00000000000000000000000000000000 0250 unused '.....'
.
.
.
!
! Dump out the page for b-tree in AREA_2, there are 5 bytes stored
!
RMU/DUMP/AREA=AREA_2
.
.
.
.          .... total B-tree node size: 430
          0031 2003 0240 line 0 (3:5:0) index: set 49
          002F FFFFFFFF FFFF 0244 owner 47:-1:-1
                   000A 024C 10 bytes of entries
                   8200 024E level 1, full suffix
                   00 05 0250 5 bytes stored, 0 byte prefix <---entry
          0100008000 0252 key '.....'
                   22B1 10 0257 pointer 47:554:0
.
.
.
```

This problem occurs when index compression is enabled. Therefore, a workaround is to create the index with compression disabled (which is the default). Once this update kit is applied, it is recommended that the index be dropped and recreated with compression enabled to rebuild the b-tree.

Note

In prior versions, the rows were successfully retrieved even though the key values were stored in the wrong storage area. This was due to the range query algorithm skipping empty partitions or scanning extra areas.

However, due to an enhancement in the algorithm for range queries on partitioned SORTED indexes in Oracle Rdb Release 7.0.2, the rows of the table which are stored in the incorrect storage areas may not be retrieved when using the partitioned index.

The optimized algorithm now only scans the relevant index areas (and no longer skips over empty areas) resulting in only those rows being returned. Therefore, it is recommended that the index be dropped and re-created. For a short term solution, another alternative is to disable the new optimization by defining the logical RDM\$INDEX_PART_CHECK to 0.

This problem has been corrected in Oracle Rdb Release 7.0.1.3.

5.1.18 Partition Clause is Optional on CREATE STORAGE MAP

Bug 642158

In the *Oracle Rdb7 SQL Reference Manual*, the syntax diagram for the CREATE STORAGE MAP statement incorrectly shows the partition clause as required syntax. The partition clause is not a required clause.

This correction will appear in the next publication of the *Oracle Rdb SQL Reference Manual*.

5.1.19 Oracle Rdb Logical Names

The *Oracle Rdb7 Guide to Database Performance and Tuning* contains a table in Chapter 2 summarizing the Oracle Rdb logical names and configuration parameters. The information in the following table supersedes the entries for the RDM\$BIND_RUJ_ALLOC_BLKCNT and RDM\$BIND_RUJ_EXTEND_BLKCNT logical names.

Logical Name Configuration Parameter	Function
RDM\$BIND_RUJ_ALLOC_BLKCNT	Allows you to override the default value of the .ruj file. The block count value can be defined between 0 and 2 billion with a default of 127.
RDM\$BIND_RUJ_EXTEND_BLKCNT	Allows you to pre-extend the .ruj files for each process using a database. The block count value can be defined between 0 and 65535 with a default of 127.

5.1.20 Documentation Error in *Oracle Rdb Guide to Database Performance and Tuning*

The *Oracle Rdb7 Guide to Database Performance and Tuning, Volume 2* contains an error in section C.7, "Displaying Sort Statistics with the R Flag".

When describing the output from this debugging flag, bullet 9 states:

- **Work File Alloc** indicates how many work files were used in the sort operation. A zero (0) value indicates that the sort was accomplished completely in memory.

This is incorrect. This statistic should be described as shown:

- *Work File Alloc* indicates how much space (in blocks) was allocated in the work files for this sort operation. A zero (0) value indicates that the sort was accomplished completely in memory.

This error will be corrected in a future release of *Oracle Rdb Guide to Database Performance and Tuning*.

5.1.21 SET FLAGS Option IGNORE_OUTLINE Not Available

Bug 510968

The *Oracle Rdb7 SQL Reference Manual* described the option IGNORE_OUTLINE in Table 7–6 of the SET FLAGS section. However, this keyword was not implemented in Oracle Rdb Release 7.0.

This has been corrected in this release of Oracle Rdb. This keyword is now recognized by the SET FLAGS statement. As a workaround the logical name RDMS\$BIND_OUTLINE_FLAGS "I" can be used to set this attribute.

5.1.22 SET FLAGS Option INTERNALS Not Described

The *Oracle Rdb7 SQL Reference Manual* does not describe the option INTERNALS in Table 7–6 in the SET FLAGS section. This keyword was available in first release of Oracle Rdb 7.0 and is used to enable debug flags output for internal queries such as constraints and triggers. It can be used in conjunction with other options such as STRATEGY, BLR, and EXECUTION. For example, the following flag settings are equivalent to defining the RDMS\$DEBUG_FLAGS as ISn and shows the strategy used by the trigger's actions on the AFTER DELETE trigger on the EMPLOYEES table.

```
SQL> SET FLAGS 'STRATEGY, INTERNAL, REQUEST_NAME';
SQL> SHOW FLAGS

Alias RDB$DBHANDLE:
Flags currently set for Oracle Rdb:
INTERNALS, STRATEGY, PREFIX, REQUEST_NAMES
SQL> DELETE FROM EMPLOYEES WHERE EMPLOYEE_ID = '00164';
~S: Trigger name EMPLOYEE_ID_CASCADE_DELETE
Get Temporary relation Retrieval by index of relation DEGREES
Index name DEG_EMP_ID [1:1]
~S: Trigger name EMPLOYEE_ID_CASCADE_DELETE
Get Temporary relation Retrieval by index of relation JOB_HISTORY
Index name JOB_HISTORY_HASH [1:1]
~S: Trigger name EMPLOYEE_ID_CASCADE_DELETE
Get Temporary relation Retrieval by index of relation SALARY_HISTORY
Index name SH_EMPLOYEE_ID [1:1]
~S: Trigger name EMPLOYEE_ID_CASCADE_DELETE
Conjunct Get Retrieval by index of relation DEPARTMENTS
Index name DEPARTMENTS_INDEX [0:0]
Temporary relation Get Retrieval by index of relation EMPLOYEES
Index name EMPLOYEES_HASH [1:1] Direct lookup
1 row deleted
```

5.1.23 Documentation for VALIDATE_ROUTINE Keyword for SET FLAGS

The SET FLAGS section of the *Oracle Rdb7 SQL Reference Manual* omitted the description of the VALIDATE_ROUTINE keyword (which can be negated as NOVALIDATE_ROUTINE). This keyword enables the re-validation of an invalidated stored procedure or function. This flag has the same action as the logical RDBMS\$VALIDATE_ROUTINE described in the *Oracle Rdb7 Guide to Database Performance and Tuning*.

This example shows the re-validation of a stored procedure. When the stored routine is successfully prepared (but not executed), the setting of VALIDATE_ROUTINE causes the entry for this routine in the RDB\$ROUTINES system table to be set as valid.

```
SQL> SET TRANSACTION READ WRITE;
SQL> SET FLAGS 'VALIDATE_ROUTINE';
SQL> SET NOEXECUTE;
SQL> CALL ADD_EMPLOYEE ('Smith');
SQL> SET EXECUTE;
SQL> COMMIT;
```

In this example, the use of the SET NOEXECUTE statement in interactive SQL allows the stored routine to be successfully compiled, but it is not executed.

5.1.24 Documentation for Defining the RDBSERVER Logical Name

Bugs 460611 and 563649.

Sections 4.3.7.1 and 4.3.7.2 in the *Oracle Rdb7 for OpenVMS Installation and Configuration Guide* provide the following examples for defining the RDBSERVER logical name:

```
$ DEFINE RDBSERVER SYS$SYSTEM:RDBSERVER70.EXE
and
$ DEFINE RDBSERVER SYS$SYSTEM:RDBSERVER61.EXE
```

These definitions are inconsistent with other command procedures that attempt to reference the RDBSERVERxx.EXE image. The following is one example where the RDBSERVER.COM procedure references SYS\$COMMON:<SYSEXE> and SYS\$COMMON:[SYSEXE], rather than SYS\$SYSTEM:

```
$ if .not. -
    ((f$locate ("SYS$COMMON:<SYSEXE>", rdbserver_image) .ne. log_len) .or. -
    (f$locate ("SYS$COMMON:[SYSEXE]", rdbserver_image) .ne. log_len))
$ then
$   say "'rdbserver_image' is not found in SYS$COMMON:<SYSEXE>"
$   say "RDBSERVER logical is 'rdbserver_image'"
$   exit
$ endif
```

In this case, if the logical name were defined as instructed in the *Oracle Rdb7 for OpenVMS Installation and Configuration Guide*, the image would not be found.

The *Oracle Rdb7 for OpenVMS Installation and Configuration Guide* should define the logical name as follows:

```
DEFINE RDBSERVER SYS$COMMON:<SYSEXE>RDBSERVER70.EXE
and
DEFINE RDBSERVER SYS$COMMON:<SYSEXE>RDBSERVER61.EXE
```

5.1.25 Undocumented SET Commands and Language Options

The following SET statements were omitted from the Oracle Rdb7 documentation.

5.1.25.1 QUIET COMMIT Option

The SET QUIET COMMIT statement (for interactive and dynamic SQL), the module header option QUIET COMMIT, the /QUIET_COMMIT (and /NOQUIET_COMMIT) qualifier for SQL module language, or the /SQLOPTIONS=QUIET_COMMIT (and NOQUIET_COMMIT) option for the SQL language precompiler allows the programmer to control the behavior of the COMMIT and ROLLBACK statements in cases where there is no active transaction.

By default, if there is no active transaction, SQL will raise an error when COMMIT or ROLLBACK is executed. This default is retained for backward compatibility for applications that may wish to detect the situation. If QUIET COMMIT is set to ON, then a COMMIT or ROLLBACK executes successfully when there is no active transaction.

Note

Within a compound statement, the COMMIT and ROLLBACK statements in this case are ignored.

Examples

In interactive or dynamic SQL, the following SET command can be used to disable or enable error reporting for COMMIT and ROLLBACK when no transaction is active. The parameter to the SET command is a string literal or host variable containing the keyword ON or OFF. The keywords may be in any case (upper, lower, or mixed).

```
SQL> COMMIT;
%SQL-F-NO_TXNOUT, No transaction outstanding
SQL> ROLLBACK;
%SQL-F-NO_TXNOUT, No transaction outstanding
SQL> SET QUIET COMMIT 'on';
SQL> ROLLBACK;
SQL> COMMIT;
SQL> SET QUIET COMMIT 'off';
SQL> COMMIT;
%SQL-F-NO_TXNOUT, No transaction outstanding
```

In the SQL module language or precompiler header, the clause QUIET COMMIT can be used to disable or enable error reporting for COMMIT and ROLLBACK when no transaction is active. The keyword ON or OFF must be used to enable or disable this feature. The following example enables QUIET COMMIT so that no error is reported if a COMMIT is executed when no transaction is active. For example:

```

MODULE TXN_CONTROL
LANGUAGE BASIC
PARAMETER COLONS
QUIET COMMIT ON

PROCEDURE S_TXN (SQLCODE);
SET TRANSACTION READ WRITE;

PROCEDURE C_TXN (SQLCODE);
COMMIT;

```

5.1.25.2 COMPOUND TRANSACTIONS Option

The SET COMPOUND TRANSACTIONS statement (for interactive and dynamic SQL) and the module header option COMPOUND TRANSACTIONS allows the programmer to control the SQL behavior for starting default transactions for compound statements.

By default, if there is no current transaction, SQL will start a transaction before executing a compound statement or stored procedure. However, this may conflict with the actions within the procedure, or may start a transaction for no reason if the procedure body does not perform any database access. This default is retained for backward compatibility for applications that may expect a transaction to be started for the procedure.

If COMPOUND TRANSACTIONS is set to EXTERNAL, then SQL starts a transaction before executing the procedure; otherwise, if it is set to INTERNAL, it allows the procedure to start a transaction as required by the procedure execution.

Examples

In interactive or dynamic SQL, the following SET command can be used to disable or enable transactions started by the SQL interface. The parameter to the SET command is a string literal or host variable containing the keyword INTERNAL or EXTERNAL. The keywords may be in any case (upper, lower, or mixed). For example:

```

SQL> SET COMPOUND TRANSACTIONS 'internal';
SQL> CALL START_TXN_AND_COMMIT ();
SQL> SET COMPOUND TRANSACTIONS 'external';
SQL> CALL UPDATE_EMPLOYEES (...);

```

In the SQL module language or precompiler header, the clause COMPOUND TRANSACTIONS can be used to disable or enable starting a transaction for procedures. The keyword INTERNAL or EXTERNAL must be used to enable or disable this feature.

```

MODULE TXN_CONTROL
LANGUAGE BASIC
PARAMETER COLONS
COMPOUND TRANSACTIONS INTERNAL

PROCEDURE S_TXN (SQLCODE);
BEGIN
SET TRANSACTION READ WRITE;
END;

PROCEDURE C_TXN (SQLCODE);

```

```
BEGIN
COMMIT;
END;
```

5.1.26 Undocumented Size Limit for Indexes with Keys Using Collating Sequences

Bug 586079

When a column is defined with a collating sequence, the index key is specially encoded to incorporate the correct ordering (collating) information. This special encoding takes more space than keys encoded for ASCII (the default when no collating sequence is used). Therefore, the encoded string uses more than the customary one byte per character of space within the index. This is true for all versions of Oracle Rdb that support collating sequences.

For all collating sequences, except Norwegian, the space required is approximately 9 bytes for every 8 characters. So, a CHAR (24) column will require approximately 27 bytes. For Norwegian collating sequences, the space required is approximately 10 bytes for every 8 characters.

The space required for encoding the string must be taken into account when calculating the size of an index key against the limit of 255 bytes. Suppose a column defined with a collating sequence of GERMAN was used in an index. The length of that column is limited to a maximum of 225 characters because the key will be encoded in 254 bytes.

The following example demonstrates how a 233 character column, defined with a German collating sequence and included in an index, exceeds the index size limit of 255 bytes, even though the column is defined as less than 255 characters in length:

```
SQL> CREATE DATABASE
cont>     FILENAME 'TESTDB.RDB'
cont>     COLLATING SEQUENCE GERMAN GERMAN;
SQL> CREATE TABLE EMPLOYEE_INFO (
cont>     EMP_NAME CHAR (233));
SQL> CREATE INDEX EMP_NAME_IDX
cont>     ON EMPLOYEE_INFO (
cont>     EMP_NAME     ASC)
cont>     TYPE IS SORTED;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-INDTOOBIG, requested index is too big
```

5.1.27 Changes to RMU/REPLICATE AFTER/BUFFERS Command

The behavior of the RMU/REPLICATE AFTER/BUFFERS command has been changed. The /BUFFERS qualifier may be used with either the CONFIGURE option or the START option.

When using local buffers, the AIJ log roll-forward server (LRS) will use a minimum of 4096 buffers. The value provided to the /BUFFERS qualifier will be accepted, but it will be ignored if it is less than 4096. In addition, further parameters will be checked and the number of buffers may be increased if the resulting

calculations are greater than the number of buffers specified by the /BUFFERS qualifier. If the database is configured to use more than 4096 AIJ request blocks (ARBs), then the number of buffers may be increased to the number of ARBs configured for the database. The LRS ensures that there are at least 10 buffers for every possible storage area in the database. Thus, if the total number of storage areas (both used and reserved) multiplied by 10 results in a greater number of buffers, that number will be used.

When global buffers are used, the number of buffers used by the AIJ log roll-forward server is determined as follows:

- If the /BUFFERS qualifier is omitted and the /ONLINE qualifier is specified, the number of buffers will default to the previously configured value, if any, or 256, whichever is larger.
- If the /BUFFERS qualifier is omitted and the /ONLINE qualifier is not specified or the /NOONLINE is specified, the number of buffers will default to the maximum number of global buffers allowed per user ("USER LIMIT"), or 256, whichever is larger.
- If the /BUFFERS qualifier is specified, that value must be at least 256, and it may not be greater than the maximum number of global buffers allowed per user ("USER LIMIT").

The /BUFFER qualifier now enforces a minimum of 256 buffers for the AIJ log roll-forward server. The maximum number of buffers allowed is still 524288 buffers.

5.1.28 Change in the Way RDMAIJ Server is Set Up in UCX

Starting with Oracle Rdb V7.0.2.1, the RDMAIJ image has become a variant image. Therefore, the information in section 2.12, "Step 10: Specify the Network Transport Protocol," of the *Oracle Rdb7 and Oracle CODASYL DBMS Guide to Hot Standby Databases* has become outdated in regards to setting up the RDMAIJSERVER object when using UCX as the network transport protocol. The UCX SET SERVICE command should now look similar to the following:

```
$ UCX SET SERVICE RDMAIJ -
  /PORT=<port_number> -
  /USER_NAME=RDMAIJ -
  /PROCESS_NAME=RDMAIJ -
  /FILE=SYS$SYSTEM:RDMAIJSERVER.com -
  /LIMIT=<limit>
```

And for Oracle Rdb multiversion, it should look similar to the following:

```
$ UCX SET SERVICE RDMAIJ70 -
  /PORT=<port_number> -
  /USER_NAME=RDMAIJ70 -
  /PROCESS_NAME=RDMAIJ70 -
  /FILE=SYS$SYSTEM:RDMAIJSERVER70.com -
  /LIMIT=<limit>
```

The installation procedure for Oracle Rdb creates a user named RDMAIJ(nn) and places a file called RDMAIJSERVER(nn).com in SYS\$SYSTEM and the RMONSTART(nn).COM command procedure will try to enable a service called RDMAIJ(nn) if UCX is installed and running.

Changing the RDMAIJ server to a multivariant image does not impact installations using DECNet since the correct DECNet object is created during the Rdb installation.

5.1.29 CREATE INDEX Supported for Hot Standby

On page 1–13 of the *Guide to Hot Standby Databases*, the add new index operation is incorrectly listed as an offline operation not supported by Hot Standby. The CREATE INDEX operation is now fully supported by Hot Standby, as long as the transaction does not span all available AIJ journals, including emergency AIJ journals.

5.1.30 Dynamic OR Optimization Formats

Bug 711643

In Table C–2 on Page C–7 of the *Oracle Rdb7 Guide to Database Performance and Tuning*, the dynamic OR optimization format is incorrectly documented as [l:h...]n. The correct formats for Oracle Rdb Release 7.0 and later are [(l:h)n] and [(l:h,l2:h2)].

Chapter 6

Known Problems and Restrictions

This chapter describes problems, restrictions, and workarounds known to exist in Oracle Rdb Release 7.0.7.1.

6.1 Oracle Rdb Considerations

6.1.1 AIJ Log Server Process May Loop Or Bugcheck

Bugs 2651475 and 1756433

Under unknown, but extremely rare conditions, on busy databases where the After Image Journal (AIJ) Log Server process is enabled, the ALS process has been observed to enter a loop condition writing AIJ information to the AIJ file(s).

In the worst case, this problem could cause all available journal files to be filled with repeating data. If no remedial action were to be taken, this condition could cause the database to be shutdown and the AIJ journals to be considered inaccessible.

The database is not corrupted by this problem.

Stopping and restarting the ALS process will clear the looping condition even if the ALS process must be stopped using the STOP/ID command.

Stopping the ALS process will not impact production as AIJ writes will automatically revert to the non-ALS behaviour.

In Oracle Rdb Release 7.0.7 the behaviour of Rdb has been changed so that should this problem be detected, the ALS process will automatically shutdown producing a bugcheck dump file. This will prevent any danger of filling all available journals and will ensure that the database remains available.

ALS may be safely restarted immediately as the conditions that cause such a loop are resolved during recovery of the ALS process.

6.1.2 Optimization of Check Constraints

Bug 1448422

When phrasing constraints using the "CHECK" syntax, a poorer strategy can be chosen by the optimizer than when the same or similar constraint is phrased using referential integrity (PRIMARY and FOREIGN KEY) constraints. Following is an example.

I have two tables T1 and T2, both with one column, and I wish to ensure that all values in table T1 exist in T2. Both tables have an index on the referenced field. I could use a PRIMARY KEY constraint on T2 and a FOREIGN KEY constraint on T1.

```
SQL> alter table t2
cont>  alter column f2 primary key not deferrable;
SQL> alter table t1
cont>  alter column f1 references t2 not deferrable;
```

When deleting from the PRIMARY KEY table, Oracle Rdb will only check for rows in the FOREIGN KEY table where the FOREIGN KEY has the deleted value. This can be seen as an index lookup on T1 in the retrieval strategy.

Oracle® Rdb for OpenVMS

```
SQL> delete from t2 where f2=1;
Get      Temporary relation      Retrieval by index of relation T2
      Index name  I2 [1:1]
Index only retrieval of relation T1
      Index name  I1 [1:1]
%RDB-E-INTEG_FAIL, violation of constraint T1_FOREIGN1 caused operation to fail
```

The failure of the constraint is not important. What is important is that Rdb efficiently detects that only those rows in T1 with the same values as the deleted row in T2 can be affected.

It is necessary sometimes to define this type of relationship using CHECK constraints. This could be necessary because the presence of NULL values in the table T2 precludes the definition of a primary key on that table. This could be done with a CHECK constraint of the form:

```
SQL> alter table t1
cont>  alter column f1
cont>  check (f1 in (select * from t2)) not deferrable;
SQL> delete from t2 where f2=1;
Get      Temporary relation      Retrieval by index of relation T2
      Index name  I2 [1:1]
Cross block of 2 entries
Cross block entry 1
      Index only retrieval of relation T1
      Index name  I1 [0:0]
Cross block entry 2
      Conjunct      Aggregate-F1      Conjunct
      Index only retrieval of relation T2
      Index name  I2 [0:0]
%RDB-E-INTEG_FAIL, violation of constraint T1_CHECK1 caused operation to fail
```

The cross block is for the constraint evaluation. This retrieval strategy indicates that to evaluate the constraint, the entire index on table T1 is being scanned and for each key, the entire index in table T2 is being scanned.

The behavior can be improved somewhat by using an equality join condition in the select clause of the constraint:

```
SQL> alter table t1
cont>  alter column f1
cont>  check (f1 in (select * from t2 where f2=f1))
cont>  not deferrable;
```

or:

```
SQL> alter table t1
cont>  alter column f1
cont>  check (f1=(select * from t2 where f2=f1))
cont>  not deferrable;
```

In both cases, the retrieval strategy will look as follows:

```
SQL> delete from t2 where f2=1;
Get      Temporary relation      Retrieval by index of relation T2
      Index name  I2 [1:1]
Cross block of 2 entries
Cross block entry 1
      Index only retrieval of relation T1
      Index name  I1 [0:0]
Cross block entry 2
```

Oracle® Rdb for OpenVMS

```
Conjunct      Aggregate-F1      Conjunct
Index only retrieval of relation T2
Index name I2 [1:1]
%RDB-E-INTEG_FAIL, violation of constraint T1_CHECK1 caused operation to fail
```

While the entire T1 index is scanned, at least the value from T1 is used to perform an index lookup on T2.

These restrictions result from semantic differences in the behavior of the "IN" and "EXISTS" operators with respect to null handling, and the complexity of dealing with non-equality join conditions.

To improve the performance of this type of integrity check on larger tables, it is possible to use a series of triggers to perform the constraint check. The following triggers perform a similar check to the constraints above.

```
SQL> create trigger t1_insert
cont> after insert on t1
cont> when (not exists (select * from t2 where f2=f1))
cont> (error) for each row;
SQL> create trigger t1_update
cont> after update on t1
cont> when (not exists (select * from t2 where f2=f1))
cont> (error) for each row;
SQL> ! A delete trigger is not needed on T1.
SQL> create trigger t2_delete
cont> before delete on t2
cont> when (exists (select * from t1 where f1=f2))
cont> (error) for each row;
SQL> create trigger t2_modify
cont> after update on t2
cont> referencing old as t2o new as t2n
cont> when (exists (select * from t1 where f1=t2o.f2))
cont> (error) for each row;
SQL> ! An insert trigger is not needed on T2.
```

The strategy for a delete on T2 is now:

```
SQL> delete from t2 where f2=1;
Aggregate-F1      Index only retrieval of relation T1
Index name I1 [1:1]
Temporary relation      Get      Retrieval by index of relation T2
Index name I2 [1:1]
%RDB-E-TRIG_INV_UPD, invalid update; encountered error condition defined for
trigger
-RDMS-E-TRIG_ERROR, trigger T2_DELETE forced an error
```

The trigger strategy is the index only retrieval displayed first. You will note that the index on T1 is used to examine only those rows that may be affected by the delete.

Care must be taken when using this workaround as there are semantic differences in the operation of the triggers, the use of "IN" and "EXISTS", and the use of referential integrity constraints.

This workaround is useful where the form of the constraint is more complex and cannot be phrased using referential integrity constraints. For example, if the application is such that the value in table T1 may be spaces or NULLs to indicate the absence of a value, the above triggers could easily be modified to allow for these semantics.

6.1.3 Dynamic Optimization Estimation Incorrect for Ranked Indices

The dynamic optimization process was incorrectly calculating the cost of scanning indices of type *SORTED RANKED*.

In the following example, the table being queried has the numbers one to one thousand in both fields. The different ranges used should result in a different estimated cost. However, in all cases the *ESTIM* phase computes the cost of scanning these indices as 680:

```
SQL> select * from t where f1 between 1 and 2 and f2 between 2 and 1000;
~S#0003
Leaf#01 FFirst T Card=1000
  BgrNdx1 T1 [1:1] Fan=17
  BgrNdx2 T2 [1:1] Fan=17
~E#0003.01(1) Estim   Ndx:Lev/Seps/DBKeys 1:34/0\680 2:34/0\680
~E#0003.01(1) BgrNdx1 EofData DBKeys=2  Fetches=2+0  RecsOut=1 #Bufs=1
~E#0003.01(1) FgrNdx  FFirst  DBKeys=1  Fetches=0+1  RecsOut=1`ABA
~E#0003.01(1) Fin      Buf      DBKeys=2  Fetches=0+0  RecsOut=1
          F1          F2
          2           2
1 row selected
SQL> select * from t where f1 between 2 and 1000 and f2 between 1 and 2;
~S#0004
Leaf#01 FFirst T Card=1000
  BgrNdx1 T1 [1:1] Fan=17
  BgrNdx2 T2 [1:1] Fan=17
~E#0004.01(1) Estim   Ndx:Lev/Seps/DBKeys 1:34/0\680 2:34/0\680
~E#0004.01(1) BgrNdx1 EofData DBKeys=999  Fetches=0+10  RecsOut=1 #Bufs=10
~E#0004.01(1) FgrNdx  FFirst  DBKeys=1  Fetches=0+11  RecsOut=1`ABA
~E#0004.01(1) Fin      Buf      DBKeys=999  Fetches=0+0  RecsOut=1
          F1          F2
          2           2
1 row selected
```

In the first example (query 3), the index T1 on field F1 is the correct index to use, as the key range is very small. In the second example (query 4), the index T2 on field F2 is the correct index to use. However, in both cases, the indices are costed the same so no index reordering takes place.

Even in this small example, significantly more work is being performed in query 4 as can be observed from the I/O counts.

This is a known problem in Oracle Rdb and it will be fixed in a future release.

The only workaround is to use indices of *TYPE IS SORTED* rather than of *TYPE IS SORTED RANKED*.

6.1.4 Running Rdb Applications With the VMS Heap Analyzer

When trying to debug an Rdb application under the OpenVMS Heap Analyzer (by defining LIBRTL as SYSS\$LIBRARY:LIBRTL_INSTRUMENTED), the software will not attach to the database, and returns

```
RDB-E-UNAVAILABLE, Oracle Rdb is not available on your system
```

as if RDB is not running.

To solve this problem, there are two executables that must be installed as known images:

```
$install add sys$share:librtl_instrumented
$install add sys$share:dgit$libshr12
```

The error is misleading. Since parts of Rdb are installed as privileged images, any shareable images it references, AND any images they, in turn, reference, must also be 'known'. By redirecting LIBRTL to SYS\$LIBRARY:LIBRTL_INSTRUMENTED, these extra images are referenced. If Rdb had directly referenced the new image, a more accurate error, such as:

```
%DCL-W-ACTIMAGE, error activating image xxxxx
```

would have been reported.

6.1.5 RMU/RECOVER/AREA Needs Area List

Bug 1778243

When doing an RMU/RECOVER/AREA, without specifying a list of area names, there will be a new version of the current active AIJ file created. This new version of the AIJ will have the next recovery sequence number. If a subsequent recovery is applied, an error is generated indicating that the original recovery sequence number cannot be found and the recovery will abort.

If a list of storage areas to be recovered is supplied, this behaviour does not occur and no new version of the journal is created. It is recommended as best practice to use a list of storage areas when recovering by area to avoid any subsequent confusion during recovery.

6.1.6 PAGE TRANSFER VIA MEMORY Disabled

Oracle internal testing has revealed that the "PAGE TRANSFER VIA MEMORY" option for global buffers is not as robust as is needed for the Mission Critical environments where Oracle Rdb is often deployed. This feature has been disabled in Oracle Rdb Version 7.0.xx. Oracle intends to re-enable this feature in a future Version 7.1 release.

6.1.7 RMU/VERIFY Reports PGSPAMENT or PGSPMCLST Errors

RMU/VERIFY may sometimes report PGSPAMENT or PGSPMCLST errors when verifying storage areas. These errors indicate that the Space Area Management ("SPAM") page fullness threshold for a particular data page does not match the actual space usage on the data page. For a further discussion of SPAM pages, consult the Oracle Rdb Guide to Database Maintenance.

In general, these errors will not cause any adverse affect on the operation of the database. There is potential for space on the data page to not be totally utilized, or for a small amount of extra I/O to be expended when searching for space in which to store new rows. But unless there are many of these errors then the impact should be negligible.

It is possible for these inconsistencies to be introduced by errors in the Oracle Rdb product. When those cases are discovered, Oracle Rdb is corrected to prevent the introduction of the inconsistencies. It is also possible for these errors to be introduced during the normal operation of the product. The following scenario can leave the SPAM pages inconsistent:

1. A process inserts a row on a page, and updates the threshold entry on the corresponding SPAM page to reflect the new space utilization of the data page. The data page and SPAM pages are not flushed to disk.
2. Another process notifies the first process that it would like to access the SPAM page being held by the process. The first process flushes the SPAM page changes to disk and releases the page. Note that it has not flushed the data page.
3. The first process then terminates abnormally (for example, from the DCL STOP/IDENTIFICATION command). Since that process never flushed the data page to disk, it never wrote the changes to the Recovery Unit Journal (RUJ) file. Since there were no changes in the RUJ file for that data page, then the Database Recovery ("DBR") process did not need to rollback any changes to the page. The SPAM page retains the threshold update change made above even though the data page was never flushed to disk.

While it would be possible to create mechanisms to ensure that SPAM pages do not become out of synch with their corresponding data pages, the performance impact would not be trivial. Since these errors are relatively rare and the impact is not significant, the introduction of these errors is considered to be part of the normal operation of the Oracle Rdb product. If it can be proven that the errors are not due to the scenario above then Oracle Product Support should be contacted.

PGSPAMENT and PGSPMCLST errors may be corrected by doing any one of the following operations:

- Recreate the database by performing:
 1. SQL EXPORT
 2. SQL DROP DATABASE
 3. SQL IMPORT
- Recreate the database by performing:
 1. RMU/BACKUP
 2. SQL DROP DATABASE
 3. RMU/RESTORE
- Repair the SPAM pages by using the RMU/REPAIR command. Note that the RMU/REPAIR command does not write its changes to an after-image journal (AIJ) file. Therefore, Oracle recommends that a full database backup be performed immediately after using the RMU/REPAIR command.

6.1.8 Behavior Change in 'With System Logical_Name Translation' Clause

The way logical name translation is performed when 'with system logical_name translation' is specified in the 'location' clause of the 'create function' or the 'create routine' statements has changed. This change occurred between OpenVMS VAX V5.5-2 and OpenVMS V7.1.

When 'with system logical_name translation' is specified, any logical name in the location string is expanded using only EXECUTIVE_MODE logical names. In OpenVMS VAX V5.5-2, the logical names are expanded from the SYSTEM logical name table only. In OpenVMS V7.1, the logical names are expanded from the first definition found when searching the logical name tables in (LNM\$FILE_DEV) order.

Thus, if a logical is only defined in the EXECUTIVE_MODE SYSTEM table (and in no other EXECUTIVE_MODE tables), then there will be no apparent change in behavior. However, if a logical name has been defined in the EXECUTIVE_MODE GROUP table and in the

EXECUTIVE_MODE SYSTEM table, then on OpenVMS VAX V5.5 the SYSTEM table translation will be used and on OpenVMS V7.1 the GROUP table translation will be used.

Oracle believes that this behavioral change is still in keeping with the secure intent of this clause for external routines. An OpenVMS user must have SYSNAM privilege to define an EXEC mode logical in any table. Therefore, it still provides a secure method of locating production sharable images for use by the Rdb server.

A future version of the Oracle Rdb SQL Reference manual will be reworded to remove the reference to the SYSTEM logical name table in the description. The keyword SECURE will be synonymous with SYSTEM in this context.

As an example, if the logical TEST_EXTRTN_1 is defined as:

```
$ show logical/access_mode=executive_mode test_extrtn_1
"TEST_EXTRTN_1" = "NOSUCHIMG9" (LNM$PROCESS_TABLE)
"TEST_EXTRTN_1" = "NOSUCHIMG9" (LNM$JOB_9D277AC0)
"TEST_EXTRTN_1" = "NOSUCHIMG9" (TEST$GROUP_LOGICALS)
"TEST_EXTRTN_1" = "DISK1:[TEST]EXTRTN.EXE" (LNM$SYSTEM_TABLE)
```

Then under OpenVMS VAX V5.5–2, TEST_EXTRTN_1 will be translated as "DISK1:[TEST]EXTRTN.EXE" whereas under OpenVMS V7.1 it will be translated as "NOSUCHIMG9".

6.1.9 Carry–Over Locks and NOWAIT Transactions Clarification

In NOWAIT transactions, the BLAST mechanism cannot be used. For the blocking user to receive the BLAST signal, the requesting user must request the locked resource with WAIT (which a NOWAIT transaction does not do). Oracle Rdb defines a resource called NOWAIT, which is used to indicate that a NOWAIT transaction has been started. When a NOWAIT transaction starts, the user requests the NOWAIT resource. All other database users hold a lock on the NOWAIT resource so that when the NOWAIT transaction starts, all other users are notified with a NOWAIT BLAST. The BLAST causes blocking users to release any carry–over locks. There can be a delay before the transactions with carry–over locks detect the presence of the NOWAIT transaction and release their carry–over locks. You can detect this condition by examining the stall messages. If the "Waiting for NOWAIT signal (CW)" stall message appears frequently, then the application is probably experiencing a decrease in performance and you should consider disabling the carry–over lock behavior.

6.1.10 Strict Partitioning May Scan Extra Partitions

When you use a WHERE clause with the less than (<) or greater than (>) operator and a value that is the same as the boundary value of a storage map, Oracle Rdb scans extra partitions. A boundary value is a value specified in the WITH LIMIT OF clause. The following example, executed while the logical name RDMS\$DEBUG_FLAGS is defined as "S", illustrates the behavior:

```
ATTACH 'FILENAME MF_PERSONNEL';
CREATE TABLE T1 (ID INTEGER, LAST_NAME CHAR(12), FIRST_NAME CHAR(12));
CREATE STORAGE MAP M FOR T1 PARTITIONING NOT UPDATABLE
STORE USING (ID)
IN EMPIDS_LOW WITH LIMIT OF (200)
```

Oracle® Rdb for OpenVMS

```
IN EMPIDS_MID WITH LIMIT OF (400)
OTHERWISE IN EMPIDS_OVER;
INSERT INTO T1 VALUES (150, 'Boney', 'MaryJean');
INSERT INTO T1 VALUES (350, 'Morley', 'Steven');
INSERT INTO T1 VALUES (300, 'Martinez', 'Nancy');
INSERT INTO T1 VALUES (450, 'Gentile', 'Russ');

SELECT * FROM T1 WHERE ID > 400;
Conjunct          Get          Retrieval sequentially of relation T1
Strict Partitioning: part          2          3
                   ID  LAST_NAME  FIRST_NAME
                   450  Gentile    Russ
1 row selected
```

In the previous example, partition 2 does not need to be scanned. This does not affect the correctness of the result. Users can avoid the extra scan by using values other than the boundary values.

6.1.11 Exclusive Access Transactions May Deadlock With RCS Process

If a table is frequently accessed by long running transactions that request READ/WRITE access reserving the table for EXCLUSIVE WRITE, and if the table has one or more indexes, you may experience deadlocks between the user process and the Row Cache Server (RCS) process.

There are at least three suggested workarounds to this problem:

1. Reserve the table for SHARED WRITE.
2. Close the database and disable row cache for the duration of the exclusive transaction.
3. Change the checkpoint interval for the RCS process to a time longer than the time required to complete the batch job and then trigger a checkpoint just before the batch job starts. Set the interval back to a smaller interval after the checkpoint completes.

6.1.12 Oracle Rdb and OpenVMS ODS-5 Volumes

The OpenVMS Version 7.2 release introduced an Extended File Specifications feature, which consists of two major components:

- ◆ A new, optional, volume structure, ODS-5, which provides support for file names that are longer and have a greater range of legal characters than in previous versions of OpenVMS.
- ◆ Support for "deep" directory trees.

ODS-5 was introduced primarily to provide enhanced file sharing capabilities for users of Advanced Server for OpenVMS 7.2 (formerly known as PATHWORKS for OpenVMS), as well as DCOM and JAVA applications.

In some cases, Oracle Rdb performs its own file and directory name parsing and explicitly requires ODS-2 (the traditional OpenVMS volume structure) file and directory name conventions to be followed. Because of this knowledge, Oracle does not support any Oracle Rdb database file components (including root files, storage area files, after image journal files, record cache backing store files, database backup files, after image journal backup files, etc.) that utilize any non-ODS-2 file naming features. For this reason, Oracle recommends that Oracle Rdb database components not be located on ODS-5 volumes.

Oracle does support Oracle Rdb database file components on ODS-5 volumes provided that all of these files and directories used by Oracle Rdb strictly follow the ODS-2 file and directory name conventions. In particular, all file names must be specified entirely in uppercase and "special" characters in file or directory names are forbidden.

6.1.13 Clarification of the USER Impersonation Provided by the Oracle Rdb Server

Bug 551240

In Oracle Rdb V6.1, a new feature was introduced which allowed a user to attach (or connect) to a database by providing a username (USER keyword) and a password (USING keyword). This functionality allows the Rdb Server to impersonate those users in two environments.

- ◆ Remote Database Access. When DECnet is used as the remote transport, the Rdb/Dispatch layer of Oracle Rdb uses the provided username and password, or proxy access to create a remote process which matches the named user. However, in a remote connection over TCP/IP, the RDBSERVER process is always logged into RDB\$REMOTE rather than a specified user account. In this case the Rdb Server impersonates the user by using the user's UIC (user identification code) during privilege checking. The UIC is assigned by the OpenVMS AUTHORIZE utility.
- ◆ SQL/Services database class services. When SQL/Services (possibly accessed by ODBC) accesses a database, it allows the user to logon to the database and the SQL/Services server then impersonates that user in the database.

When a database has access control established using OpenVMS rights identifiers, then access checking in these two environments does not work as expected. For example, if a user JONES was granted the rights identifier PAYROLL_ACCESS, then you would expect a table in the database with SELECT access granted to PAYROLL_ACCESS to be accessible to JONES. This does not currently work because the Rdb Server does not have the full OpenVMS security profile loaded, just the UIC. So only access granted to JONES is allowed.

This problem results in an error being reported such as the following from ODBC:

```
[Oracle][ODBC][Rdb]RDB-E-NO_PRIV privileged by database facility (#-1028)
```

This is currently a restriction in this release of Oracle Rdb. In the Rdb V7.1 release, support is provided to inherit the users full security profile into the database.

6.1.14 Index STORE Clause WITH LIMIT OF Not Enforced in Single Partition Map

Bug 413410

An index which has a STORE clause with a single WITH LIMIT OF clause and no OTHERWISE clause doesn't validate the inserted values against the high limit. Normally values beyond the last WITH LIMIT OF clause are rejected during INSERT and UPDATE statements.

Consider this example:

```

create table PTABLE (
  NR
    INTEGER,
  A
    CHAR (2));
create index NR_IDX
  on PTABLE (
  NR)
  type is HASHED
  store using (NR)
  in EMPIDS_LOW
  with limit of (10);

```

When a value is inserted for NR that exceeds the value 10, then an error such as "%RDMS-E-EXCMAPLIMIT, exceeded limit on last partition in storage map for NR_IDX" should be generated. However, this error is only reported if the index has two or more partitions.

A workaround for this problem is to create a CHECK constraint on the column to restrict the upper limit. e.g. CHECK (NR <= 10). This check constraint should be defined as NOT DEFERRABLE and will be solved using an index lookup.

This problem will be corrected in a future version of Oracle Rdb.

6.1.15 Unexpected NO_META_UPDATE Error Generated by DROP MODULE ... CASCADE When Attached by PATHNAME

Bug 755182

The SQL statement DROP MODULE ... CASCADE may sometimes generate an unexpected NO_META_UPDATE error. This occurs when the session attaches to a database by PATHNAME.

```

SQL> drop module m1 cascade;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-OBJ_INUSE, object "M1P1" is referenced by M2.M2P1 (usage: Procedure)
-RDMS-E-MODNOTDEL, module "M1" has not been deleted

```

This error occurs because the CASCADE option is ignored because the Oracle CDD/Repository does not support CASCADE. The workaround is to attach by FILENAME and perform the metadata operation.

In a future version of Oracle Rdb, an informational message will be issued describing the downgrade from CASCADE to RESTRICT in such cases.

6.1.16 Application and Oracle Rdb Both Using SYS\$HIBER

In application processes that use Oracle Rdb and the \$HIBER system service (possibly via RTL routines such as LIB\$WAIT), it is important that the application ensures that the event being waited for has actually occurred. Oracle Rdb uses \$HIBER/\$WAKE sequences for interprocess communications particularly when the ALS (AIJ Log Server) or the Row Cache features are enabled.

Oracle Rdb's use of the \$WAKE system service can interfere with other users of \$HIBER (such as the routine LIB\$WAIT) that do not check for event completion, possibly causing a \$HIBER to be unexpectedly resumed without waiting at all.

To avoid these situations, consider altering the application to use a code sequence that avoids continuing without a check for the operation (such as a delay or a timer firing) being complete.

The following pseudo-code shows one example of how a flag can be used to indicate that a timed-wait has completed correctly. The wait does not complete until the timer has actually fired and set TIMER_FLAG to TRUE. This code relies on ASTs being enabled.

```
ROUTINE TIMER_WAIT:
  BEGIN
    ! Clear the timer flag
    TIMER_FLAG = FALSE

    ! Schedule an AST for sometime in the future
    STAT = SYS$SETIMR (TIMADR = DELTATIME, ASTRTN = TIMER_AST)
    IF STAT <> SS$_NORMAL THEN LIB$SIGNAL (STAT)

    ! Hibernate. When the $HIBER completes, check to make
    ! sure that TIMER_FLAG is set indicating that the wait
    ! has finished.
    WHILE TIMER_FLAG = FALSE
      DO SYS$HIBER()
    END

ROUTINE TIMER_AST:
  BEGIN
    ! Set the flag indicating that the timer has expired
    TIMER_FLAG = TRUE

    ! Wake the main-line code
    STAT = SYS$WAKE ( )
    IF STAT <> SS$_NORMAL THEN LIB$SIGNAL (STAT)
  END
```

Starting with OpenVMS V7.1, the LIB\$WAIT routine has been enhanced via the FLAGS argument (with the LIB\$_NOWAKE flag set) to allow an alternate wait scheme (using the \$SYNCH system service) that can avoid potential problems with multiple code sequences using the \$HIBER system service. See the OpenVMS RTL Library (LIB\$) Manual for more information about the LIB\$WAIT routine.

6.1.17 IMPORT Unable to Import Some View Definitions

Bug 520651

View definitions that reference SQL functions, that is functions defined by the CREATE MODULE statement, cannot currently be imported by the SQL IMPORT statement. This is because the views are defined before the functions themselves exist.

The following example shows the errors from IMPORT.

```
IMPORTing view TVIEW
%SQL-F-NOVIERES, unable to import view TVIEW
```

Oracle® Rdb for OpenVMS

```
%RDB-E-NO_META_UPDATE, metadata update failed
-RDB-E-OBSOLETE_METADA, request references metadata objects that no
longer exist
-RDMS-E-RTNNEXTS, routine FORMAT_OUT does not exist in this database
%RDB-E-OBSOLETE_METADA, request references metadata objects that no
longer exist
-RDMS-F-TABNOTDEF, relation TVIEW is not defined in database
```

The following script can be used to demonstrate the problem.

```
create database filename badimp;
create table t (sex char);

create module TFORMAT
  language SQL

  function FORMAT_OUT (:s char)
  returns char(4);
  return (case :s
          when 'F' then 'Female'
          when 'M' then 'Male'
          else NULL
          end);
end module;

create view TVIEW (m_f) as
  select FORMAT_OUT (sex) from t;

commit;

export database filename badimp into exp;
drop database filename badimp;
import database from exp filename badimp;
```

This restriction will be lifted in a future release of Oracle Rdb. Currently the workaround is to save the view definitions and reapply them after the IMPORT completes.

This restriction does not apply to external functions, created using the CREATE FUNCTION statement, as these database objects are defined before tables and views.

6.1.18 AIJSERVER Privileges

For security reasons, the AIJSERVER account ("RDMAIJSERVER") is created with only NETMBX and TMPMBX privileges. These privileges are sufficient to start Hot Standby, in most cases.

However, for production Hot Standby systems, these privileges are not adequate to ensure continued replication in all environments and workload situations. Therefore, Oracle recommends that the DBA provide the following additional privileges for the AIJSERVER account:

- ◆ ALTPRI
This privilege allows the AIJSERVER to adjust its own priority to ensure adequate quorum (CPU utilization) to prompt message processing.
- ◆ PSWAPM
This privilege allows the AIJSERVER to enable and disable process swapping, also necessary to ensure prompt message processing.

- ◆ SETPRV
This privilege allows the AIJSERVER to temporarily set any additional privileges it may need to access the standby database or its server processes.
- ◆ SYSPRV
This privilege allows the AIJSERVER to access the standby database rootfile, if necessary.
- ◆ WORLD
This privilege allows the AIJSERVER to more accurately detect standby database server process failure and handle network failure more reliably.

6.1.19 Lock Remastering and Hot Standby

When using the Hot Standby feature, Oracle recommends that the VMS distributed lock manager resource tree be mastered on the standby node where Hot Standby is started. This can be using any of the following methods:

- ◆ Disable dynamic lock remastering. This can be done dynamically by setting the SYSGEN parameter PE1 to the value 1.
When using this option, be sure that Hot Standby is started on the node where the standby database is first opened.
- ◆ Increasing the LOCKDIRWT value for the LRS node higher than any other node in the same cluster. However, this is not a dynamic SYSGEN parameter, and a node re-boot is required.

Failure to prevent dynamic lock remastering may cause severe performance degradation for the standby database, which ultimately may be reflected by decreased master database transaction throughput.

6.1.20 RDB_SETUP Privilege Error

Rdb Web Agent V3.0 exposes a privilege problem with Rdb V7.0 and later. This will be fixed in the next Rdb 7.0 release.

The RDB_SETUP function fails with %RDB-E-NO_PRIV, privilege denied by database facility.

It appears that the only workaround is to give users DBADM privilege. Oracle Corporation does not recommend giving users the DBADM privilege.

6.1.21 Starting Hot Standby on Restored Standby Database May Corrupt Database

If a standby database is modified outside of Hot Standby, then backed up and restored, Hot Standby will appear to start up successfully but will corrupt the standby database. A subsequent query of the database will return unpredictable results, possibly in a bugcheck in DIOFETCH\$FETCH_ONE_LINE. When the standby database is restored from a backup of itself, the database is marked as unmodified. Therefore, Hot Standby cannot tell whether the database had been modified before the backup was taken.

WORKAROUND: None.

6.1.22 Restriction on Compound Statement Nesting Levels

The use of multiple nesting levels of compound statements such as CASE or IF–THEN–ELSE within multistatement procedures can result in excessive memory usage during the compile of the procedure. Virtual memory problems have been reported with 10 or 11 levels of nesting. The following example shows an outline of the type of nesting that can lead to this problem.

```

CREATE MODULE MY_MOD LANGUAGE SQL
PROCEDURE MY PROCEDURE
  ( PARAMETERS .....);

BEGIN
  DECLARE ....;

  SET :VARS = 0;

  SELECT .....;
  GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE;
  CASE :FLAG
    ! Case #1
    WHEN 100 THEN SET ...;
    WHEN -811 THEN SET ...;
    WHEN 0 THEN
      SET ...; SELECT ...;
      GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE;
      CASE :FLAG
        ! Case #2
        WHEN 0 THEN SET ...;
        WHEN -811 THEN SET ...;
        WHEN 100 THEN
          UPDATE...; SET ...;
          GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE;
          IF :FLAG= 100 THEN SET ...;           ! #1
          ELSE
            IF :FLAG < 0 THEN SET...;         ! #2
          ELSE
            DELETE ...
            GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE;
            IF :FLAG= 100 THEN SET...;         ! #3
            SET ...;
          ELSE
            IF :FLAG < 0 THEN SET...;         ! #4
          ELSE
            IF IN_CHAR_PARAM = 'S' THEN      ! #5
              UPDATE ...
              GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE;
              IF :FLAG= 100 THEN SET ...;     ! #6
              ELSE
                IF :FLAG < 0 THEN SET...;    ! #7
              END IF;                         ! #7
            END IF;                           ! #6
          END IF;                             ! #5
        END IF;
      END IF;
    END IF;
  IF :FLAG = 0 THEN                          ! #5
    UPDATE ...
    GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE;
    IF :FLAG= 100 THEN SET ...;             ! #6
    ELSE
      IF :FLAG < 0 THEN SET ...;           ! #7
    ELSE
      DELETE ...

```



```

GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE:
IF :FLAG= 100 THEN SET ...;          ! #8
ELSE
  IF :FLAG < 0 THEN SET ...;          ! #9
  ELSE
    DELETE ...;
    GET DIAGNOSTICS EXCEPTION 1 :FLAG = RETURNED_SQLCODE;
    IF :FLAG= 100 THEN SET ...;      ! #10
    SET ...;
    ELSE
      IF :FLAG < 0 THEN SET ...;     ! #11
      END IF; (11 end if's for #11 - #1)
    ELSE SET ...;
    END CASE;                          ! Case #2
  ELSE SET ...;
  END CASE;                              ! Case #1
END;
END MODULE;

```

Workaround: Reduce the complexity of the multistatement procedure. Use fewer levels of compound statement nesting by breaking the multistatement procedure into smaller procedures or by using the CALL statement to execute nested stored procedures.

6.1.23 Back Up All AIJ Journals Before Performing a Hot Standby Switchover Operation

Prior to performing a proper Hot Standby switchover operation from the old master database to the new master database (old standby database), be sure to back up ALL AIJ journals.

If you do not back up the AIJ journals on the old master database prior to switchover, they will be initialized by the Hot Standby startup operation, and you will not have a backup of those AIJ journals.

Failure to back up these journals may place your new master database at risk of not being able to be recovered, requiring another fail-over in the event of system failure.

6.1.24 Concurrent DDL and Read-Only Transaction on the Same Table Not Compatible

It is possible that a read-only transaction could generate a bugcheck at DIOBND\$FETCH_AIP_ENT + 1C4 if there is an active, uncommitted transaction that is making metadata changes to the same table. Analysis shows that the snapshot transaction is picking up stale metadata information. Depending on what metadata modifications are taking place, it is possible for metadata information to be removed from the system tables but still exist in the snapshot file. When the read-only transaction tries to use that information, it no longer exists and causes a bugcheck.

The following example shows the actions of the two transactions:

```

A:                                     B:
attach                                 attach
set transaction read write             set transaction read only

```

```
drop index emp_last_name  
  
select * from employees  
...bugcheck...
```

The only workaround is to avoid running the two transactions together.

6.1.25 Oracle Rdb and the SRM_CHECK Tool

The Alpha Architecture Reference Manual, Third Edition (AARM) describes strict rules for using interlocked memory instructions. The HP Alpha 21264 (EV6) processor and all future Alpha processors are more stringent than their predecessors in their requirement that these rules be followed. As a result, code that has worked in the past despite noncompliance may now fail when executed on systems featuring the new 21264 processor.

Oracle Rdb Release 7.0.3 supports the HP Alpha 21264 (EV6) processor. Oracle has performed extensive testing and analysis of the Rdb code to ensure that it is compliant with the rules for using interlocked memory instructions.

However, customers using the HP supplied SRM_CHECK tool may find that several of the Oracle Rdb images cause the tool to report potential alpha architecture violations. Although SRM_CHECK can normally identify a code section in an image by the section's attributes, it is possible for OpenVMS images to contain data sections with those same attributes. As a result, SRM_CHECK may scan data as if it were code, and occasionally, a block of data may look like a noncompliant code sequence. This is the case with the Oracle Rdb supplied images. There is no actual instruction stream violation.

However, customers must use the SRM_CHECK tool on their own application executable image files. It is possible that applications linked with very old version of Oracle Rdb (versions prior to Oracle Rdb Release 6.0–05) could have included illegal interlocked memory instruction sequences produced by very old versions of compilers. This code was included in the Oracle Rdb object library files for some very old versions of Oracle Rdb.

If errant instruction sequences are detected in the objects supplied by the Oracle Rdb object libraries, the correct action is to relink the application with a more-current version of Oracle Rdb.

Additional information about the HP Alpha 21264 (EV6) processor interlocked memory instructions issues is available at:

http://www.openvms.digital.com/openvms/21264_considerations.html

6.1.26 Oracle RMU Checksum_Verification Qualifier

The Oracle Rdb RMU BACKUP database backup command includes a Checksum_Verification qualifier.

Specifying Checksum_Verification requests that the RMU Backup command verify the checksum stored on each database page before it is backed up, thereby providing end-to-end error detection on the database I/O.

The Checksum_Verification qualifier uses additional CPU resources but can provide an extra measure of confidence in the quality of the data backed up. Use of the Checksum_Verification qualifier offers an additional level of data security and use of the Checksum_Verification qualifier permits Oracle RMU to detect the possibility that the data it is reading from these disks has only been partially updated.

Note, however, that if you specify the Nochecksum_Verification qualifier, and undetected corruptions exist in your database, the corruptions are included in your backup file and restored when you restore the backup file. Such a corruption might be difficult to recover from, especially if it is not detected until weeks or months after the restore operation is performed.

Oracle Corporation recommends that you use the Checksum_Verification qualifier with all database backup operations because of the improved data integrity this qualifier provides.

Unfortunately, due to an oversight, for versions of Oracle Rdb prior to Version 8.0, the default for online backups is the Nochecksum_Verification qualifier. When you do not specify the Checksum_Verification qualifier on all of your RMU database backup commands.

6.1.27 Do Not Use HYPERSORT with RMU/OPTIMIZE/AFTER_JOURNAL (Alpha)

OpenVMS Alpha V7.1 introduced the high-performance Sort/Merge utility (also known as HYPERSORT). This utility takes advantage of the Alpha architecture to provide better performance for most sort and merge operations.

The high-performance Sort/Merge utility supports a subset of the SOR routines. Unfortunately, the high-performance Sort/Merge utility does not support several of the interfaces used by the RMU/OPTIMIZE/AFTER_JOURNAL command. In addition, the high-performance Sort/Merge utility reports no error or warning when being called with the unsupported options used by the RMU/OPTIMIZE/AFTER_JOURNAL command.

For this reason, the use of the high-performance Sort/Merge utility is not supported for the RMU/OPTIMIZE/AFTER_JOURNAL command. Do not define the logical name SORTSHR to reference HYPERSORT.EXE.

6.1.28 Restriction on Using /NOONLINE with Hot Standby

When a user process is performing a read-only transaction on a standby database, an attempt to start replication on the standby database with the /NOONLINE qualifier will fail with the following error, and the database will be closed cluster-wide:

```
%RDMS-F-OPERCLOSE, database operator requested database shutdown
```

In a previous release, the following error was returned and the process doing the read-only transaction was not affected:

```
%RDMS-F-STBYDBINUSE, standby database cannot be exclusively accessed for replication
```

As a workaround, if exclusive access is necessary to the standby database, terminate any user processes before starting replication with the /NOONLINE qualifier.

This restriction is due to another bug fix and will be lifted in a future release of Oracle Rdb.

6.1.29 SELECT Query May Bugcheck with PSII2SCANGETNEXTBBCDUPLICATE Error

Bug 683916

A bugcheck could occur when a ranked B–tree index is used in a query after a database has been upgraded to Release 7.0.1.3. This is a result of index corruption that was introduced in previous versions of Oracle Rdb. This corruption has been fixed and indexes created using Release 7.0.1.3 will not be impacted.

As a workaround, delete the affected index and re–create it under Oracle Rdb Release 7.0.1.3 or later.

6.1.30 DBAPack for Windows 3.1 is Deprecated

Oracle Enterprise Manager DBAPack will no longer be supported for use on Windows 3.1.

6.1.31 Determining Mode for SQL Non–Stored Procedures

Bug 506464.

Although stored procedures allow parameters to be defined with the modes IN, OUT, and INOUT, there is no similar mechanism provided for SQL module language or SQL precompiled procedures. However, SQL still associates a mode with a parameter using the following rules:

Any parameter which is the target of an assignment is considered an OUT parameter. Assignments consist of the following:

- ◆ The parameter is assigned a value with the SET or GET DIAGNOSTICS statement. For example:

```
set :p1 = 0;
get diagnostics :p2 = TRANSACTION_ACTIVE;
```

- ◆ The parameter is assigned a value with the INTO clause of an INSERT, UPDATE, or SELECT statement. For example:

```
insert into T (col1, col2)
  values (...)
  returning dbkey into :p1;

update accounts
  set account_balance = account_balance + :amount
  where account_number = :p1
  returning account_balance
  into :current_balance;

select last_name
```

Oracle® Rdb for OpenVMS

```
into :p1
from employees
where employee_id = '00164';
```

- ◆ The parameter is passed on a CALL statement as an OUT or INOUT argument. For example:

```
begin
call GET_CURRENT_BALANCE (:p1);
end;
```

Any parameter that is the source for a query is considered an IN parameter. Query references include:

- ◆ The parameter appears in the SELECT list, WHERE or HAVING clauses of a SELECT, or DELETE statement. For example:

```
select :p1 || last_name, count(*)
from T
where last_name like 'Smith%'
group by last_name
having count(*) > :p2;
```

```
delete from T
where posting_date < :p1;
```

- ◆ The parameter appears on the right side of the assignment in a SET statement or SET clause of an UPDATE statement. For example:

```
set :p1 = (select avg(salary)
from T
where department = :p2);
update T
set coll = :p1
where ...;
```

- ◆ The parameter is used to provide a value to a column in an INSERT statement. For example:

```
insert into T (coll, col2)
values (:p1, :p2);
```

- ◆ The parameter is referenced by an expression in a TRACE, CASE, IF/ELSEIF, WHILE statement, or by the DEFAULT clause of a variable declaration. For example:

```
begin
declare :v integer default :p1;
DO_LOOP:
while :p2 > :p1
loop
if :p1 is null then
leave DO_LOOP;
end if;
set :p2 = :p2 + 1;
...;
trace 'Loop at ', :p2;
end loop;
end;
```

- ◆ The parameter is passed on a CALL statement as an INOUT or IN argument. For example:

```
begin
call SET_LINE_SPEED (:p1);
end;
```

SQL only copies values from the client (application parameters) to the procedure running in the database server if it is marked as either an IN or INOUT parameter. SQL only returns values from the

server to the client application parameter variables if the parameter is an OUT or INOUT parameter.

If a parameter is considered an OUT only parameter, then it must be assigned a value within the procedure, otherwise the result returned to the application is considered undefined. This could occur if the parameter is used within a conditional statement such as CASE or IF/ELSEIF. In the following example, the value returned by :p2 would be undefined if :p1 were negative or zero:

```
begin
if :p1 > 0 then
    set :p2 = (select count(*)
              from T
              where coll = :p1);
end if;
end;
```

It is the responsibility of the application programmer to ensure that the parameter is correctly assigned values within the procedure. A workaround is to either explicitly initialize the OUT parameter, or make it an INOUT parameter. For example:

```
begin
if :p1 > 0 then
    set :p2 = (select count(*)
              from T
              where coll = :p1);
elseif :p2 is null then
    begin
    end;
end if;
end;
```

The empty statement will include a reference to the parameter to make it an IN parameter as well as an OUT parameter.

6.1.32 DROP TABLE CASCADE Results in %RDB-E-NO_META_UPDATE Error

An error could result when a DROP TABLE CASCADE statement is issued. This occurs when the following conditions apply:

- ◆ A table is created with an index defined on the table.
- ◆ A storage map is created with a placement via index.
- ◆ The storage map is a vertical record partition storage map with two or more STORE COLUMNS clauses.

The error message given is %RDB-E-NO_META_UPDATE, metadata update failed.

The following example shows a table, index, and storage map definition followed by a DROP TABLE CASCADE statement and the resulting error message:

```
SQL> CREATE TABLE VRP_TABLE ( ID INT, ID2 INT);
SQL> COMMIT;
SQL> CREATE UNIQUE INDEX VRP_IDX ON VRP_TABLE (ID)
SQL> STORE IN EMPIDS_LOW;
SQL> COMMIT;
SQL> CREATE STORAGE MAP VRP_MAP
```

Oracle® Rdb for OpenVMS

```
cont> FOR VRP_TABLE
cont> PLACEMENT VIA INDEX VRP_IDX
cont> ENABLE COMPRESSION
cont> STORE COLUMNS (ID)
cont> IN EMPIDS_LOW
cont> STORE COLUMNS (ID2)
cont> IN EMPIDS_MID;
SQL> COMMIT;
SQL>
SQL> DROP TABLE VRP_TABLE CASCADE;
SQL> -- Index VRP_IDX is also being dropped.
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-WISH_LIST, feature not implemented yet
-RDMS-E-VRPINVALID, invalid operation for storage map "VRP_MAP"
```

The workaround to this problem is to first delete the storage map, and then delete the table using the CASCADE option. The following example shows the workaround. The SHOW statement indicates that the table, index, and storage map were deleted:

```
SQL> DROP STORAGE MAP VRP_MAP;
SQL> DROP TABLE VRP_TABLE CASCADE;
SQL> -- Index VRP_IDX is also being dropped.
SQL> COMMIT;
SQL> SHOW TABLE VRP_TABLE
No tables found
SQL> SHOW INDEX VRP_IDX
No indexes found
SQL> SHOW STORAGE MAP VRP_MAP
No Storage Maps Found
```

This problem will be corrected in a future version of Oracle Rdb.

6.1.33 Bugcheck Dump Files with Exceptions at COSI_CHF_SIGNAL

In certain situations, Oracle Rdb bugcheck dump files will indicate an exception at COSI_CHF_SIGNAL. This location is, however, not the address of the actual exception. The actual exception occurred at the previous call frame on the stack (the one listed as the next "Saved PC" after the exception).

For example, consider the following bugcheck file stack information:

```
$ SEARCH RDSBUGCHK.DMP "EXCEPTION", "SAVED PC", "-F-", "-E-"

***** Exception at 00EFA828 : COSI_CHF_SIGNAL + 00000140
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 00C386F0 : PSIINDEX2JOINSCR + 00000318
Saved PC = 00C0BE6C : PSII2BALANCE + 0000105C
Saved PC = 00C0F4D4 : PSII2INSERTT + 000005CC
Saved PC = 00C10640 : PSII2INSERTTREE + 000001A0
.
.
.
```

In this example, the exception actually occurred at PSIINDEX2JOINSCR offset 00000318. If you have a bugcheck dump with an exception at COSI_CHF_SIGNAL, it is important to note the next

"Saved PC" because it will be needed when working with Oracle Rdb Support Services.

6.1.34 Interruptions Possible when Using Multistatement or Stored Procedures

Long running multistatement or stored procedures can cause other users in the database to be interrupted by holding resources needed by those other users. Some resources obtained by the execution of a multistatement or stored procedure will not be released until the multistatement or stored procedure finishes. This problem can be encountered even if the statement contains COMMIT or ROLLBACK statements.

The following example demonstrates the problem. The first session enters an endless loop; the second session attempts to backup the database, but it is permanently interrupted:

Session 1

```
SQL> ATTACH 'FILE MF_PERSONNEL';
SQL> CREATE FUNCTION LIB$WAIT (IN REAL BY REFERENCE)
cont> RETURNS INT;
cont> EXTERNAL NAME LIB$WAIT
cont> LOCATION 'SYS$SHARE:LIBRTL.EXE'
cont> LANGUAGE GENERAL
cont> GENERAL PARAMETER STYLE
cont> VARIANT;
SQL> COMMIT;
SQL> EXIT;
```

```
$ SQL
SQL> ATTACH 'FILE MF_PERSONNEL';
SQL> BEGIN
cont> DECLARE :LAST_NAME LAST_NAME_DOM;
cont> DECLARE :WAIT_STATUS INTEGER;
cont> LOOP
cont> SELECT LAST_NAME INTO :LAST_NAME
cont> FROM EMPLOYEES WHERE EMPLOYEE_ID = '00164';
cont> ROLLBACK;
cont> SET :WAIT_STATUS = LIB$WAIT (5.0);
cont> SET TRANSACTION READ ONLY;
cont> END LOOP;
cont> END;
```

Session 2

```
$ RMU/BACKUP/LOG/ONLINE MF_PERSONNEL MF_PERSONNEL
```

From a third session we can see that the backup process is waiting for a lock held in the first session:

```
$ RMU/SHOW LOCKS /MODE=BLOCKING MF_PERSONNEL
=====
SHOW LOCKS/BLOCKING Information
=====
-----
Resource: nowait signal
```


Oracle® Rdb for OpenVMS

	ProcessID	Process Name	Lock ID	System ID	Requested	Granted
	-----	-----	-----	-----	-----	-----
Waiting:	20204383	RMU BACKUP.....	5600A476	00010001	CW	NL
Blocker:	2020437B	SQL.....	3B00A35C	00010001	PR	PR

\$

There is no workaround for this restriction. When the multistatement or stored procedure finishes execution, the resources needed by other processes will be released.

6.1.35 Row Cache Not Allowed on Standby Database While Hot Standby Replication Is Active

The row cache feature may not be active on a Hot Standby database while replication is taking place. The Hot Standby feature will not start if row cache is active on the standby database.

This restriction exists because rows in the row cache are accessed using logical dbkeys. However, information transferred to the Hot Standby database from the after-image journal facility only contains physical dbkeys. Because there is no way to maintain rows in the cache using the Hot Standby processing, the row cache must be disabled on the standby database when the standby database is open and replication is active. The master database is not affected; the row cache feature and the Hot Standby feature may be used together on a master database.

The row cache feature should be identically configured on the master and standby databases in the event failover occurs, but the row cache feature must not be activated on the standby database until it becomes the master.

A new command qualifier, `ROW_CACHE=DISABLED`, has been added to the `RMU/OPEN` command to disable the row cache feature on the standby database. To open the Hot Standby database prior to starting replication, use the `ROW_CACHE=DISABLED` qualifier on the `RMU/OPEN` command.

6.1.36 Hot Standby Replication Waits when Starting if Read-Only Transactions Running

Hot Standby replication will wait to start if there are read-only (snapshot) transactions running on the standby database. The log roll-forward server (LRS) will wait until the read-only transactions commit, and then replication will continue.

This is an existing restriction of the Hot Standby software. This release note is intended to complement the Hot Standby documentation.

6.1.37 Error when Using the `SYS$LIBRARY:SQL_FUNCTIONS70.SQL` Oracle Functions Script

If your programming environment is not set up correctly, you may encounter problems running the `SYS$LIBRARY:SQL_FUNCTIONS70.SQL` script used to set up the Oracle7 functions being supplied with Oracle Rdb.

The following example shows the error:

```
%RDB-E-EXTFUN_FAIL, external routine failed to compile or execute successfully
-RDMS-E-INVRTNUSE, routine RDB$ORACLE_SQLFUNC_INTRO can not be used, image
"SQL$FUNCTIONS" not activated
-RDMS-I-TEXT, Error activating image
DISK:[DIR]SQL$FUNCTIONS.;; File not found
```

To resolve this problem, use the @SYS\$LIBRARY:RDB\$SETVER to set up the appropriate logical names. This will be necessary for programs that use the functions as well.

In a standard environment, use the command shown in the following example:

```
$ @SYS$LIBRARY:RDB$SETVER S
```

In a multiversion environment, use the command shown in the following example:

```
$ @SYS$LIBRARY:RDB$SETVER 70
```

6.1.38 DEC C and Use of the /STANDARD Switch

Bug 394451

The SQL\$PRE compiler examines the system to know which dialect of C to generate. That default can be overwritten by using the /CC=[DECC/VAXC] switch. The /STANDARD switch should not be used to choose the dialect of C.

Support for DEC C was added to the product with V6.0 and this note is meant to clarify that support, not to indicate a change. It is possible to use /STANDARD=RELAXED_ANSI89 or /STANDARD=VAXC correctly, but this is not recommended.

The following example shows both the right and wrong way to compile an Oracle Rdb SQL program. Assume a symbol SQL\$PRE has been defined, and DEC C is the default C compiler on the system:

```
$ SQL$PRE/CC ! This is correct.
$ SQL$PRE/CC=DECC ! This is correct.
$ SQL$PRE/CC=VAXC ! This is correct.

$ SQL$PRE/CC/STANDARD=VAXC ! This is incorrect.
```

Notice that the /STANDARD switch has other options in addition to RELAXED_ANSI89 and :VAX C. Those are also not supported.

6.1.39 Excessive Process Page Faults and Other Performance Considerations During Oracle Rdb Sorts

Excessive hard or soft page faulting can be a limiting factor of process performance. Sometimes this page faulting occurs during Oracle Rdb sort operations. This note describes how page faulting can occur and some ways to help control, or at least understand, it.

One factor contributing to Oracle Rdb process page faulting is sorting operations. Common causes of sorts include the SQL GROUP BY, ORDER BY, UNION, and DISTINCT clauses specified for query and index creation operations. Defining the logical name RDMS\$DEBUG_FLAGS to "RS" can help determine when Oracle Rdb sort operations are occurring and to display the sort keys and statistics.

Oracle Rdb includes its own copy of the OpenVMS SORT32 code within the Oracle Rdb images and does not generally call the routines in the OpenVMS run-time library. A copy of the SORT32 code is used to provide stability between versions of Oracle Rdb and OpenVMS and because Oracle Rdb calls the sort routines from executive processor mode which is difficult to do using the SORT32 sharable image. Database import and RMU load operations call the OpenVMS sort run-time library.

At the beginning of a sort operation, the sort code allocates some memory for working space. The sort code uses this space for buffers, in-memory copies of the data, and sorting trees.

Sort code does not directly consider the process quotas or parameters when allocating memory. The effects of WSQUOTA and WSEXTENT are indirect. At the beginning of each sort operation, the sort code attempts to adjust the process' working set to the maximum possible size using the \$ADJWSL system service specifying a requested working set limit of %X7FFFFFFF pages (the maximum possible). Sort code then uses a value of 75% of the returned working set for virtual memory scratch space. The scratch space is then initialized and the sort begins.

The initialization of the scratch space generally causes page faults to access the pages newly added to the working set. Pages that were in the working set already may be faulted out as new pages are faulted in. Once the sort operation completes, the pages that may have been faulted out of the working set are likely to be faulted back into the working set.

When a process' working set is limited by the working set quota (WSQUOTA) parameter and the working set extent (WSEXTENT) parameter is a much larger value, the first call to the sort routines can cause many page faults as the working set grows. Using a value of WSEXTENT that is closer to WSQUOTA can help reduce the impact of this case.

With some OpenVMS versions, AUTOGEN sets the SYSGEN parameter PQL_MWSEXTENT equal to the WSMAX parameter. This means that all processes on the system end up with WSEXTENT the same as WSMAX. Because WSMAX might be quite high, sorting might result in excessive page faulting. You may want to explicitly set PQL_MWSEXTENT to a lower value if this is the case on your system.

Sort work files are another factor to consider when tuning Oracle Rdb sort operations. When the operation cannot be done in available memory, sort code will use temporary disk files to hold the data as it is being sorted. The *Oracle Rdb Guide to Performance and Tuning* contains more detailed information about sort work files.

The logical name RDMS\$BIND_SORT_WORKFILES specifies how many work files sort code is to use if work files are required. The default is 2, and the maximum number is 10. The work files can be individually controlled by the SORTWORKn logical names (where n is from 0 through 9). You can increase the efficiency of sort operations by assigning the location of the temporary sort work files to different disks. These assignments are made by using up to 10 logical names, SORTWORK0 through

SORTWORK9.

Normally, sort code places work files in the user's SYSSCRATCH directory. By default, SYSSCRATCH is the same device and directory as the SYS\$LOGIN location. Spreading the I/O load over many disks improves efficiency as well as performance by taking advantage of the system resources and helps prevent disk I/O bottlenecks. Specifying that a user's work files will reside on separate disks permits overlap of the sort read/write cycle. You may also encounter cases where insufficient space exists on the SYSSCRATCH disk device, such as when Oracle Rdb builds indexes for a very large table. Using the SORTWORK0 through SORTWORK9 logical names can help you avoid this problem.

Note that sort code uses the work files for different sorted runs, and then merges the sorted runs into larger groups. If the source data is mostly sorted, then not every sort work file may need to be accessed. This is a possible source of confusion because even with 10 sort work files, it is possible to exceed the capacity of the first sort file, and the sort operation will fail never having accessed the remaining 9 sort work files.

Note that the logical names RDMS\$BIND_WORK_VM and RDMS\$BIND_WORK_FILE do not affect or control the operation of sort. These logical names are used to control other temporary space allocations within Oracle Rdb.

6.1.40 Performance Monitor Column Mislabeled

The File IO Overview statistics screen, in the Rdb Performance Monitor, contains a column labeled Pages Checked. The column should be labeled Pages Discarded to correctly reflect the statistic displayed.

6.1.41 Restriction Using Backup Files Created Later than Oracle Rdb Release 7.0.1

Bug 521583

Backup files created using Oracle Rdb releases later than 7.0.1 cannot be restored using Oracle Rdb Release 7.0.1. To fix a problem in a previous release, some internal backup file data structures were changed. These changes are not backward compatible with Oracle Rdb Release 7.0.1.

If you restore the database using such a backup file, then any attempt to access the restored database may result in unpredictable behavior, even though a verify operation may indicate no problems.

There is no workaround to this problem. For this reason, Oracle Corporation strongly recommends performing a full and complete backup both before and after the upgrade from Release 7.0.1 to later releases of Oracle Rdb.

6.1.42 RMU Backup Operations and Tape Drive Types

When using more than one tape drive for an RMU backup operation, all the tape drives must be of the same type. For example, all the tape drives must be either TA90s or TZ87s or TK50s. Using different tape drive types (one TK50 and one TA90) for a single database backup operation may make database restoration difficult or impossible.

Oracle RMU attempts to prevent using different tape drive densities during a backup operation, but is not able to detect all invalid cases and expects that all tape drives for a backup are of the same type.

As long as all the tapes used during a backup operation can be read by the same type of tape drive during a restore operation, the backup is likely to be valid. This may be the case, for example, when using a TA90 and a TA90E.

Oracle recommends that, on a regular basis, you test your backup and recovery procedures and environment using a test system. You should restore the databases and then recover them using AIJs to simulate failure recovery of the production system.

Consult the *Oracle Rdb Guide to Database Maintenance*, the *Oracle Rdb Guide to Database Design and Definition*, and the *Oracle RMU Reference Manual* for additional information about Oracle Rdb backup and restore operations.

6.1.43 Use of Oracle Rdb from Shared Images

Bug 470946

If code in the image initialization routine of a shared image makes any calls into Oracle Rdb, through SQL or any other means, access violations or other unexpected behavior may occur if Oracle Rdb's images have not had a chance to do their own initialization.

To avoid this problem, applications must do one of the following:

- ◆ Do not make Oracle Rdb calls from the initialization routines of shared images.
- ◆ Link in such a way that the RDBSHR.EXE image initializes first. This can be done by placing the reference to RDBSHR.EXE and any other Oracle Rdb shared images last in the linker options file.

6.1.44 Restriction Added for CREATE STORAGE MAP on Table with Data

Oracle Rdb Release 7.0 added support that allows a storage map to be added to an existing table which contains data. The restrictions listed for Oracle Rdb Release 7.0 were:

- ◆ The storage map must be a simple map that references only the default storage area and represents the current (default) mapping for the table. The default storage area is either RDB\$SYSTEM or the area name provided by the CREATE DATABASE...DEFAULT STORAGE AREA clause.
- ◆ The new map cannot change THRESHOLDS or COMPRESSION for the table, nor can it use the PLACEMENT VIA INDEX clause. It can only contain one area and cannot be vertically partitioned. This new map simply describes the mapping as it exists by default for the table.

This release of Rdb adds the additional restriction that the storage map may not include a WITH LIMIT clause for the storage area. The following example shows the reported error:

```
SQL> CREATE TABLE MAP_TEST1 (A INTEGER, B CHAR(10));
SQL> CREATE INDEX MAP_TEST1_INDEX ON MAP_TEST1 (A);
SQL> INSERT INTO MAP_TEST1 (A, B) VALUES (3, 'Third');
1 row inserted
```

```
SQL> CREATE STORAGE MAP MAP_TEST1_MAP FOR MAP_TEST1
cont> STORE USING (A) IN RDB$SYSTEM
cont> WITH LIMIT OF (10); -- can't use WITH LIMIT clause
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-RELNOTEEMPTY, table "MAP_TEST1" has data in it
-RDMS-E-NOCMPLEXMAP, can not use complex map for non-empty table
```

6.1.45 Oracle Rdb Workload Collection Can Stop Hot Standby Replication

If you are replicating your Oracle Rdb database using the Oracle Hot Standby option, you must not use the workload collection option. By default, workload collection is disabled. However, if you enabled workload collection, you must disable it on the master database prior to performing a backup operation on that master database if it will be used to create the standby database for replication purposes. If you do not disable workload collection, it could write workload information to the standby database and prevent replication operations from occurring.

The workaround included at the end of this section describes how to disable workload collection on the master database and allow the Hot Standby software to propagate the change to the standby database automatically during replication operations.

Background Information

By default, workload collection and cardinality collection are automatically disabled when Hot Standby replication operations are occurring on the standby database. However, if replication stops (even for a brief network failure), Oracle Rdb potentially can start a read/write transaction on the standby database to write workload collection information. Then, because the standby database is no longer synchronized transactionally with the master database, replication operations cannot restart.

Note

The Oracle Rdb optimizer can update workload collection information in the RDB\$WORKLOAD system table even though the standby database is opened exclusively for read-only queries. A read/write transaction is started during the disconnection from the standby database to flush the workload and cardinality statistics to the system tables.

If the standby database is modified, you receive the following messages when you try to restart Hot Standby replication operations:

```
%RDMS-F-DBMODIFIED, database has been modified; AIJ roll-forward not possible
%RMU-F-FATALRDB, Fatal error while accessing Oracle Rdb.
```

Workaround

To work around this problem, perform the following:

- ◆ On the master database, disable workload collection using the SQL clause `WORKLOAD COLLECTION IS DISABLED` on the `ALTER DATABASE` statement. For example:

Oracle® Rdb for OpenVMS

```
SQL> ALTER DATABASE FILE mf_personnel  
cont> WORKLOAD COLLECTION IS DISABLED;
```

This change is propagated to the standby database automatically when you restore the standby database and restart replication operations. Note that, by default, the workload collection feature is disabled. You need to disable workload collection only if you previously enabled workload collection with the `WORKLOAD COLLECTION IS ENABLED` clause.

- ◆ On the standby database, include the `Transaction_Mode` qualifier on the `RMU/Restore` command when you restore the standby database. You should set this qualifier to `read-only` to prevent modifications to the standby database when replication operations are not active. The following example shows the `Transaction_Mode` qualifier used in a typical `RMU/Restore` command:

```
$ RMU/RESTORE /TRANSACTION_MODE=READ_ONLY  
/NOCCD  
/NOLOG  
/ROOT=DISK1:[DIR]standby_personnel.rdb  
/AIJ_OPT=aij_opt.dat  
DISK1:[DIR]standby_personnel.rbf
```

If, in the future, you fail over processing to the standby database (so that the standby database becomes the master database), you can re-enable updates to the "new" master database. For example, to re-enable updates, use the `SQL` statement `ALTER DATABASE` and include the `SET TRANSACTION MODES (ALL)` clause. The following example shows this statement used on the new master database:

```
SQL> ALTER DATABASE FILE mf_personnel  
cont> SET TRANSACTION MODES (ALL);
```

6.1.46 RMU Convert Command and System Tables

When the `RMU Convert` command converts a database from a previous version to Oracle Rdb V7.0 or higher, it sets the `RDB$CREATED` and `RDB$LAST ALTERED` columns to the timestamp of the convert operation.

The `RDB$xxx_CREATOR` columns are set to the current user name (which is space filled) of the converter. Here `xxx` represents the object name, such as in `RDB$TRIGGER_CREATOR`.

The `RMU Convert` command also creates the new index on `RDB$TRANSFER_RELATIONS` if the database is transfer enabled.

6.1.47 Converting Single-File Databases

Because of a substantial increase in the database root file information for Release 7.0, you should ensure that you have adequate disk space before you use the `RMU Convert` command with single-file databases and Release 7.0 or higher.

The size of the database root file of any given database will increase a minimum of 13 blocks and a maximum of 597 blocks. The actual increase depends mostly on the maximum number of users specified for the database.

6.1.48 Restriction when Adding Storage Areas with Users Attached to Database

If you try to interactively add a new storage area where the page size is smaller than the smallest existing page size and the database has been manually opened or users are active, the add operation fails with the following error:

```
%RDB-F-SYS_REQUEST, error from system services request
-RDMS-F-FILACCERR, error opening database root DKA0:[RDB]TEST.RDB;1
-SYSTEM-W-ACCONFLICT, file access conflict
```

You can make this change only when no users are attached to the database and, if the database is set to OPEN IS MANUAL, the database is closed. Several internal Oracle Rdb data structures are based on the minimum page size, and these structures cannot be resized if users are attached to the database.

Furthermore, because this particular change is not recorded in the AIJ file, any recovery scenario will fail. Note also that if you use .aij files, you must backup the database and restart after-image journaling because this change invalidates the current AIJ recovery.

6.1.49 Support for Single-File Databases to be Dropped in a Future Release

Oracle Rdb currently supports both single-file and multifile databases on OpenVMS. However, single-file databases will not be supported in a future release of Oracle Rdb. At that time, Oracle Rdb will provide the means to easily convert single-file databases to multifile databases.

Oracle recommends that users with single-file databases perform the following actions:

- ◆ Use the Oracle RMU commands, such as Backup and Restore, to make copies, back up, or move single-file databases. Do not use operating system commands to copy, back up, or move databases.
- ◆ Create new databases as multifile databases even though single-file databases are supported in Oracle Rdb release 6.1 and release 7.0.

6.1.50 DECdtm Log Stalls

Resource managers using the DECdtm services can sometimes suddenly stop being able to commit transactions. If Oracle Rdb is installed and transactions are being run, an RMU Show command on the affected database will show transactions as being "stalled, waiting to commit".

Refer to the DECdtm documentation and release notes for information on symptoms, fixes, and workarounds for this problem. One workaround, for OpenVMS V5.5-x, is provided here.

On the affected node while the log stall is in progress, type the following command from a privileged account:

```
$ MCR LMCP SET NOTIMEZONE
```


Oracle® Rdb for OpenVMS

This should force the log to restart.

This stall occurs only when a particular bit in a pointer field becomes set. To see the value of the pointer field, enter the following command from a privileged account (where <nodename> is the SCS node name of the node in question).

```
$ MCR LMCP DUMP/ACTIVE/NOFORM SYSTEM$<nodename>
```

This command displays output similar to the following:

```
Dump of transaction log SYS$COMMON:[SYSEXE]SYSTEM$<nodename>.LM$JOURNAL;1
End of file block 4002 / Allocated 4002
Log Version 1.0
Transaction log UID:    29551FC0-CBB7-11CC-8001-AA000400B7A5
Penultimate Checkpoint: 000013FD4479 0079
Last Checkpoint:       000013FD4479 0079
```

Total of 2 transactions active, 0 prepared and 2 committed.

The stall will occur when bit 31 of the checkpoint address becomes set, as this excerpt from the previous example shows:

```
Last Checkpoint:       000013FD4479 0079
                        ^
                        |
```

When the number indicated in the example becomes 8, the log will stall. Check this number and observe how quickly it grows. When it is at 7FFF, frequently use the following command:

```
$ MCR LMCP SHOW LOG /CURRENT
```

If this command shows a stall in progress, use the workaround to restart the log.

See your HP Computer Corporation representative for information about patches to DECdtm.

6.1.51 Cannot Run Distributed Transactions on Systems with DECnet/OSI and OpenVMS Alpha Version 6.1 or OpenVMS VAX Version 6.0

If you have DECnet/OSI installed on a system with OpenVMS Alpha Version 6.1 or OpenVMS VAX Version 6.0, you cannot run Oracle Rdb operations that require the two-phase commit protocol. The two-phase commit protocol guarantees that if one operation in a distributed transaction cannot be completed, none of the operations is completed.

If you have DECnet/OSI installed on a system running OpenVMS VAX Version 6.1 or higher or OpenVMS Alpha Version 6.2 or higher, you can run Oracle Rdb operations that require the two-phase commit protocol.

For more information about the two-phase commit protocol, see the *Oracle Rdb Guide to Distributed Transactions*.

6.1.52 Multiblock Page Writes May Require Restore Operation

If a node fails while a multiblock page is being written to disk, the page in the disk becomes inconsistent and is detected immediately during failover. (Failover is the recovery of an application by restarting it on another computer.) The problem is rare and occurs because only single-block I/O operations are guaranteed by OpenVMS to be written atomically. This problem has never been reported by any customer and was detected only during stress tests in our labs.

Correct the page by an area-level restore operation. Database integrity is not compromised, but the affected area will not be available until the restore operation completes.

A future release of Oracle Rdb will provide a solution that guarantees multiblock atomic write operations. Cluster failovers will automatically cause the recovery of multiblock pages, and no manual intervention will be required.

6.1.53 Replication Option Copy Processes Do Not Process Database Pages Ahead of an Application

When a group of copy processes initiated by the Replication Option (formerly Data Distributor) begins running after an application has begun modifying the database, the copy processes will catch up to the application and will not be able to process database pages that are logically ahead of the application in the RDB\$CHANGES system table. The copy processes all align waiting for the same database page and do not move on until the application has released it. The performance of each copy process degrades because it is being paced by the application.

When a copy process completes updates to its respective remote database, it updates the RDB\$TRANSFERS system table and then tries to delete any RDB\$CHANGES rows not needed by any transfers. During this process, the RDB\$CHANGES table cannot be updated by any application process, holding up any database updates until the deletion process is complete. The application stalls while waiting for the RDB\$CHANGES table. The resulting contention for RDB\$CHANGES SPAM pages and data pages severely impacts performance throughput, requiring user intervention with normal processing.

This is a known restriction in Release 4.0 and higher. Oracle Rdb uses page locks as latches. These latches are held only for the duration of an action on the page and not to the end of transaction. The page locks also have blocking asynchronous system traps (ASTs) associated with them. Therefore, whenever a process requests a page lock, the process holding that page lock is sent a blocking AST (BLAST) by OpenVMS. The process that receives such a blocking AST queues the fact that the page lock should be released as soon as possible. However, the page lock cannot be released immediately.

Such work requests to release page locks are handled at verb commit time. An Oracle Rdb verb is an Oracle Rdb query that executes atomically, within a transaction. Therefore, verbs that require the scan of a large table, for example, can be quite long. An updating application does not release page locks until its verb has completed.

The reasons for holding on to the page locks until the end of the verb are fundamental to the database management system.

6.1.54 SQL Does Not Display Storage Map Definition After Cascading Delete of Storage Area

When you delete a storage area using the CASCADE keyword and that storage area is not the only area to which the storage map refers, the SHOW STORAGE MAP statement no longer shows the placement definition for that storage map.

The following example demonstrates this restriction:

```
SQL> SHOW STORAGE MAP DEGREES_MAP1
      DEGREES_MAP1
For Table:           DEGREES1
Compression is:     ENABLED
Partitioning is:    NOT UPDATABLE
Store clause:       STORE USING (EMPLOYEE_ID)
                   IN DEG_AREA WITH LIMIT OF ('00250')
                   OTHERWISE IN DEG_AREA2

SQL> DISCONNECT DEFAULT;
SQL> -- Drop the storage area, using the CASCADE keyword.
SQL> ALTER DATABASE FILENAME MF_PERSONNEL
cont> DROP STORAGE AREA DEG_AREA CASCADE;
SQL> --
SQL> -- Display the storage map definition.
SQL> ATTACH 'FILENAME MF_PERSONNEL';
SQL> SHOW STORAGE MAP DEGREES_MAP1
      DEGREES_MAP1
For Table:           DEGREES1
Compression is:     ENABLED
Partitioning is:    NOT UPDATABLE

SQL>
```

The other storage area, DEG_AREA2, still exists, even though the SHOW STORAGE MAP statement does not display it.

A workaround is to use the RMU Extract command with the Items=Storage_Map qualifier to see the mapping.

6.1.55 ARITH_EXCEPT or Incorrect Results Using LIKE IGNORE CASE

When you use LIKE...IGNORE CASE, programs linked under Oracle Rdb Release 4.2 and Release 5.1, but run under higher versions of Oracle Rdb, may result in incorrect results or %RDB-E-ARITH_EXCEPT exceptions.

To work around the problem, avoid using IGNORE CASE with LIKE, or recompile and relink under a higher version (Release 6.0 or higher.)

6.1.56 Different Methods of Limiting Returned Rows from Queries

Oracle® Rdb for OpenVMS

You can establish the query governor for rows returned from a query by using the SQL SET QUERY LIMIT statement, a logical name, or a configuration parameter. This note describes the differences between the mechanisms.

- ◆ If you define the RDMS\$BIND_QG_REC_LIMIT logical name or RDB_BIND_QG_REC_LIMIT configuration parameter to a small value, the query will often fail with no rows returned. The following example demonstrates setting the limit to 10 rows and the resulting failure:

```
$ DEFINE RDMS$BIND_QG_REC_LIMIT 10
$ SQL$
SQL> ATTACH 'FILENAME MF_PERSONNEL';
SQL> SELECT EMPLOYEE_ID FROM EMPLOYEES;
%RDB-F-EXQUOTA, Oracle Rdb runtime quota exceeded
-RDMS-E-MAXRECLIM, query governor maximum limit of rows has been reached
```

Interactive SQL must load its metadata cache for the table before it can process the SELECT statement. In this example, interactive SQL loads its metadata cache to allow it to check that the column EMPLOYEE_ID really exists for the table. The queries on the Oracle Rdb system tables RDB\$RELATIONS and RDB\$RELATION_FIELDS exceed the limit of rows.

Oracle Rdb does not prepare the SELECT statement, let alone execute it. Raising the limit to a number less than 100 (the cardinality of EMPLOYEES) but more than the number of columns in EMPLOYEES (that is, the number of rows to read from the RDB\$RELATION_FIELDS system table) is sufficient to read each column definition. To see an indication of the queries executed against the system tables, define the RDMS\$DEBUG_FLAGS logical name or the RDB_DEBUG_FLAGS configuration parameter as S or B.

- ◆ If you set the row limit using the SQL SET QUERY statement and run the same query, it returns the number of rows specified by the SQL SET QUERY statement before failing:

```
SQL> ATTACH 'FILENAME MF_PERSONNEL';
SQL> SET QUERY LIMIT ROWS 10;
SQL> SELECT EMPLOYEE_ID FROM EMPLOYEES;
EMPLOYEE_ID
00164
00165
.
.
.
00173
%RDB-E-EXQUOTA, Oracle Rdb runtime quota exceeded
-RDMS-E-MAXRECLIM, query governor maximum limit of rows has been reached
```

The SET QUERY LIMIT specifies that only user queries be limited to 10 rows. Therefore, the queries used to load the metadata cache are not restricted in any way.

Like the SET QUERY LIMIT statement, the SQL precompiler and module processor command line qualifiers (QUERY_MAX_ROWS and SLOPTIONS=QUERY_MAX_ROWS) only limit user queries.

Keep the differences in mind when limiting returned rows using the logical name RDMS\$BIND_QG_REC_LIMIT or the configuration parameter RDB_BIND_QG_REC_LIMIT. They may limit more queries than are obvious. This is important when using 4GL tools, the SQL precompiler, the SQL module processor, and other interfaces that read the Oracle Rdb system tables

as part of query processing.

6.1.57 Suggestions for Optimal Usage of the SHARED DATA DEFINITION Clause for Parallel Index Creation

The CREATE INDEX process involves the following steps:

1. Process the metadata.
2. Lock the index name.
Because new metadata (which includes the index name) is not written to disk until the end of the index process, Oracle Rdb must ensure index name uniqueness across the database during this time by taking a special lock on the provided index name.
3. Read the table for sorting by selected index columns and ordering.
4. Sort the key data.
5. Build the index (includes partitioning across storage areas).
6. Write new metadata to disk.

Step 6 is the point of conflict with other index definers because the system table and indexes are locked like any other updated table.

Multiple users can create indexes on the same table by using the RESERVING table_name FOR SHARED DATA DEFINITION clause of the SET TRANSACTION statement. For optimal usage of this capability, Oracle Rdb suggests the following guidelines:

- ◆ You should commit the transaction immediately after the CREATE INDEX statement so that locks on the table are released. This avoids lock conflicts with other index definers and improves overall concurrency.
- ◆ By assigning the location of the temporary sort work files SORTWORK0, SORTWORK1, ..., SORTWORK9 to different disks for each parallel process that issues the SHARED DATA DEFINITION statement, you can increase the efficiency of sort operations. This minimizes any possible disk I/O bottlenecks and allows overlap of the SORT read/write cycle.
- ◆ If possible, enable global buffers and specify a buffer number large enough to hold a sufficient amount of table data. However, do not define global buffers larger than the available system physical memory. Global buffers allow sharing of database pages and thus result in disk I/O savings. That is, pages are read from disk by one of the processes and then shared by the other index definers for the same table, reducing the I/O load on the table.
- ◆ If global buffers are not used, ensure that enough local buffers exist to keep much of the index cached (use the RDM\$BIND_BUFFERS logical name or RDB_BIND_BUFFERS configuration parameter or the NUMBER OF BUFFERS IS clause in SQL to change the number of buffers).
- ◆ To distribute the disk I/O load, place the storage areas for the indexes on separate disk drives. Note that using the same storage area for multiple indexes will result in contention during the index creation (Step 5) for SPAM pages.
- ◆ Consider placing the .ruj file for each parallel definer on its own disk or an infrequently used disk.
- ◆ Even though snapshot I/O should be minimal, consider disabling snapshots during parallel index creation.
- ◆ Refer to the *Oracle Rdb Guide to Performance and Tuning* to determine the appropriate working set values for each process to minimize excessive paging activity. In particular, avoid using working set parameters where the difference between WSQUOTA and WSEXTENT is large. The SORT utility uses the difference between these two values to

allocate scratch virtual memory. A large difference (that is, the requested virtual memory grossly exceeds the available physical memory) may lead to excessive page faulting.

- ◆ The performance benefits of using SHARED DATA DEFINITION can best be observed when creating many indexes in parallel. The benefit is in the average elapsed time, not in CPU or I/O usage. For example, when two indexes are created in parallel using the SHARED DATA DEFINITION clause, the database must be attached twice, and the two attaches each use separate system resources.
- ◆ Using the SHARED DATA DEFINITION clause on a single-file database or for indexes defined in the RDB\$SYSTEM storage area is not recommended.

The following table displays the elapsed time benefit when creating multiple indexes in parallel with the SHARED DATA DEFINITION clause. The table shows the elapsed time for 10 parallel process index creations (Index1, Index2,...Index10) and one process with 10 sequential index creations (All10). In this example, global buffers are enabled and the number of buffers is 500. The longest time for a parallel index creation is Index7 with an elapsed time of 00:02:34.64, compared to creating 10 indexes sequentially with an elapsed time of 00:03:26.66. The longest single parallel create index elapsed time is shorter than the elapsed time of creating all 10 of the indexes serially.

Index Create Job	Elapsed Time
Index1	00:02:22.50
Index2	00:01:57.94
Index3	00:02:06.27
Index4	00:01:34.53
Index5	00:01:51.96
Index6	00:01:27.57
Index7	00:02:34.64
Index8	00:01:40.56
Index9	00:01:34.43
Index10	00:01:47.44
All 10	00:03:26.66

6.1.58 Side Effect when Calling Stored Routines

When calling a stored routine, you must not use the same routine to calculate argument values by a stored function. For example, if the routine being called is also called by a stored function during the calculation of an argument value, passed arguments to the routine may be incorrect.

The following example shows a stored procedure P being called during the calculation of the arguments for another invocation of the stored procedure P:

```
SQL> CREATE MODULE M
cont>     LANGUAGE SQL
cont>
cont>     PROCEDURE P (IN :A INTEGER, IN :B INTEGER, OUT :C INTEGER);
cont>     BEGIN
cont>     SET :C = :A + :B;
cont>     END;
cont>
cont>     FUNCTION F ( ) RETURNS INTEGER
```

Oracle® Rdb for OpenVMS

```
cont> COMMENT IS 'expect F to always return 2';
cont> BEGIN
cont> DECLARE :B INTEGER;
cont> CALL P (1, 1, :B);
cont> TRACE 'RETURNING ', :B;
cont> RETURN :B;
cont> END;
cont> END MODULE;
SQL>
SQL> SET FLAGS 'TRACE';
SQL> BEGIN
cont> DECLARE :CC INTEGER;
cont> CALL P (2, F(), :CC);
cont> TRACE 'Expected 4, got ', :CC;
cont> END;
~Xt: returning 2
~Xt: Expected 4, got 3
```

The result as shown above is incorrect. The routine argument values are written to the called routine's parameter area before complex expression values are calculated. These calculations may (as in the example) overwrite previously copied data.

The workaround is to assign the argument expression (in this example calling the stored function F) to a temporary variable and pass this variable as the input for the routine. The following example shows the workaround:

```
SQL> BEGIN
cont> DECLARE :BB, :CC INTEGER;
cont> SET :BB = F();
cont> CALL P (2, :BB, :CC);
cont> TRACE 'Expected 4, got ', :CC;
cont> END;
~Xt: returning 2
~Xt: Expected 4, got 4
```

This problem will be corrected in a future version of Oracle Rdb.

6.1.59 Considerations when Using Holdable Cursors

If your applications use holdable cursors, be aware that after a COMMIT or ROLLBACK statement is executed, the result set selected by the cursor may not remain stable. That is, rows may be inserted, updated, and deleted by other users because no locks are held on the rows selected by the holdable cursor after a commit or rollback occurs. Moreover, depending on the access strategy, rows not yet fetched may change before Oracle Rdb actually fetches them.

As a result, you may see the following anomalies when using holdable cursors in a concurrent user environment:

- ◆ If the access strategy forces Oracle Rdb to take a data snapshot, the data read and cached may be inaccurate by the time the cursor fetches the data.
For example, user 1 opens a cursor and commits the transaction. User 2 deletes rows read by user 1 (this is possible because the read locks are released). It is possible for user 1 to report data now deleted and committed.
- ◆ If the access strategy uses indexes that allow duplicates, updates to the duplicates chain may

cause rows to be skipped, or even revisited.

Oracle Rdb keeps track of the dbkey in the duplicate chain pointing to the data that was fetched. However, the duplicates chain could be revised by the time Oracle Rdb returns to using it.

Holdable cursors are a very powerful feature for read-only or predominantly read-only environments. However, in concurrent update environments, the instability of the cursor may not be acceptable. The stability of holdable cursors for update environments will be addressed in future versions of Oracle Rdb.

You can define the logical name RDMS\$BIND_HOLD_CURSOR_SNAP or configuration parameter RDB_BIND_HOLD_CURSOR_SNAP to the value 1 to force all hold cursors to fetch the result set into a cached data area. (The cached data area appears as a "Temporary Relation" in the optimizer strategy displayed by the SET FLAGS STRATEGY statement or the RDMS\$DEBUG_FLAGS S flag.) This logical name or configuration parameter helps to stabilize the cursor to some degree.

6.1.60 INCLUDE SQLDA2 Statement Is Not Supported for SQL Precompiler for PL/I in Oracle Rdb Release 5.0 or Higher

The SQL statement INCLUDE SQLDA2 is not supported for use with the PL/I precompiler in Oracle Rdb Release 5.0 or higher.

There is no workaround. This problem will be fixed in a future version of Oracle Rdb.

6.1.61 SQL Pascal Precompiler Processes ARRAY OF RECORD Declarations Incorrectly

The Pascal precompiler for SQL gives an incorrect %SQL-I-UNMATEND error when it parses a declaration of an array of records. The precompiler does not associate the END statement with the record definition, and the resulting confusion in host variable scoping causes a fatal error.

A workaround for the problem is to declare the record as a type and then define your array of that type. For example:

```
main.spa:

    program main (input,output);

    type
    exec sql include 'bad_def.pin';      !gives error
    exec sql include 'good_def.pin';    !ok
    var
        a : char;

    begin
    end.
```

```
bad_def.pin
```



```

x_record = record
y  : char;
variable_a: array [1..50] of record
    a_fld1 : char;
    b_fld2 : record;
        t : record
            v : integer;
        end;
    end;
end;
end;
-----

good_def.pin

good_rec = record
    a_fld1 : char;
    b_fld2 : record
        t : record
            v: integer;
        end;
    end;
end;

x_record = record
    y  : char
    variable_a : array [1..50] of good_rec;
end;

```

6.1.62 RMU Parallel Backup Command Not Supported for Use with SLS

The RMU Parallel Backup command is not supported for use with the Storage Library System (SLS) for OpenVMS.

6.2 Oracle CDD/Repository Restrictions

This section describes known problems and restrictions in Oracle CDD/Repository Release 7.0 and earlier.

6.2.1 Oracle CDD/Repository Compatibility with Oracle Rdb Features

Some Oracle Rdb features are not fully supported by all versions of Oracle CDD/Repository. [Table 6-1](#) shows which versions of Oracle CDD/Repository support Oracle Rdb features and the extent of support.

In [Table 6-1](#), repository support for Oracle Rdb features can vary as follows:

- ◆ **Explicit support**—The repository recognizes and integrates the feature, and you can use the repository to manipulate the item.
- ◆ **Implicit support**—The repository recognizes and integrates the feature, but you cannot use any repository interface to manipulate the item.
- ◆ **Pass-through support**—The repository does not recognize or integrate the feature, but allows the Oracle Rdb operation to complete without aborting or overwriting metadata. With pass-through support, a CDD-I-MBLRSYNINFO informational message may be returned.

Table 6-1 Oracle CDD/Repository Compatibility for Oracle Rdb Features

Oracle Rdb Feature	Minimum Release of Oracle Rdb	Minimum Release of Oracle CDD/Repository	Support
CASE, NULLIF, and COALESCE expressions	6.0	6.1	Implicit
CAST function	4.1	7.0	Explicit
Character data types to support character sets	4.2	6.1	Implicit
Collating sequences	3.1	6.1	Explicit
Constraints (PRIMARY KEY, UNIQUE, NOT NULL, CHECK, FOREIGN KEY)	3.1	5.2	Explicit
CURRENT_DATE, CURRENT_TIME, and CURRENT_TIMESTAMP functions	4.1	7.0	Explicit
CURRENT_USER, SESSION_USER, SYSTEM_USER functions	6.0	7.0	Explicit
Date arithmetic	4.1	6.1	Pass-through
DATE ANSI, TIME, TIMESTAMP, and INTERVAL data types	4.1	6.1	Explicit
Delimited identifiers	4.2	6.1 ¹	Explicit
External functions	6.0	6.1	Pass-through
External procedures	7.0	6.1	Pass-through

EXTRACT, CHAR_LENGTH, and OCTET_LENGTH functions	4.1	6.1	Explicit
GRANT/REVOKE privileges	4.0	5.0 accepts but does not store information	Pass-through
Indexes	1.0	5.2	Explicit
INTEGRATE DOMAIN	6.1	6.1	Explicit
INTEGRATE TABLE	6.1	6.1	Explicit
Logical area thresholds for storage maps and indexes	4.1	5.2	Pass-through
Multinational character set	3.1	4.0	Explicit
Multiversion environment (multiple Rdb versions)	4.1	5.1	Explicit
NULL keyword	2.2	7.0	Explicit
Oracle7 compatibility functions, such as CONCAT, CONVERT, DECODE, and SYSDATE	7.0	7.0	Explicit
Outer joins, derived tables	6.0	7.0	Pass-through
Query outlines	6.0	6.1	Pass-through
Storage map definitions correctly restored	3.0	5.1	Explicit
Stored functions	7.0	6.1	Pass-through
Stored procedures	6.0	6.1	Pass-through
SUBSTRING function	4.0	7.0 supports all features 5.0 supports all but 4.2 MIA features ²	Explicit
Temporary tables	7.0	6.1	Pass-through
Triggers	3.1	5.2	Pass-through
TRUNCATE TABLE	7.0	6.1	Pass-through
TRIM and POSITION functions	6.1	7.0	Explicit
UPPER, LOWER, TRANSLATE functions	4.2	7.0	Explicit
USER function	2.2	7.0	Explicit

¹The repository does not preserve the distinction between uppercase and lowercase identifiers. If you use delimited identifiers with Oracle Rdb, the repository ensures that the record definition does not include objects with names that are duplicates except for case.

²Multivendor Integration Architecture (MIA) features include the CHAR_LENGTH clause and the TRANSLATE function.

6.2.2 Multischema Databases and CDD/Repository

You cannot use multischema databases with CDD/Repository and Oracle Rdb release 7.0 and earlier. This problem will be corrected in a future release of Oracle Rdb.

6.2.3 Interaction of Oracle CDD/Repository Release 5.1 and Oracle RMU Privileges Access Control Lists

Oracle Rdb provides special Oracle RMU privileges that use the unused portion of the OpenVMS access control list (ACL) to manage access to Oracle RMU operations.

You can use the RMU Set Privilege and RMU Show Privilege commands to set and show the Oracle RMU privileges. The DCL SHOW ACL and DIRECTORY/ACL commands also show the added access control information; however, these tools cannot translate the names defined by Oracle Rdb.

Note

The RMU Convert command propagates the database internal ACL to the root file for access control entries (ACEs) that possess the SECURITY and DBADM (ADMINISTRATOR) privileges.

Oracle CDD/Repository protects its repository (dictionary) by placing the CDD\$SYSTEM rights identifier on each file created within the anchor directory. CDD\$SYSTEM is a special, reserved rights identifier created by Oracle CDD/Repository.

When Oracle CDD/Repository executes the DEFINE REPOSITORY command, it adds (or augments) an OpenVMS default ACL to the anchor directory. Typically, this ACL allows access to the repository files for CDD\$SYSTEM and denies access to everyone else. All files created in the anchor directory inherit this default ACL, including the repository database.

Unfortunately, there is an interaction between the default ACL placed on the repository database by Oracle CDD/Repository and the Oracle RMU privileges ACL processing.

Within the ACL on the repository database, the default access control entries (ACEs) that were inherited from the anchor directory will precede the ACEs added by RMU Restore. As a result, the CDD\$SYSTEM identifier will not have any Oracle RMU privileges granted to it. Without these privileges, if the user does not have the OpenVMS SYSPRV privilege enabled, Oracle RMU operations, such as Convert and Restore, will not be allowed on the repository database.

The following problems may be observed by users who do not have the SYSPRV privilege enabled:

- ◆ While executing a CDO DEFINE REPOSITORY or DEFINE DICTIONARY command:
 - ◇ If the CDD\$TEMPLATEDB backup (.rbf) file was created by a previous version of Oracle Rdb, the automatic RMU Convert operation that will be carried out on the .rbf file will fail because SYSPRV privilege is required.
 - ◇ If the CDD\$TEMPLATEDB backup (.rbf) file was created by the current version of Oracle Rdb, the restore of the repository database will fail because the default ACEs that already existed on the repository file that was backed up will take precedence, preventing RMU\$CONVERT and RMU\$RESTORE privileges from being granted to CDD\$SYSTEM or the user.
 - ◇ If no CDD\$TEMPLATEDB is available, the repository database will be created without a template, inheriting the default ACL from the parent directory. The ACE containing all the required Oracle RMU privileges will be added to the end of the ACL; however, the preexisting default ACEs will prevent any Oracle RMU privilege

from being granted.

- ◆ You must use the RMU Convert command to upgrade the database disk format to Oracle Rdb after installing Release 7.0. This operation requires the SYSPRV privilege. During the conversion, RMU Convert adds the ACE containing the Oracle RMU privileges at the end of the ACL. Because the repository database already has the default Oracle CDD/Repository ACL associated with it, the Oracle CDD/Repository ACL will take precedence, preventing the granting of the Oracle RMU privileges.
- ◆ During a CDO MOVE REPOSITORY command, the Oracle RMU privilege checking may prevent the move, as the RMU\$COPY privilege has not been granted on the repository database.
- ◆ When you execute the CDD template builder CDD_BUILD_TEMPLATE, the step involving RMU Backup privilege has not been granted.

Oracle CDD/Repository Releases 5.2 and higher correct this problem. A version of the Oracle CDD/Repository software that corrects this problem and allows new repositories to be created using Oracle Rdb is provided on the Oracle Rdb kit for use on OpenVMS VAX systems. See [Section 6.2.3.1](#) for details.

6.2.3.1 Installing the Corrected CDDSHR Images

OpenVMS VAX Systems

Note

The following procedure must be carried out if you have installed or plan to install Oracle Rdb and have already installed CDD/Repository Release 5.1 software on your system.

Due to the enhanced security checking associated with Oracle RMU commands in Oracle Rdb on OpenVMS VAX, existing CDDSHR images for CDD/Repository Release 5.1 must be upgraded to ensure that the correct Oracle RMU privileges are applied to newly created or copied repository databases.

Included in the Oracle Rdb for OpenVMS VAX distribution kit is a CDD upgraded image kit, called CDDRDB042, that must be installed after you have installed the Oracle Rdb for OpenVMS VAX kit.

This upgrade kit should be installed by using VMSINSTAL. It automatically checks which version of CDDSHR you have installed and replaces the existing CDDSHR.EXE with the corrected image file. The existing CDDSHR.EXE will be renamed SYS\$LIBRARY:OLD_CDDSHR.EXE.

The upgrade installation will also place a new CDD_BUILD_TEMPLATE.COM procedure in SYS\$LIBRARY for use with CDD/Repository V5.1.

Note

If you upgrade your repository to CDD/Repository V5.1 after you install Oracle Rdb V7.0, you must install the corrected CDDSHR image again to ensure that the correct CDDSHR images have been made available.

The CDD/Repository upgrade kit determines which version of CDD/Repository is

installed and replaces the existing CDDSHR.EXE with the appropriate version of the corrected image.

6.2.3.2 CDD Conversion Procedure

OpenVMS VAX Systems

Oracle Rdb provides RDB\$CONVERT_CDD\$DATABASE.COM, a command procedure that both corrects the anchor directory ACL and performs the RMU Convert operation. The command procedure is located in SYS\$LIBRARY.

Note

You must have SYSPRV enabled before you execute the procedure RDB\$CONVERT_CDD\$DATABASE.COM because the procedure performs an RMU Convert operation.

Use the procedure RDB\$CONVERT_CDD\$DATABASE.COM to process the anchor directory and update the ACLs for both the directory and, if available, the repository database.

This procedure accepts one parameter: the name of the anchor directory that contains, or will contain, the repository files. For example:

```
$ @SYS$LIBRARY:DECRDB$CONVERT_CDD$DATABASE [PROJECT.CDD_REP]
```

If many repositories exist on a system, you may want to create a DCL command procedure to locate them, set the Oracle RMU privileges ACL, and convert the databases. Use DCL commands similar to the following:

```
$ LOOP:
$   REP_SPEC = F$SEARCH("[000000...]CDD$DATABASE.RDB")
$   IF REP_SPEC .NES. ""
$   THEN
$       @SYS$LIBRARY:DECRDB$CONVERT_CDD$DATABASE -
$       'F$PARSE(REP_SPEC,, "DIRECTORY")'
$       GOTO LOOP
$   ENDIF
```

[| Contents](#)