

# **Oracle® Rdb for OpenVMS**

# Table of Contents

<b><u>Oracle® Rdb for OpenVMS</u></b> .....	1
<b><u>Release Notes</u></b> .....	2
<b><u>September 2009</u></b> .....	3
<b><u>Contents</u></b> .....	4
<b><u>Preface</u></b> .....	5
<b><u>Purpose of This Manual</u></b> .....	6
<b><u>Intended Audience</u></b> .....	7
<b><u>Document Structure</u></b> .....	8
<b><u>Chapter 1 Installing Oracle Rdb Release 7.2.4.0</u></b> .....	9
<b><u>1.1 Oracle Rdb on HP OpenVMS Industry Standard 64</u></b> .....	10
<b><u>1.2 Requirements</u></b> .....	11
<u>1.2.1 Ensure No Processes Have RDMSHRP Image Activated</u> .....	11
<b><u>1.3 Intel Itanium Processor 9100 "Montvale" Support</u></b> .....	13
<b><u>1.4 Maximum OpenVMS Version Check</u></b> .....	14
<b><u>1.5 Database Format Changed</u></b> .....	15
<b><u>1.6 Using Databases from Releases Earlier than V7.0</u></b> .....	16
<b><u>1.7 Invoking the VMSINSTAL Procedure</u></b> .....	17
<b><u>1.8 Stopping the Installation</u></b> .....	18
<b><u>1.9 After Installing Oracle Rdb</u></b> .....	19
<b><u>1.10 VMS\$MEM RESIDENT USER Rights Identifier Required</u></b> .....	20
<b><u>1.11 Installation, Configuration, Migration, Upgrade Suggestions</u></b> .....	21
<b><u>Chapter 2 Software Errors Fixed in Oracle Rdb Release 7.2.4.0</u></b> .....	24
<b><u>2.1 Software Errors Fixed That Apply to All Interfaces</u></b> .....	25
<u>2.1.1 Bugcheck With SYSTEM-F-ROPRAND in KODTXN\$POST TSNBLK UPDATE</u> .....	25
<u>2.1.2 Memory Leak On Systems With RAD Support Enabled</u> .....	26
<u>2.1.3 Incorrect Messages From RMU /MONITOR START When RDM\$MON USERNAME</u> <u>Specifies Non-existent Account</u> .....	26

# Table of Contents

<b>2.1 Software Errors Fixed That Apply to All Interfaces</b>	
2.1.4 Potential INVEXCEPTN System Crash Using LARGE MEMORY IS ENABLED or SHARED MEMORY IS SYSTEM on Itanium	26
2.1.5 Wrong Result From Query With Derived Table	27
2.1.6 Application Loses Virtual Memory (Memory Leak)	33
2.1.7 VMSS\$BUFFER OBJECT USER Not Always Checked	33
2.1.8 Deleted Space in Uniform Areas Not Reclaimed by Other Users When Erased Rows Moved From Row Cache to Disk	33
2.1.9 LOWER Function Problem When Using ISOLATINCYRILLIC Character Set	34
2.1.10 Query Bugchecks with SYSTEM-F-FLTINV	34
2.1.11 Bugcheck Altering Storage Map	35
2.1.12 Wrong Results when Expression Shared Between Aggregate Filters	36
2.1.13 Optimize For Sequential Access Using Index	36
2.1.14 Query With Match Strategy for Left Outer Join Bugchecks	37
2.1.15 Bugchecks at DIOFETCH\$FETCH SNAP_SEG or DIOCCH\$FETCH SNAP_SEG	37
2.1.16 Bugcheck in PSII2INSERTDUPBBC	38
<b>2.2 SQL Errors Fixed</b>	<b>39</b>
2.2.1 Unexpected Failures of TRUNC and ROUND Functions	39
2.2.2 Unexpected Restrictions in Date/Time Subtraction	39
2.2.3 Unexpected DEFVALINC Error When Altering a Column's Domain	40
2.2.4 Unexpected FLDNOTCRS Error When Referencing DBKEY or ROWID in Quotes	40
2.2.5 Unexpected Bugcheck When Assigning Value to FOR Cursor Variables	41
2.2.6 Unexpected Bugcheck When Executing External Routine	42
2.2.7 Unexpected Error When Importing a Database With Profile Definitions	42
2.2.8 Unexpected Bugcheck When Table Partitions Used in RESERVING Clause	43
2.2.9 Unexpected Failure of GRANT and REVOKE When Using Synonyms	44
2.2.10 Unexpected Correlation Names Appearing as Column Headers	44
2.2.11 Wrong Results When CONCAT Used in an Aggregate Expression with DISTINCT Clause	45
2.2.12 UPDATE ONLY CURSOR Clause Quietly Ignored by SQL Module Language Processor and SQL Precompiler	45
<b>2.3 RMU Errors Fixed</b>	<b>46</b>
2.3.1 Application Hangs Using Automatic AIJ Backups	46
2.3.2 Failed RMU MOVE or COPY Deletes Wrong Files	46
2.3.3 RMU Extract Now Extracts Comments for Constraints	46
2.3.4 Bugcheck in RMU /COLLECT OPTIMIZER STATISTICS /STATISTICS=WORKLOAD	47
2.3.5 Increased Cardinality Limit in RMU /COLLECT OPTIMIZER STATISTICS /STATISTICS=WORKLOAD	47
2.3.6 Incorrect LOCK TIMEOUT With RMU /BACKUP /ONLINE /LOCK TIMEOUT /LOG	48
2.3.7 RMU /RESTORE /ONLINE /JUST CORRUPT Bugcheck	49
2.3.8 RMU /LOAD /PARALLEL COSI-F-READERR Without VMS BYPASS Privilege	50
<b>2.4 LogMiner Errors Fixed</b>	<b>51</b>
2.4.1 Continuous LogMiner Startup Serialized With AIJ Backups	51
2.4.2 RMU /UNLOAD /AFTER JOURNAL Allows /RESTART Without /CONTINUOUS	51

# Table of Contents

<b><u>2.5 Row Cache Errors Fixed</u></b> .....	<b>52</b>
<u>2.5.1 Incremental Backup With Row Cache</u> .....	52
<b><u>2.6 Hot Standby Errors Fixed</u></b> .....	<b>53</b>
<u>2.6.1 Hot Standby LRS Database Prefetch Count Limit</u> .....	53
<b><u>Chapter 3 Software Errors Fixed in Oracle Rdb Release 7.2.3.5</u></b> .....	<b>54</b>
<b><u>3.1 Software Errors Fixed That Apply to All Interfaces</u></b> .....	<b>55</b>
<u>3.1.1 Patch Required When Using VMS V8.3 or V8.2-1 on Itanium</u> .....	55
<u>3.1.2 Wrong Result From Outer Zigzag Match</u> .....	55
<u>3.1.3 Cannot Disable After-Image Journaling when Fast-Commit is Enabled</u> .....	57
<u>3.1.4 Alignment Faults Related to Remote Database Access</u> .....	57
<u>3.1.5 RDMS-E-NOSOL FOUND Error From Full OJ Query</u> .....	57
<b><u>3.2 SQL Errors Fixed</u></b> .....	<b>60</b>
<u>3.2.1 Wrong Result from CONCAT Query Under ORACLE Dialects</u> .....	60
<u>3.2.2 Unexpected Bugcheck During RENAME Operation</u> .....	60
<u>3.2.3 Unexpected DDL CONFLICT Error Generated by ALTER TABLE ... ENABLE     CONSTRAINT</u> .....	60
<b><u>3.3 RMU Errors Fixed</u></b> .....	<b>62</b>
<u>3.3.1 Unexpected RMU/REPAIR Error if Index RDB\$REL REL ID NDX is Sorted Ranked</u> .....	62
<b><u>3.4 LogMiner Errors Fixed</u></b> .....	<b>63</b>
<u>3.4.1 Continuous LogMiner Shutdown Via SYS\$FORCEX Causes DBR</u> .....	63
<b><u>3.5 RMU Show Statistics Errors Fixed</u></b> .....	<b>64</b>
<u>3.5.1 RMU/SHOW STATISTICS Row Cache Overview Invalid Percentage</u> .....	64
<b><u>Chapter 4 Software Errors Fixed in Oracle Rdb Release 7.2.3.4</u></b> .....	<b>65</b>
<b><u>4.1 Software Errors Fixed That Apply to All Interfaces</u></b> .....	<b>66</b>
<u>4.1.1 Single Byte String Comparison With Some Relational Operators Returns Different Result     on Itanium</u> .....	66
<u>4.1.2 Wrong Results Queries When Using DISTINCT</u> .....	66
<u>4.1.3 Potential System Crash Using LARGE MEMORY IS ENABLED or SHARED MEMORY     IS SYSTEM on Itanium</u> .....	67
<u>4.1.4 Bugcheck from Nested GROUP BY Query with OR Predicate</u> .....	67
<u>4.1.5 Sort Query Returns Incorrect Data</u> .....	68
<u>4.1.6 Wrong Result from Query with UNION and OR Predicates</u> .....	69
<b><u>4.2 SQL Errors Fixed</u></b> .....	<b>72</b>
<u>4.2.1 Unexpected Data Type Restrictions for Some SQL Operations</u> .....	72
<u>4.2.2 Unexpected Bugcheck from DECLARE Cursor Through Dynamic SQL</u> .....	72
<u>4.2.3 Unexpected INVQUAL Error When Using DCL SET PROCESS/PARSE STYLE Qualifier</u> .....	73
<u>4.2.4 Improved Diagnostics for INVVALIS Error</u> .....	73

# Table of Contents

<b><u>Chapter 5 Software Errors Fixed in Oracle Rdb Release 7.2.3.3</u></b> .....	<b>74</b>
<b><u>Chapter 6 Software Errors Fixed in Oracle Rdb Release 7.2.3.2</u></b> .....	<b>75</b>
<b><u>6.1 Software Errors Fixed That Apply to All Interfaces</u></b> .....	<b>76</b>
<u>6.1.1 Result Data Type From NEGATE was Incorrect</u> .....	76
<u>6.1.2 Floating Multiply May Produce Lower Precision Result</u> .....	76
<u>6.1.3 Zero–Length String Compare on Itanium Using Run–Time Native Compiler Could Cause Bugcheck</u> .....	76
<u>6.1.4 Query Slows Down Using Seq Scan After Reaching ThreLim</u> .....	77
<u>6.1.5 String Compares of More Than 255 Bytes Could Return False Matches</u> .....	78
<b><u>6.2 SQL Errors Fixed</u></b> .....	<b>80</b>
<u>6.2.1 Incorrect COMPRESSION Attributes Applied to Vertical Row Partitioned Table</u> .....	80
<u>6.2.2 Unexpected Bugcheck During ALTER INDEX ... BUILD PARTITION on Sorted Ranked Index</u> .....	81
<u>6.2.3 Change in Handling of Dynamic Cursor Names</u> .....	81
<u>6.2.4 Unexpected DISABLED Indices After Database Created by IMPORT DATABASE Statement</u> .....	81
<b><u>6.3 RMU Errors Fixed</u></b> .....	<b>84</b>
<u>6.3.1 Use of External Media Manager or Tape Librarian Causes a Hang or an ACCVIO</u> .....	84
<u>6.3.2 RMU/BACKUP Sometimes Dismounts Last Volume in a Multi–tape Drive Case</u> .....	84
<u>6.3.3 Parallel Backup Hangs</u> .....	85
<u>6.3.4 Users Were Audited Who Held Identifiers Not Enabled For Auditing</u> .....	85
<u>6.3.5 RMU–W–AIJBCKFAIL – RMU/RECOVER Reports Incomplete AIJ Backup</u> .....	86
<b><u>6.4 LogMiner Errors Fixed</u></b> .....	<b>88</b>
<u>6.4.1 LogMiner Incorrectly Returns DBK Multiple Times Within Transaction</u> .....	88
<b><u>Chapter 7 Software Errors Fixed in Oracle Rdb Release 7.2.3.1</u></b> .....	<b>89</b>
<b><u>7.1 Software Errors Fixed That Apply to All Interfaces</u></b> .....	<b>90</b>
<u>7.1.1 Page Transfer Via Memory Lost Database Update</u> .....	90
<u>7.1.2 Rollback Performance Improvement with Row Cache</u> .....	90
<u>7.1.3 Wrong Result from Query with Shared NULL Conjoint in OR Predicate</u> .....	90
<b><u>7.2 SQL Errors Fixed</u></b> .....	<b>92</b>
<u>7.2.1 Unexpected SQL–F–DATTYPUNK Error from Dynamic SQL Applications</u> .....	92
<u>7.2.2 Updated Support for SQL Pseudo Types for C Language</u> .....	92
<u>7.2.3 Unexpected SQL–F–NUMXPRES Error While Processing ABS Function</u> .....	93
<u>7.2.4 Unexpected Truncation of Application Parameter Data in Dynamic SQL</u> .....	93
<u>7.2.5 Unexpected Prompt When Adding a Comment to an Index</u> .....	94
<b><u>7.3 RMU Errors Fixed</u></b> .....	<b>95</b>
<u>7.3.1 RMU/BACKUP Exits with ACCVIO in ENCRYPTSHR</u> .....	95

# Table of Contents

<b><u>7.4 RMU Show Statistics Errors Fixed</u></b> .....	<b>96</b>
<u>7.4.1 RMU/SHOW STATISTICS Does Not Show BLOB Areas On Logical Area Information Screen</u> .....	96
<b><u>Chapter 8 Software Errors Fixed in Oracle Rdb Release 7.2.3.0</u></b> .....	<b>97</b>
<b><u>8.1 Software Errors Fixed That Apply to All Interfaces</u></b> .....	<b>98</b>
<u>8.1.1 Inability to Cause Free Space Scan for First Insert After Attach</u> .....	98
<u>8.1.2 Insert ACCVIO Bugcheck on IA64</u> .....	98
<u>8.1.3 Processes With Multiple Database Attaches Stall "Waiting for RTUPB list" on IA64</u> .....	99
<u>8.1.4 Multi-Threaded Processes Hang in HIB Status</u> .....	99
<u>8.1.5 Tape Savesets With More Than 999,999 Blocks Report Wrong Block Number</u> .....	99
<u>8.1.6 Possible DBR Failure Due to Address Space Being Exhausted</u> .....	100
<u>8.1.7 Control Number and Size of AIJ Initialization I/O Buffers</u> .....	100
<u>8.1.8 Some Sort Operations May Bugcheck With More Than 134 Million Rows</u> .....	101
<u>8.1.9 DBR Hangs on Page Lock After Node Crash</u> .....	101
<u>8.1.10 Alter Storage Map ... Reorganize Pages Corrupts SORTED Indices</u> .....	101
<u>8.1.11 RDMALS Bugcheck at AIJUTL\$FORMAT ARBS + 00000DA4</u> .....	102
<u>8.1.12 Bugcheck Due to Command Line Too Long</u> .....	102
<b><u>8.2 SQL Errors Fixed</u></b> .....	<b>104</b>
<u>8.2.1 ALTER PROFILE Generates a SQL Bugcheck Dump</u> .....	104
<u>8.2.2 Incorrect File Specification Reported for RDB\$STORAGE AREAS Information Table</u> .....	104
<u>8.2.3 Unexpected Bugcheck When Processing Zero Length Strings</u> .....	105
<u>8.2.4 Unexpected Bugcheck from RENAME TABLE Statement</u> .....	106
<u>8.2.5 Unexpected RDMS-F-NODBK Reported by DROP USER Statement</u> .....	106
<u>8.2.6 Unexpected COSI-F-VASFULL Reported by Long Running SQL Application</u> .....	107
<u>8.2.7 Diagnostics Lost Following Procedure Call in Dynamic SQL</u> .....	107
<u>8.2.8 Unexpected Bugcheck from SET TRANSACTION Statement</u> .....	108
<b><u>8.3 RMU Errors Fixed</u></b> .....	<b>110</b>
<u>8.3.1 RMU Extract Reports BAD_CODE Error for TRANSLATE Function</u> .....	110
<u>8.3.2 Restore of Selected Area Fails With ACCVIO</u> .....	110
<u>8.3.3 RMU Load Quietly Truncated String Data During Insert</u> .....	110
<u>8.3.4 RMU /UNLOAD /AFTER JOURNAL Sort Performance</u> .....	111
<u>8.3.5 Possible RMU/MOVE AREA/ONLINE Database Corruption</u> .....	111
<u>8.3.6 RMU /BACKUP Performance Enhanced On Some Types Of Tape Drives</u> .....	113
<u>8.3.7 RMU Analyze Reports Incorrect Compression Ratio for Some Tables</u> .....	113
<u>8.3.8 Restore May Miss Storage Area Without Warning</u> .....	113
<u>8.3.9 Incorrect View Syntax Generated by RMU Extract Item=VIEW</u> .....	114
<u>8.3.10 Revised Format for the RMU/UNLOAD/AFTER IMAGE CONTROL File</u> .....	115
<u>8.3.11 RMU Created an Incorrect Restore Options File for Single File Databases</u> .....	115
<b><u>8.4 RMU Show Statistics Errors Fixed</u></b> .....	<b>117</b>
<u>8.4.1 RMU/SHOW STATISTICS Checkpoint Information Screen Sort By Oldest Transaction Did Not Work Correctly</u> .....	117
<u>8.4.2 Date and Time Included In Lock Deadlock Log File</u> .....	117
<u>8.4.3 Incorrect RMU/SHOW STATISTICS Layout File RMU\$SHOW_STATISTICS72.CDQ</u> .....	117

# Table of Contents

<b><u>Chapter 9 Enhancements And Changes Provided in Oracle Rdb Release 7.2.4</u></b> .....	<b>118</b>
<b><u>9.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.4</u></b> .....	<b>119</b>
<u>9.1.1 Date/Time Arithmetic Enhancements</u> .....	119
<u>9.1.2 New DEFAULT PROFILE Feature</u> .....	119
<u>9.1.3 RMU /DUMP/BACKUP/OPTIONS=ROOT /HEADER ONLY Displays the Header Information Only</u> .....	122
<u>9.1.4 GET ENVIRONMENT Now Supports SQLCODE and SQLSTATE Capture</u> .....	124
<u>9.1.5 Timestamp Added to Messages For RMU LOAD and UNLOAD</u> .....	125
<u>9.1.6 New SET SOLDA Statement</u> .....	125
<u>9.1.7 RMU /SHOW VERSION Displays System Architecture and Version</u> .....	129
<u>9.1.8 New IDENT Option for SQL Module Language PRAGMA Clause</u> .....	129
<u>9.1.9 New Keyword for SQL Module Language /PRAGMA Qualifier</u> .....	132
<u>9.1.10 New IDENT Option for SQL Precompiler DECLARE MODULE Statement</u> .....	132
<u>9.1.11 New Keyword for SQL Precompiler PRAGMA Option</u> .....	134
<u>9.1.12 RDB_STATS_DATABASE Example Program</u> .....	134
<u>9.1.13 RCS Time-Based Cache Sweeping</u> .....	136
<u>9.1.14 RMU Command TSN Keyword and Qualifier Value</u> .....	137
<u>9.1.15 New Support for RENAME and CREATE SYNONYM Commands</u> .....	137
<b><u>Chapter 10 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.5</u></b> .....	<b>140</b>
<b><u>10.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.5</u></b> .....	<b>141</b>
<u>10.1.1 AIJ Extend Additional Information In Operator Notification New Feature</u> .....	141
<u>10.1.2 Default Behavior Change, New Syntax for RMU/RECOVER/CONFIRM</u> .....	141
<b><u>Chapter 11 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.2</u></b> .....	<b>144</b>
<b><u>11.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.2</u></b> .....	<b>145</b>
<u>11.1.1 New SQL Functions Added</u> .....	145
<u>11.1.2 RMU /SHOW STATISTICS Enhanced LogMiner Information Display</u> .....	147
<u>11.1.3 RMU /SHOW STATISTICS "Checkpoint Statistics" New Counters</u> .....	147
<u>11.1.4 SQL Enhancements: Allowing Optional Correlation Names</u> .....	148
<u>11.1.5 Performance Enhancements With Internal Lock Data Structures</u> .....	149
<u>11.1.6 Change in Frequency of Cardinality Updates Might be Observed</u> .....	149
<u>11.1.7 New Interactive SQL Statements</u> .....	149
<b><u>Chapter 12 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.0</u></b> .....	<b>151</b>
<b><u>12.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.0</u></b> .....	<b>152</b>
<u>12.1.1 Optional Run-Time Routine Native Compiler on I64 Enabled By Default</u> .....	152
<u>12.1.2 Temporary Table Improvements</u> .....	153
<u>12.1.3 New SIGN Builtin Function</u> .....	153
<u>12.1.4 Enhanced Simple CASE Expression</u> .....	154
<u>12.1.5 Changes in Generated Query Outline ID</u> .....	158
<u>12.1.6 ALTER INDEX ... MAINTENANCE IS ENABLED DEFERRED Syntax is Now Active</u> ... 159	
<u>12.1.7 SQL Precompiler and Module Language Compiler /ARCHITECTURE Command Line Qualifier</u> .....	160

# Table of Contents

<b><u>12.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.0</u></b>	
<u>12.1.8 PERFT4 RDB Example Program</u> .....	162
<u>12.1.9 RMU Load Quietly Truncated String Data During Insert</u> .....	162
<u>12.1.10 New FETCH FIRST and OFFSET Clauses for Select Expression</u> .....	163
<u>12.1.11 RMU Unload Now Creates SQL*Loader Control Files</u> .....	168
<b><u>Chapter 13 Documentation Corrections, Additions and Changes</u></b> .....	<b>170</b>
<b><u>13.1 Documentation Corrections</u></b> .....	<b>171</b>
<u>13.1.1 RMU /SET ROW CACHE Command Updates</u> .....	171
<u>13.1.2 Documentation for the DEBUG OPTIONS Qualifier of RMU Unload</u> .....	172
<u>13.1.3 SQL\$MSGxx.DOC Is Not Alphabetical</u> .....	173
<u>13.1.4 LOCK TIMEOUT Documentation Error in RMU Reference Manual Release 7.2</u> .....	174
<u>13.1.5 Back Out SQL PROTOCOL RETRY Configuration Parameter</u> .....	174
<u>13.1.6 Revised Example for SET OPTIMIZATION LEVEL Statement</u> .....	174
<u>13.1.7 RMU /VERIFY Process Quotas and Limits Clarification</u> .....	176
<u>13.1.8 Online Backup Can Be Performed With Transfer Via Memory</u> .....	176
<u>13.1.9 Missing Example for CREATE STORAGE MAP</u> .....	176
<u>13.1.10 RDM\$BIND MAX DBR COUNT Documentation Clarification</u> .....	178
<u>13.1.11 Database Server Process Priority Clarification</u> .....	179
<u>13.1.12 Explanation of SQL\$INT in a SQL Multiversion Environment and How to Redefine SQL\$INT</u> .....	180
<u>13.1.13 Clarification of PREPARE Statement Behavior</u> .....	181
<u>13.1.14 RDM\$BIND LOCK TIMEOUT INTERVAL Overrides the Database Parameter</u> .....	182
<b><u>13.2 Address and Phone Number Correction for Documentation</u></b> .....	<b>183</b>
<b><u>13.3 Online Document Format and Ordering Information</u></b> .....	<b>184</b>
<b><u>Chapter 14 Known Problems and Restrictions</u></b> .....	<b>185</b>
<b><u>14.1 Known Problems and Restrictions in All Interfaces</u></b> .....	<b>186</b>
<u>14.1.1 Possible Incorrect Results When Using Partitioned Descending Indexes</u> .....	186
<u>14.1.2 Remote Attach Stalls Before Detecting a Node is Unreachable</u> .....	187
<u>14.1.3 Case Sensitive Values in RDB\$CLIENT DEFAULTS.DAT</u> .....	188
<u>14.1.4 Standalone WITH Clause in Compound Statements Now Deprecated</u> .....	188
<u>14.1.5 Calling DECC\$CTRL INIT</u> .....	189
<u>14.1.6 Application and Oracle Rdb Both Using SYSS\$HIBER</u> .....	189
<u>14.1.7 Unexpected RCS Termination</u> .....	191
<u>14.1.8 Possible Incorrect Results When Using Partitioned Descending Indexes on I64</u> .....	191
<u>14.1.9 Response 'QUIT' to RMU Restart Prompt Loops</u> .....	192
<u>14.1.10 Changes for Processing Existence Logical Names</u> .....	192
<u>14.1.11 Patch Required When Using VMS V8.3 and Dedicated CPU Lock Manager</u> .....	193
<u>14.1.12 SQL Module or Program Fails with %SQL-F-IGNCASE BAD</u> .....	193
<u>14.1.13 External Routine Images Linked with PTHREAD\$RTL</u> .....	194
<u>14.1.14 SQL Procedure External Location Should Be Upper Case</u> .....	194
<u>14.1.15 Using Databases from Releases Earlier than V7.0</u> .....	195
<u>14.1.16 Partitioned Index with Descending Column and Collating Sequence</u> .....	195



# Table of Contents

<b>14.1 Known Problems and Restrictions in All Interfaces</b>	
14.1.17 Domain–Qualified TCP/IP Node Names in Distributed Transactions	196
14.1.18 ILINK–E–INVOVRINI Error on I64	197
14.1.19 New Attributes Saved by RMU/LOAD Incompatible With Prior Versions	198
14.1.20 SYSTEM–F–INSMEM Fatal Error With SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED in Galaxy Environment	198
14.1.21 Oracle Rdb and OpenVMS ODS–5 Volumes	199
14.1.22 Optimization of Check Constraints	199
14.1.23 Carryover Locks and NOWAIT Transaction Clarification	201
14.1.24 Unexpected Results Occur During Read–Only Transactions on a Hot Standby Database	202
14.1.25 Row Cache Not Allowed While Hot Standby Replication is Active	202
14.1.26 Excessive Process Page Faults and Other Performance Considerations During Oracle Rdb Sorts	203
14.1.27 Control of Sort Work Memory Allocation	204
14.1.28 The Halloween Problem	205
<b>14.2 SQL Known Problems and Restrictions</b>	<b>207</b>
14.2.1 SET FLAGS CRONO FLAG Removed	207
14.2.2 Interchange File (RBR) Created by Oracle Rdb Release 7.2 Not Compatible With Previous Releases	207
14.2.3 Single Statement LOCK TABLE is Not Supported for SQL Module Language and SQL Precompiler	207
14.2.4 Multistatement or Stored Procedures May Cause Hangs	208
14.2.5 Use of Oracle Rdb from Shareable Images	209
<b>14.3 Oracle RMU Known Problems and Restrictions</b>	<b>210</b>
14.3.1 RMU Convert Fails When Maximum Relation ID is Exceeded	210
14.3.2 RMU Unload /After Journal Requires Accurate AIP Logical Area Information	210
14.3.3 Do Not Use HYPERSORT with RMU Optimize After Journal Command	211
14.3.4 Changes in EXCLUDE and INCLUDE Qualifiers for RMU Backup	212
14.3.5 RMU Backup Operations Should Use Only One Type of Tape Drive	212
14.3.6 RMU/VERIFY Reports PGSPAMENT or PGSPMCLST Errors	213
<b>14.4 Known Problems and Restrictions in All Interfaces for Release 7.0 and Earlier</b>	<b>215</b>
14.4.1 Converting Single–File Databases	215
14.4.2 Row Caches and Exclusive Access	215
14.4.3 Exclusive Access Transactions May Deadlock with RCS Process	215
14.4.4 Strict Partitioning May Scan Extra Partitions	215
14.4.5 Restriction When Adding Storage Areas with Users Attached to Database	216
14.4.6 Multiblock Page Writes May Require Restore Operation	216
14.4.7 Replication Option Copy Processes Do Not Process Database Pages Ahead of an Application	217
<b>14.5 SQL Known Problems and Restrictions for Oracle Rdb Release 7.0 and Earlier</b>	<b>218</b>
14.5.1 ARITH EXCEPT or Incorrect Results Using LIKE IGNORE CASE	218
14.5.2 Different Methods of Limiting Returned Rows from Queries	218
14.5.3 Suggestions for Optimal Use of SHARED DATA DEFINITION Clause for Parallel Index Creation	219

# Table of Contents

<b><u>14.5 SQL Known Problems and Restrictions for Oracle Rdb Release 7.0 and Earlier</u></b>	
<u>14.5.4 Side Effect When Calling Stored Routines</u> .....	221
<u>14.5.5 Considerations When Using Holdable Cursors</u> .....	222
<u>14.5.6 ALLSERVER Privileges</u> .....	222

# Oracle® Rdb for OpenVMS

# Release Notes

Release 7.2.4.0

---

# September 2009

Oracle Rdb Release Notes, Release 7.2.4.0 for OpenVMS

Copyright © 1984, 2009 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.

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:

**US GOVERNMENT RIGHTS** Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software – Restricted Rights (June, 1987). Oracle Corporation, Inc., 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 7.2.4.0. The notes describe changed and enhanced features; upgrade and compatibility information; new and existing software problems and restrictions; and software and documentation corrections.



# 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 7.2.4.0.

# Document Structure

This manual consists of the following chapters:

<a href="#">Chapter 1</a>	Describes how to install Oracle Rdb Release 7.2.4.0.
<a href="#">Chapter 2</a>	Describes problems corrected in Oracle Rdb Release 7.2.4.0.
<a href="#">Chapter 3</a>	Describes problems corrected in Oracle Rdb Release 7.2.3.5.
<a href="#">Chapter 4</a>	Describes problems corrected in Oracle Rdb Release 7.2.3.4.
<a href="#">Chapter 5</a>	Describes problems corrected in Oracle Rdb Release 7.2.3.3.
<a href="#">Chapter 6</a>	Describes problems corrected in Oracle Rdb Release 7.2.3.2.
<a href="#">Chapter 7</a>	Describes problems corrected in Oracle Rdb Release 7.2.3.1.
<a href="#">Chapter 8</a>	Describes problems corrected in Oracle Rdb Release 7.2.3.0.
<a href="#">Chapter 9</a>	Describes enhancements introduced in Oracle Rdb Release 7.2.4.0.
<a href="#">Chapter 10</a>	Describes enhancements introduced in Oracle Rdb Release 7.2.3.5.
<a href="#">Chapter 11</a>	Describes enhancements introduced in Oracle Rdb Release 7.2.3.2.
<a href="#">Chapter 12</a>	Describes enhancements introduced in Oracle Rdb Release 7.2.3.0.
<a href="#">Chapter 13</a>	Provides information not currently available in the Oracle Rdb documentation set.
<a href="#">Chapter 14</a>	Describes problems, restrictions, and workarounds known to exist in Oracle Rdb Release 7.2.4.0.

---

# Chapter 1

## Installing Oracle Rdb Release 7.2.4.0

This software update is installed using the OpenVMS VMSINSTAL utility.

---

### NOTE

*Oracle Rdb Release 7.2 kits are full kits. There is no requirement to install any prior release of Oracle Rdb when installing new Rdb Release 7.2 kits.*

---

# 1.1 Oracle Rdb on HP OpenVMS Industry Standard 64

The Oracle Rdb product family is available on the HP OpenVMS Industry Standard 64 platform and the OpenVMS AlphaServer platform. In general, the functionality present in Oracle Rdb on OpenVMS Alpha will be available in Oracle Rdb on OpenVMS Industry Standard 64. However, certain differences between the platforms may result in minor capability and functionality differences.

The database format for Oracle Rdb Release 7.2 is the same on both I64 and Alpha platforms and databases may be accessed simultaneously from both architectures in a cluster environment. Access to an Oracle Rdb Release 7.2 database from prior Rdb versions (on Alpha or VAX platforms) or from other systems on the network is available via the Oracle Rdb remote database server.

## 1.2 Requirements

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

- This Oracle Rdb release requires the following OpenVMS environments:
  - ◆ OpenVMS Alpha V8.2 to V8.3–x.
  - ◆ OpenVMS Industry Standard 64 V8.2–1 to V8.3–x.
- Oracle Rdb must be shutdown before you install this update kit. That is, the command file SYS\$STARTUP:RMONSTOP72.COM should be executed before proceeding with this installation. If you have an OpenVMS cluster, you must shutdown the Rdb Release 7.2 monitor on all nodes in the cluster before proceeding.
- After executing RMONSTOP72.COM, no process on any system in the cluster should have any existing RDMSHRP72.EXE image activated. See [Section 1.2.1](#) for additional information.
- The installation requires approximately 280,000 blocks for OpenVMS Alpha systems.
- The installation requires approximately 500,000 blocks for OpenVMS I64 systems.
- Oracle strongly recommends that all available OpenVMS patches are installed on all systems prior to installing Oracle Rdb. Contact your HP support representative for more information and assistance.

### 1.2.1 Ensure No Processes Have RDMSHRP Image Activated

The Oracle Rdb installation procedure checks to make sure that the Oracle Rdb Monitor (RDMMON) process is not running. However, it is also important to make sure that there are no processes on the cluster that share the system disk that have image activated a prior version RDMSHRP image. Such processes may not be currently attached to a database but may do so in the future and could cause problems by using an older RDMSHRP image with a later Rdb installation.

The following command procedure can be used on each cluster node that shares the system disk to determine if there are any processes that have activated the RDMSHRP72.EXE image. This procedure should be executed by a privileged account after RMONSTOP72 has been run. Any processes that have RDMSHRP72.EXE activated at this point should be terminated prior to starting the Rdb installation procedure.

```
$ RDB$TMP = "SYS$SCRATCH:'F$UNIQUE()' .TMP
$ DEFINE /NOLOG /USER RDB$TMP 'RDB$TMP
$ ANALYZE /SYSTEM
    SET OUTPUT RDB$TMP
    SHOW PROCESS /CHANNELS ALL
    EXIT
$ SEARCH /OUTPUT='RDB$TMP' 'RDB$TMP';-1 RDMSHRP72.EXE,"PID:"
$ SEARCH 'RDB$TMP' RDMSHRP72.EXE /WINDOW=(1,0)
$ DELETE /NOLOG 'RDB$TMP';*{text}
```

In the following example, the process 2729F16D named "FOO\$SERVER" has the image RDMSHRP72.EXE activated even after RMONSTOP72.COM has been executed and this process should be terminated prior to starting the Rdb installation procedure:

```
$ @SYS$STARTUP:RMONSTOP72.COM
.
.
```

## Oracle® Rdb for OpenVMS

```
$ @FIND_RDMSHRP72_PROC.COM
```

```
OpenVMS system analyzer
```

```
Process index: 016D   Name: FOO$SERVER   Extended PID: 2729F16D  
0240 7FEF4460 8384F300 $1$DGA2:[VMS$COMMON.SYSLIB]RDMSHRP72.EXE;722  
$
```

## **1.3 Intel Itanium Processor 9100 "Montvale" Support**

For this release of Oracle Rdb on HP Integrity servers, the Intel Dual-Core Itanium 2 Processor 9100 series, code named "Montvale", is the newest processor supported.

## 1.4 Maximum OpenVMS Version Check

OpenVMS Version 8.3–x is the maximum supported version of OpenVMS for this release of Oracle Rdb.

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.



## 1.5 Database Format Changed

The Oracle Rdb on-disk database format has been incremented to 721. An RMU /CONVERT operation is required for databases created by or accessed by Oracle Rdb V7.0 or V7.1 to be accessed with Rdb Release 7.2.

Prior to upgrading to Oracle Rdb Release 7.2 and prior to converting an existing database to Oracle Rdb Release 7.2 format, Oracle strongly recommends that you perform a full database verification (with the "RMU /VERIFY /ALL" command) along with a full database backup (with the "RMU /BACKUP" command) to ensure a valid and protected database copy.

## 1.6 Using Databases from Releases Earlier than V7.0

You cannot convert or restore databases earlier than the Oracle Rdb V7.0 format directly to Oracle Rdb V7.2 format. The RMU Convert command for Oracle Rdb V7.2 supports conversions from Oracle Rdb V7.0 and V7.1 format databases only. If you have an Oracle Rdb V3.0 through V6.1 format database or database backup, you must convert it to at least Oracle Rdb V7.0 format and then convert it to Oracle Rdb V7.2 format. For example, if you have a V4.2 format database, you must convert it first to at least Oracle Rdb V7.0 format, then convert it to Oracle Rdb V7.2 format.

If you attempt to convert or restore a database that is prior to Oracle Rdb V7.0 format directly to Oracle Rdb V7.2 format, Oracle RMU generates an error.

## 1.7 Invoking the VMSINSTAL Procedure

The installation procedure for Oracle Rdb has been simplified as compared with prior Oracle Rdb major releases. All Oracle Rdb components are always installed and the number of prompts during the installation has been reduced. The installation procedure is the same for Oracle Rdb for OpenVMS Alpha and Oracle Rdb for OpenVMS I64.

To start the installation procedure, invoke the VMSINSTAL command procedure as in the following examples.

- To install the Oracle Rdb for OpenVMS I64 kit that is performance targeted for I64 platforms:

```
@SYS$UPDATE:VMSINSTAL RDBV72400IM device-name
```

- To install the Oracle Rdb for OpenVMS Alpha kit that is compiled to run on all Alpha platforms:

```
@SYS$UPDATE:VMSINSTAL RDBV72400AM device-name
```

- To install the Oracle Rdb for OpenVMS Alpha kit that is performance targeted for Alpha EV56 and later platforms:

```
@SYS$UPDATE:VMSINSTAL RDBV72401AM device-name
```

*device-name*

Use the name of the device on which the media is mounted. If the device is a disk-type drive, you also need to specify a directory. For example: *DKA400:[RDB.KIT]*

## 1.8 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.

## 1.9 After Installing Oracle Rdb

This update provides a new Oracle TRACE facility definition for Oracle Rdb. 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, "RDBVMSV7.2".

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=RDBVMSV7.2 -  
_ $ /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 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.

## 1.10 VMS\$MEM\_RESIDENT\_USER Rights Identifier Required

Oracle Rdb Version 7.1 introduced additional privilege enforcement for the database or row cache attributes RESIDENT, SHARED MEMORY IS SYSTEM and LARGE MEMORY IS ENABLED. If a database utilizes any of these features, then the user account that opens the database must be granted the VMS\$MEM\_RESIDENT\_USER rights identifier.

Oracle recommends that the RMU/OPEN command be used when utilizing these features.

# 1.11 Installation, Configuration, Migration, Upgrade Suggestions

Oracle Rdb Release 7.2 fully supports mixed–architecture clusters for AlphaServer systems and HP Integrity servers.

In certain development environments, it may be helpful to incorporate a VAX system into the AlphaServer systems and HP Integrity servers cluster. While HP and Oracle believe that in most cases this will not cause problems to the computing environment, we have not tested it extensively enough to provide support. It is possible that VAX systems in a cluster may cause a problem with the cluster performance or stability. Should this happen, the VAX systems in the cluster which are causing the difficulty should be removed.

Oracle continues to support mixed architecture clusters of VAX systems and AlphaServer systems with direct database access using Rdb V7.0. Oracle Rdb V7.1 runs natively on Alpha systems and clusters. All Rdb versions include a built–in remote network database server allowing cross–architecture and cross–version application and database access.

When moving applications from existing Alpha or VAX configurations to new environments containing Integrity Server systems, there are numerous possible paths depending on the requirements of individual sites. In general, this can be as straightforward as adding a new node to an already existing AlphaServer systems cluster or standalone system, except the node is an HP Integrity server. [Table 1–1, Migration Suggestions](#), considers several possible situations and recommended steps to take.

**Table 1–1 Migration Suggestions**

Case	You Wish To...	You should...
1	Add an Integrity server to an existing cluster of Alpha servers	<ol style="list-style-type: none"> <li>1. Verify database(s) using RMU/VERIFY/ALL.</li> <li>2. Backup database(s) using RMU/BACKUP.</li> <li>3. Install Rdb 7.2 on Integrity and Alpha nodes.</li> <li>4. Convert database(s) to the Rdb 7.2 structure level using RMU/CONVERT.</li> <li>5. Verify database(s) again using RMU/VERIFY/ALL.</li> <li>6. Backup database(s) using RMU/BACKUP.</li> <li>7. Access database(s) from Alpha and Integrity directly by specifying database root file specification(s) in SQL ATTACH statements.</li> </ol>
2	Add an Integrity server to an existing mixed cluster of VAX and Alpha nodes and access an Rdb database	<ol style="list-style-type: none"> <li>1. Verify database(s) using</li> </ol>

	<p>from all nodes. Disks used for the database are accessible from all nodes.</p>	<ol style="list-style-type: none"> <li>RMU/VERIFY/ALL.</li> <li>2. Backup database(s) using RMU/BACKUP.</li> <li>3. Install Rdb 7.2 on Integrity and Alpha nodes.</li> <li>4. Convert database(s) to the Rdb 7.2 structure level using RMU/CONVERT.</li> <li>5. Verify database(s) again using RMU/VERIFY/ALL.</li> <li>6. Backup database(s) using RMU/BACKUP.</li> <li>7. Access database(s) from Alpha and Integrity nodes directly by specifying database root file specification(s) in SQL ATTACH statements.</li> <li>8. Access the database from VAX node(s) using the Rdb built-in network server (remote database) by specifying one of the Alpha or Integrity node names in SQL ATTACH statements.</li> <li>9. After thorough testing, remove VAX nodes from the cluster.</li> </ol>
3	<p>Move database(s) to new disks and add an Integrity server to an existing cluster.</p>	<ol style="list-style-type: none"> <li>1. Use RMU/COPY with an options file to move the database files to the new disks.</li> <li>2. Follow the steps for case 1 or case 2.</li> </ol>
4	<p>Continue to use Rdb primarily from VAX or Alpha nodes using earlier releases. Add an Integrity server for application testing purposes.</p>	<ol style="list-style-type: none"> <li>1. Install Rdb 7.2 on Integrity node.</li> <li>2. Access existing database(s) from Integrity node by specifying one of the Alpha or VAX node names in the SQL ATTACH statements.</li> <li>3. When testing is complete, follow the steps in case 1 or case 2.</li> </ol>
5	<p>Add an Integrity server to an existing cluster of Alpha servers or Create a new cluster from an existing stand-alone Alpha server by adding one or more new Integrity servers.</p>	<ol style="list-style-type: none"> <li>1. Verify database(s) using RMU/VERIFY/ALL.</li> <li>2. Backup database(s) using RMU/BACKUP.</li> <li>3. Install Rdb 7.2 on Integrity and Alpha nodes.</li> </ol>



		<ol style="list-style-type: none"> <li>4. Convert database(s) to the Rdb 7.2 structure level using RMU/CONVERT.</li> <li>5. Verify database(s) again using RMU/VERIFY/ALL.</li> <li>6. Backup database(s) using RMU/BACKUP.</li> <li>7. Access database(s) from Alpha and Integrity directly by specifying database root file specification in the SQL ATTACH statements.</li> </ol>
6	<p>Create a new stand-alone Integrity Server system or cluster of Integrity Servers and move database(s) to the new environment.</p>	<ol style="list-style-type: none"> <li>1. Verify database(s) using RMU/VERIFY/ALL.</li> <li>2. Install Rdb 7.2 on new system(s).</li> <li>3. Back up database(s) on the existing cluster using RMU/BACKUP.</li> <li>4. Copy backup file(s) to the new system (or, if using tape media, make the tapes available to the new system).</li> <li>5. Restore database(s) on the new system using RMU/RESTORE specifying the location of each database file in an options file.</li> <li>6. Verify the new database using RMU/VERIFY/ALL.</li> </ol>

Refer to the Oracle Rdb documentation set for additional information and detailed instructions for using RMU and remote databases.

Note that database parameters might need to be altered in the case of accessing a database from a larger number of systems in a cluster.

---

# **Chapter 2**

## **Software Errors Fixed in Oracle Rdb Release 7.2.4.0**

This chapter describes software errors that are fixed by Oracle Rdb Release 7.2.4.

## 2.1 Software Errors Fixed That Apply to All Interfaces

### 2.1.1 Bugcheck With SYSTEM-F-ROPRAND in KODTXN\$POST\_TSNBLK\_UPDATE

Bugs 7193991 and 8796832

It is possible for a process to bugcheck due to a corrupt internal queue header with a reserved operand fault in the routine KODTXN\$POST\_TSNBLK\_UPDATE. Correcting the memory corruption requires closing and re-opening the database. This problem will generally present itself with a bugcheck exception "footprint" similar to the following:

```
***** Exception at 00000000815C89D2 : RDMSHRP72\KODTXN$POST_TSNBLK_UPDATE
+ 000000D2
%SYSTEM-F-ROPRAND, reserved operand fault at PC=00000000815C89D2, PS=00000009
Saved PC = 00000000815C3EF0 : RDMSHRP72\KOD$START + 00000D40
```

Analysis of the bugcheck dump file will often indicate that at least one of the "TUPB\_RELQHD" queue headers will contain an entry containing "00000000:00000001" as in the following example where entry 13 has been corrupted:

```
TUPB_RELQHD_VEC[1.] @04910A00 = 00000000:00000000 (04910A00:04910A00)
TUPB_RELQHD_VEC[2.] @04910A08 = 00000000:00000000 (04910A08:04910A08)
TUPB_RELQHD_VEC[3.] @04910A10 = 00000000:00000000 (04910A10:04910A10)
TUPB_RELQHD_VEC[4.] @04910A18 = 00000000:00000000 (04910A18:04910A18)
TUPB_RELQHD_VEC[5.] @04910A20 = 00000000:00000000 (04910A20:04910A20)
TUPB_RELQHD_VEC[6.] @04910A28 = 00000000:00000000 (04910A28:04910A28)
TUPB_RELQHD_VEC[7.] @04910A30 = 00000000:00000000 (04910A30:04910A30)
TUPB_RELQHD_VEC[8.] @04910A38 = 00000000:00000000 (04910A38:04910A38)
TUPB_RELQHD_VEC[9.] @04910A40 = 00000000:00000000 (04910A40:04910A40)
TUPB_RELQHD_VEC[10.] @04910A48 = 00000000:00000000 (04910A48:04910A48)
TUPB_RELQHD_VEC[11.] @04910A50 = 00000000:00000000 (04910A50:04910A50)
TUPB_RELQHD_VEC[12.] @04910A58 = 00000000:00000000 (04910A58:04910A58)
TUPB_RELQHD_VEC[13.] @04910A60 = 00000000:00000001 (04910A60:04910A61)
```

The cause of the problem was related to an incorrect synchronization between processes manipulating a relative memory queue within the database global section. This (or related) problem may impact all Rdb databases that perform transactions by more than one database user. The possibility of corruption increases with higher transaction rates and when there are more database users performing transactions. Generally, databases configured for less than 168 users will never see this symptom and databases with fewer than 168 simultaneous users will never see this problem.

Oracle recommends that all Rdb installations upgrade to at least Oracle Rdb Release 7.2.4 to implement the correction to this problem.

This problem has been corrected in Oracle Rdb Release 7.2.4. The shared memory access to "TUPB\_RELQHD" queue headers is now correctly synchronized.

## 2.1.2 Memory Leak On Systems With RAD Support Enabled

Bug 8410893

When running on an Alpha system with RAD support enabled, a memory leak was possible during database detach/attach sequences. The size of the leak would be related to the size of the per-RAD database statistics global section.

As a possible workaround, the system parameter RAD\_SUPPORT can be set to zero.

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.1.3 Incorrect Messages From RMU /MONITOR START When RDM\$MON\_USERNAME Specifies Non-existent Account

Bug 8420114

In prior releases, when the logical name RDM\$MON\_USERNAME was defined to specify a non-existent username, the RMU /MONITOR START command would return incorrect or misleading messages. The following example shows such output:

```
$ RMU /MONITOR START
%RMU-F-CANTCREMON, unable to start database monitor process
%F
%?-RESULTOVF
%F-CTRLERR
%I-POWERFAIL
-PLI
%NOMSG, Message number 0053474E
%BADPARAM
%NORMAL
Message number 00081C48
-INPSMB
-JBC, normal successful completion
-TRACE, Message number 0009A930
%F-NOMSG, Message number 4F525024
```

As a workaround, either do not define RDM\$MON\_USERNAME or make sure that it specifies a valid username.

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.1.4 Potential INVEXCEPTN System Crash Using LARGE MEMORY IS ENABLED or SHARED MEMORY IS SYSTEM on Itanium

Bug 8541571

When closing a database when using the SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED feature on some Itanium systems, it was possible for Oracle Rdb to cause an OpenVMS system crash with a bugcheck type of "INVEXCEPTN, Exception while above ASTDEL". The INVEXCEPTN can be triggered by an invalid access to the OpenVMS PTE database.

The following OpenVMS crash "footprint" is a result of this problem:

```

Bugcheck Type:      INVEXCEPTN, Exception while above ASTDEL
Node:               RDBI64      (Cluster)
CPU Type:           HP rx8640   (1.60GHz/12.0MB)
VMS Version:        V8.3-1H1
Current Process:    RDMS_MONITOR72
Current Image:      $1$DGA4000:[SYS2.SYSCOMMON.][SYSEXE]RDMMON72.EXE;3
Failing PC:         00000000.0057DDD0      DEALLOC_PHY_BUFFER_C+001C0
Failing PS:         00000000.00000803
Module:             RDMPRV72
Offset:             000A7DD0
    
```

This problem only effects those Itanium systems with physical memory accessible to the system of PFNs greater than 268,435,455 (Hex 0FFFFFFF). The SDA command CLUE CONFIG can be used to display the system memory configuration.

The workaround to this problem for such systems with PFNs greater than 268,435,455, is to discontinue use of the LARGE MEMORY IS ENABLED and SHARED MEMORY IS SYSTEM features, if enabled, by using one, or both, of the following SQL commands:

- ALTER DATABASE ... GLOBAL BUFFERS ...LARGE MEMORY IS DISABLED
- ALTER DATABASE ... SHARED MEMORY IS PROCESS RESIDENT

This problem has been corrected in Oracle Rdb Release 7.2.4. The Oracle Rdb VLM feature correctly processes 64-bit PFN values when closing databases using the SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED feature on Itanium systems.

## 2.1.5 Wrong Result From Query With Derived Table

Bug 8475230

The following query returns the wrong result: 4 rows instead of 5 rows.

```

SELECT TRANS_DATE
FROM
(
SELECT TRANS_DATE
FROM T1 T1
WHERE
T1.TRANS_DATE <= '20081231' AND
T1.TRANS_DATE > '20081222' AND
T1.TRANS_DATE NOT IN
(
SELECT TRANS_DATE
FROM T2
WHERE
T2_FLAG = 'Y' AND
CODE = 'FBAR' AND
TRANS_DATE > '20081222' AND
TRANS_DATE <= '20081231'
    
```

```

)
AND
T1.TRANS_DATE NOT IN
(
  SELECT TRANS_DATE
    FROM V1          -- another view V1
   WHERE
     IS_FLAG = 'Y' AND
     CODE = 'FBAR' AND
     TRANS_DATE > '20081222' AND
     TRANS_DATE <= '20081231'
)
ORDER BY T1.TRANS_DATE DESC) DT;
Tables:
  0 = T1
  1 = T2
  2 = T2
  3 = T3
  4 = T4
  5 = T3
Merge of 1 entries  Q1
Merge block entry 1  Q2
Cross block of 3 entries  Q2
  Cross block entry 1
    Get      Retrieval by index of relation 0:T1
      Index name  T1_IDX [1:1]          Reverse Scan
      Keys: (0.TRANS_DATE > '20081222') AND (0.TRANS_DATE <=
        '20081231')
  Cross block entry 2
    Conjunct: <agg0> = 0
    Aggregate-F1: 0:COUNT-ANY (<subselect>) Q3
    Conjunct: (1.T2_FLAG = 'Y') AND (1.CODE = 'FBAR') AND (MISSING (
      0.TRANS_DATE) OR MISSING (1.TRANS_DATE) OR (
      0.TRANS_DATE = 1.TRANS_DATE))
    Get      Retrieval by index of relation 1:T2
      Index name  T2_IDX [1:1]
      Keys: (1.TRANS_DATE > '20081222') AND (1.TRANS_DATE <=
        '20081231')
  Cross block entry 3
    Conjunct: <agg1> = 0
    Aggregate-F1: 1:COUNT-ANY (<subselect>) Q4
    Conjunct: (CASE (WHEN (Q6-CAGG"=AGG_TOTAL"<var0> = Q6-CAGG"=AGG_TOTAL"
      <var1>) THEN 'Y' ELSE 'N') = 'Y') AND (2.CODE = 'FBAR')
      AND (2.TRANS_DATE > '20081222') AND (2.TRANS_DATE <=
      '20081231') AND (MISSING (0.TRANS_DATE) OR MISSING (
      2.TRANS_DATE) OR (0.TRANS_DATE = 2.TRANS_DATE))
    Aggregate: 2:SUM (<var1>) Q6
      3:SUM (<var0>) Q6
  Cross block of 3 entries  Q6
    Cross block entry 1
      Get      Retrieval by index of relation 2:T2  <--See NOTE
        Index name  T2_IDX [1:1]
        Keys: (2.TRANS_DATE > '20081222') AND (2.TRANS_DATE <=
          '20081231')
    Cross block entry 2
      Merge of 1 entries  Q6
      Merge block entry 1  Q9
      Aggregate: 4:COUNT (*) Q10
      Conjunct: (5.DONE = 'Y') AND (5.STATUS = 'N')
      Get      Retrieval by index of relation 5:T3
        Index name  T3_IDX [3:3]
        Keys: (2.TRANS_DATE = 5.TRANS_DATE) AND (2.CITY = 5.CITY) AND

```

## Oracle® Rdb for OpenVMS

```
(2.TOWN = 5.TOWN)
Cross block entry 3
Merge of 1 entries Q6
Merge block entry 1 Q7
Aggregate: 5:COUNT (*) Q8
Cross block of 2 entries Q8
Cross block entry 1
  Conjunct: 4.NON_FLAG = 'N'
  Conjunct: 4.NON_REASON = 'P'
  Get      Retrieval by index of relation 4:T4
           Index name  T4_IDX [1:1]
           Keys: 2.TRANS_DATE = 4.CURR_DATE
Cross block entry 2
  Conjunct: (2.MARKET = 3.MARKET) AND (2.CITY = 3.CITY) AND
           (2.TOWN = 3.TOWN) AND (3.DONE = 'Y') AND
           (3.STATUS = 'N')
  Get      Retrieval by index of relation 3:T3
           Index name  T3_IDX [2:2]   Direct lookup
           Keys: (2.TRANS_DATE = 3.TRANS_DATE) AND (3.PRODUCT =
                 4.PRODUCT)

TRANS_DATE
20081230
20081229
20081224
20081223
4 rows selected
```

NOTE:: The conjuncts are missing to filter the retrieval of context "2:T2" under Cross block entry 1 and causes the query to return the wrong result.

The query works if the derived table is removed from the query, as in the following example.

```
SELECT TRANS_DATE
FROM    T1 T1
WHERE
  T1.TRANS_DATE <= '20081231' AND
  T1.TRANS_DATE > '20081222' AND
  T1.TRANS_DATE NOT IN
  (
    SELECT TRANS_DATE
    FROM T2
    WHERE
      T2_FLAG = 'Y' AND
      CODE = 'FBAR' AND
      TRANS_DATE > '20081222' AND
      TRANS_DATE <= '20081231'
  )
AND
T1.TRANS_DATE NOT IN
(
  SELECT TRANS_DATE
  FROM V1          -- another view V1
  WHERE
    IS_FLAG = 'Y' AND
    CODE = 'FBAR' AND
    TRANS_DATE > '20081222' AND
    TRANS_DATE <= '20081231'
)
ORDER BY T1.TRANS_DATE DESC;
Tables:
0 = T1
```

## Oracle® Rdb for OpenVMS

```
1 = T2
2 = T2
3 = T3
4 = T4
5 = T3
Cross block of 3 entries Q1
Cross block entry 1
  Get      Retrieval by index of relation 0:T1
          Index name  T1_IDX [1:1]      Reverse Scan
          Keys: (0.TRANS_DATE > '20081222') AND (0.TRANS_DATE <= '20081231')
Cross block entry 2
  Conjunct: <agg0> = 0
  Aggregate-F1: 0:COUNT-ANY (<subselect>) Q2
  Conjunct: (1.T2_FLAG = 'Y') AND (1.CODE = 'FBAR') AND (MISSING (
            0.TRANS_DATE) OR MISSING (1.TRANS_DATE) OR (0.TRANS_DATE
            = 1.TRANS_DATE))
  Get      Retrieval by index of relation 1:T2
          Index name  T2_IDX [1:1]
          Keys: (1.TRANS_DATE > '20081222') AND (1.TRANS_DATE <= '20081231')
Cross block entry 3
  Conjunct: <agg1> = 0
  Aggregate-F1: 1:COUNT-ANY (<subselect>) Q3
  Conjunct: (CASE (WHEN (Q5-CAGG"=AGG_TOTAL"<var0> = Q5-CAGG"=AGG_TOTAL"<var1>
                    ) THEN 'Y' ELSE 'N') = 'Y') AND (MISSING (0.TRANS_DATE) OR
                    MISSING (2.TRANS_DATE) OR (0.TRANS_DATE = 2.TRANS_DATE))
  Aggregate: 2:SUM (<var1>) Q5
            3:SUM (<var0>) Q5
Cross block of 3 entries Q5
  Cross block entry 1
    Conjunct: (2.CODE = 'FBAR') AND (2.TRANS_DATE > '20081222')
              AND (2.TRANS_DATE <= '20081231')      <--- See NOTE
  Get      Retrieval by index of relation 2:T2
          Index name  T2_IDX [1:1]
          Keys: (2.TRANS_DATE > '20081222') AND (2.TRANS_DATE <=
              '20081231')
  Cross block entry 2
    Merge of 1 entries Q5
    Merge block entry 1 Q8
    Aggregate: 4:COUNT (*) Q9
    Conjunct: (5.DONE = 'Y') AND (5.STATUS = 'N')
    Get      Retrieval by index of relation 5:T3
          Index name  T3_IDX [3:3]
          Keys: (2.TRANS_DATE = 5.TRANS_DATE) AND (2.CITY = 5.CITY) AND
              (2.TOWN = 5.TOWN)
  Cross block entry 3
    Merge of 1 entries Q5
    Merge block entry 1 Q6
    Aggregate: 5:COUNT (*) Q7
  Cross block of 2 entries Q7
    Cross block entry 1
      Conjunct: 4.NON_FLAG = 'N'
      Conjunct: 4.NON_REASON = 'P'
      Get      Retrieval by index of relation 4:T4
            Index name  T4_IDX [1:1]
            Keys: 2.TRANS_DATE = 4.CURR_DATE
    Cross block entry 2
      Conjunct: (2.MARKET = 3.MARKET) AND (2.CITY = 3.CITY) AND
              (2.TOWN = 3.TOWN) AND (3.DONE = 'Y') AND
              (3.STATUS = 'N')
      Get      Retrieval by index of relation 3:T3
            Index name  T3_IDX [2:2]      Direct lookup
            Keys: (2.TRANS_DATE = 3.TRANS_DATE) AND (3.PRODUCT =
```



```

                                4.PRODUCT)
TRANS_DATE
20081231
20081230
20081229
20081224
20081223
5 rows selected

```

NOTE:: The conjuncts appear correctly in the retrieval of context "2:T2" under Cross block entry 1.

The workaround is to replace the "NOT IN" clause with "NOT EXISTS" in the query, as in the following example.

```

SELECT TRANS_DATE
FROM
(SELECT TRANS_DATE
FROM T1 T1
WHERE
T1.TRANS_DATE <= '20081231' AND
T1.TRANS_DATE > '20081222' AND
T1.TRANS_DATE NOT IN
(
SELECT TRANS_DATE
FROM T2
WHERE
T2_FLAG = 'Y' AND
CODE = 'FBAR' AND
TRANS_DATE > '20081222' AND
TRANS_DATE <= '20081231'
)
AND
T1.TRANS_DATE NOT IN
(
SELECT TRANS_DATE
FROM V1 -- another view V1
WHERE
IS_FLAG = 'Y' AND
CODE = 'FBAR' AND
TRANS_DATE > '20081222' AND
TRANS_DATE <= '20081231'
)
ORDER BY T1.TRANS_DATE DESC) DT;

```

Tables:

```

0 = T1
1 = T2
2 = T2
3 = T3
4 = T4
5 = T3

```

Cross block of 3 entries Q1

Cross block entry 1

Aggregate-F1: 0:COUNT-ANY (<subselect>) Q4

Conjunct: CASE (WHEN (Q6-CAGG"=AGG\_TOTAL"<var0> = Q6-CAGG"=AGG\_TOTAL"<var1>) THEN 'Y' ELSE 'N') = 'Y'

Aggregate: 1:SUM (<var1>) Q6

2:SUM (<var0>) Q6

Cross block of 3 entries Q6

Cross block entry 1

Conjunct: (2.CODE = 'FBAR') AND (2.TRANS\_DATE > '20081222') AND (2.TRANS\_DATE <= '20081231') <-- See NOTE

## Oracle® Rdb for OpenVMS

```
Get      Retrieval by index of relation 2:T2
Index name  T2_IDX [1:1]
Keys: (2.TRANS_DATE > '20081222') AND (2.TRANS_DATE <=
      '20081231')
Cross block entry 2
Merge of 1 entries  Q6
Merge block entry 1  Q9
Aggregate: 3:COUNT (*) Q10
Conjunct: (5.T2_FLAG = 'Y') AND (5.STATUS = 'N')
Get      Retrieval by index of relation 5:T3
Index name  T3_IDX [3:3]
Keys: (2.TRANS_DATE = 5.TRANS_DATE) AND (2.CITY =
      5.CITY) AND (2.TOWN = 5.TOWN)
Cross block entry 3
Merge of 1 entries  Q6
Merge block entry 1  Q7
Aggregate: 4:COUNT (*) Q8
Cross block of 2 entries  Q8
Cross block entry 1
Conjunct: 4.NON_FLAG = 'N'
Conjunct: 4.NON_REASON = 'P'
Get      Retrieval by index of relation 4:T4
Index name  T4_IDX [1:1]
Keys: 2.TRANS_DATE = 4.CCUR_DATE
Cross block entry 2
Conjunct: (2.MARKET = 3.MARKET) AND (2.CITY =
      3.CITY) AND (2.TOWN = 3.TOWN
      ) AND (3.T2_FLAG = 'Y') AND (3.STATUS = 'N')
Get      Retrieval by index of relation 3:T3
Index name  T3_IDX [2:2]  Direct lookup
Keys: (2.TRANS_DATE = 3.TRANS_DATE) AND (3.PRODUCT =
      4.PRODUCT)
Cross block entry 2
Aggregate-F1: 5:COUNT-ANY (<subselect>) Q3
Conjunct: (1.HOLIDAY = 'Y') AND (1.CODE = 'FBAR')    <-- See NOTE
Get      Retrieval by index of relation 1:T2
Index name  T2_IDX [1:1]
Keys: (1.TRANS_DATE > '20081222') AND (1.TRANS_DATE <= '20081231')
Cross block entry 3
Merge of 1 entries  Q1
Merge block entry 1  Q2
Conjunct: (<agg5> = 0) AND (<agg0> = 0)
Get      Retrieval by index of relation 0:T1
Index name  T1_IDX [1:1]  Reverse Scan
Keys: (0.TRANS_DATE > '20081222') AND (0.TRANS_DATE <=
      '20081231')

TRANS_DATE
20081231
20081230
20081229
20081224
20081223
5 rows selected
```

NOTE:: The conjuncts appear correctly in the retrieval of context "2:T2" and "1:T2".

This problem was introduced in Oracle Rdb Release 7.2.3.0.

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.1.6 Application Looses Virtual Memory (Memory Leak)

Bugs 7697133 and 8449605

An application could run out of virtual memory or pagefile quota when executing multiple queries within a single attach to a database and when these queries (or stored procedures) accessed indices.

This happened only on Integrity systems and was caused by allocating some small index key structures from a memory pool which is only freed at the end of the session (during database disconnect).

This loss of virtual memory increases if SAMPLED SELECTIVITY is used (OPTIMIZE clause of a query, SET FLAGS 'SELECTIVITY(2), or the SET OPTIMIZATION LEVEL statement).

These problems have been corrected in Oracle Rdb Release 7.2.4.

## 2.1.7 VMS\$BUFFER\_OBJECT\_USER Not Always Checked

Bug 8464087

In some cases, the OpenVMS rights identifier VMS\$BUFFER\_OBJECT\_USER was not being required for database buffer object use. This could allow users to utilize the buffer objects feature even though they did not have the identifier granted.

This problem has been corrected in Oracle Rdb Release 7.2.4. Users attempting to utilize buffer object features now must always have the VMS\$BUFFER\_OBJECT\_USER rights identifier granted.

It is possible that applications that worked previously when using the buffer objects feature with processes that did not have the VMS\$BUFFER\_OBJECT\_USER rights identifier granted may now correctly fail with messages similar to the following:

```
%SQL-F-ERRATTDEF, Could not use database file specified by SQL$DATABASE
-RDB-F-SYS_REQUEST, error from system services request
-RDMS-F-CANTCREBOB, error creating Buffer Object
-SYSTEM-E-NOBUFOBJID, requires rights identifier VMS$BUFFER_OBJECT_USER
```

These messages indicate that the process must be granted the VMS\$BUFFER\_OBJECT\_USER rights identifier before attempting to use the buffer objects feature. In many cases, granting the identifier to the user account and then logging in again will resolve the issue.

## 2.1.8 Deleted Space in Uniform Areas Not Reclaimed by Other Users When Erased Rows Moved From Row Cache to Disk

Bug 8522094

Oracle Rdb Release 7.2 introduced a mechanism that allows database users on the same cluster node to share information regarding the availability of free space. When a user chooses a location to store new rows, the location is stored in the database global section so that other users can use that location as a starting point when searching for available space. When a user deletes rows from a table, if the location of the deleted rows

is closer to the beginning of the storage area than the last page used for an insert, then the starting page for the next insert is updated to the location of the lowest page that had rows deleted.

In some cases when using the Row Cache feature, when an erased row was moved from cache back to a page in the database, the shared information regarding the availability of free space was not being updated. This caused users to be unaware of possible free space and unexpected storage area extension could result.

This problem has been corrected in Oracle Rdb Release 7.2.4. The shared information regarding the availability of free space is now updated when erased rows are moved from a row cache back to the database page.

## 2.1.9 LOWER Function Problem When Using ISOLATINCYRILLIC Character Set

Bug 8605022

A problem in the casing tables for ISOLATINCYRILLIC would produce wrong lowercase values for the following two cyrillic characters:

- UPPERCASE ER ( hex C0 )
- UPPERCASE SHA ( hex C8 )

This problem has been corrected in Oracle Rdb Release 7.2.4. The LOWER function now produces the correct lowercase characters for ISOLATINCYRILLIC.

## 2.1.10 Query Bugchecks with SYSTEM-F-FLTINV

Bugs 8580585 and 7649113

The following query bugchecks with SYSTEM-F-FLTINV error when sampled selectivity is applied.

```

SELECT DISTINCT 7
  FROM R_ILS_STATUS S , R_LS L
 WHERE S.A_DATE >= 20090531 AND S.A_DATE <= 20090531
 AND S.A_NLS = L.A_NLS
 AND NOT EXISTS
      ( SELECT 1 FROM R_ILS I where I.A_NLS=S.a_nls
        AND I.A_DOP=s.a_date )
 OPTIMIZE WITH SAMPLED SELECTIVITY ;
%RDB-E-ARITH_EXCEPT, truncation of a numeric value at runtime
-SYSTEM-F-HPARITH, high performance arithmetic trap, Imask=00000000,
Fmask=00000001, summary=02, PC=0000000008D0FB8, PS=0000001B
-SYSTEM-F-FLTINV, floating invalid operation, PC=0000000008D0FB8, PS=0000001B
    
```

The query works if sampled selectivity is NOT applied, as in the following example.

```

SELECT DISTINCT 7
  FROM R_ILS_STATUS S , R_LS L
 WHERE S.A_DATE >= 20090531 AND S.A_DATE <= 20090531
 AND S.A_NLS = L.A_NLS
 AND NOT EXISTS
      ( SELECT 1 FROM R_ILS I where I.A_NLS=S.a_nls
        AND I.A_DOP=s.a_date )
    
```

```

! OPTIMIZE WITH SAMPLED SELECTIVITY
;
Tables:
  0 = R_ILS_STATUS
  1 = R_LS
  2 = R_ILS
Reduce: 7
Sort: 7(a)
Cross block of 3 entries  Q1
  Cross block entry 1
    Index only retrieval of relation 1:R_LS
    Index name  SI_LS_9 [0:0]
  Cross block entry 2
    Leaf#01 BgrOnly 0:R_ILS_STATUS Card=608708
    Bool: (0.A_DATE >= 20090531) AND (0.A_DATE <= 20090531) AND (0.A_NLS =
      1.A_NLS)
    BgrNdx1 I_R_ILS_STATUS [1:1] Fan=13
    Keys: 0.A_NLS = 1.A_NLS
    BgrNdx2 I_R_ILS_STATUS2 [1:1] Fan=17
    Keys: (0.A_DATE >= 20090531) AND (0.A_DATE <= 20090531)
  Cross block entry 3
    Conjunct: <agg0> = 0
    Aggregate-F1: 0:COUNT-ANY (<subselect>) Q2
    Index only retrieval of relation 2:R_ILS
    Index name  SI_ILS_NDK [2:2]  Index counts lookup
    Keys: (2.A_NLS = 0.A_NLS) AND (2.A_DOP = 0.A_DATE)
0 rows selected

```

This problem occurs when the database contains empty tables or indices and the query uses the optimizer with sampled selectivity.

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.1.11 Bugcheck Altering Storage Map

Bug 8573340

A bugcheck could occur when executing an alter storage map using thresholds and compression.

This problem only occurred under OpenVMS Itanium.

A typical alter storage map SQL statement:

```

ALTER STORAGE MAP STORE_MAP
  STORE IN DATA001_AREA
    (THRESHOLDS ARE (80,90,95))
  ENABLE COMPRESSION
  REORGANIZE;

```

```

***** Exception at 000000008095F620 : RDMSHRP721\RDMS$$EXE_NEXT + 00002710
%COSI-F-BUGCHECK, internal consistency failure

```

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.1.12 Wrong Results when Expression Shared Between Aggregate Filters

Bug 8553416

In prior versions of Oracle Rdb V7.2, queries that included shared expressions between aggregate filters might return the wrong results.

The following example shows such a query with the expression "substring(postal\_code from 1 for 3)" shared between the two count filters.

```
select
  count(*) filter
    (where
      substring(postal_code from 1 for 1) <> '0'
      or substring(postal_code from 3 for 1) <> '3'
      or substring(postal_code from 1 for 3) = '034'
    ) as field_a
, count(*) filter
    (where
      employee_id < '00300' and
      substring(postal_code from 1 for 3) <> '034'
    ) as field_b
from
  employees e, departments d
where
  e.employee_id = d.manager_id
  and d.department_code starting with 'SU';
```

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.1.13 Optimize For Sequential Access Using Index

Bug 8746157

In prior releases of Oracle Rdb, some UPDATE statements would ignore the "Optimize for Sequential Access" clause. See the following example.

```
UPDATE T1 SET VALUE = 'A'
FROM T1
  WHERE ID = 1
  AND COUNTRY = 12
  AND MARKET = 20
  AND GROUP = 6
  AND TRAN_NO = 5710175
OPTIMIZE FOR SEQUENTIAL ACCESS;
Get      Temporary relation      Retrieval by index of relation 0:T1
  Index name  T1_IDX [5:5]      Direct lookup
~S: Full compliance with the outline was not possible
```

This problem has been corrected in Oracle Rdb Release 7.2.4. Rdb now correctly processes the "Optimize for Sequential Access" in such cases.

## 2.1.14 Query With Match Strategy for Left Outer Join Bugchecks

Bug 8709430

In prior releases of Oracle Rdb, it was possible for queries with match strategy for a left outer join operation to generate a bugcheck dump with the following footprint.

```
***** Exception at 000000000688627C : RDMSHRP72\RDMS$$EXE_CREATE_TTBL_FILE +
0000214C
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual address=..., PC=
..., PS=...
Saved PC = 0000000006880CEC : RDMSHRP72\RDMS$$EXE_NEXT + 0000110C
Saved PC = 000000000687FFD0 : RDMSHRP72\RDMS$$EXE_NEXT + 000003F0
Saved PC = 000000000688445C : RDMSHRP72\RDMS$$EXE_CREATE_TTBL_FILE + 0000032C
Saved PC = 000000000688451C : RDMSHRP72\RDMS$$EXE_CREATE_TTBL_FILE + 000003EC
Saved PC = 00000000068800D8 : RDMSHRP72\RDMS$$EXE_NEXT + 000004F8
Saved PC = 0000000006880808 : RDMSHRP72\RDMS$$EXE_NEXT + 00000C28
```

An example query with these similar characteristics follows.

```
select *
from
  (select
    (select count(*)
     from T1 t1 left outer join T2 t2
     on (t2.COL1 = t1.COL1) and (t2.COL2 = t1.COL2)
     where (t1.COL3 = c2.COL3) and (t1.COL4 = 0)
    )
    ,c2.*
    ,0
  from T5 c2
 )
--as c1 (f0, f2, f3)
where
  (1 = 1) and (COL5 = 1);
%RDMS-I-BUGCHKDMP, generating bugcheck dump file [directory]RDSBUGCHK.DMP;
```

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.1.15 Bugchecks at DIOFETCH\$FETCH\_SNAP\_SEG or DIOCCH\$FETCH\_SNAP\_SEG

In some cases, when using read-only transactions along with read-write transactions that request exclusive access to a table, the read-only transaction may bugcheck with a "footprint" similar to the following:

```
***** Exception at 0000000001948194 : RDMSHRP72\DIOFETCH$FETCH_SNAP_SEG
+ 000003C4
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 00000000019488B4 : RDMSHRP72\DIOFETCH$FETCH_VISIBLE_SEG + 00000684
Saved PC = 00000000019492FC : RDMSHRP72\DIOFETCH$FETCH_ONE_LINE + 000008AC
Saved PC = 0000000001949CCC : RDMSHRP72\DIOFETCH$SCAN_ONE_PAGE + 000002AC
Saved PC = 000000000194A228 : RDMSHRP72\DIO$FETCH_AREA + 00000228
Saved PC = 000000000174616C : RDMSHRP72\RDMS$$EXE_NEXT + 00001D9C
Saved PC = 0000000001744A24 : RDMSHRP72\RDMS$$EXE_NEXT + 00000654
```

## Oracle® Rdb for OpenVMS

```
Saved PC = 0000000001746748 : RDMSHRP72\RDMS$$EXE_NEXT + 00002378
Saved PC = 0000000002081F04 : symbol not found
Saved PC = 000000000175F9C8 : RDMSHRP72\RDMS$TOP_START_REQUEST + 00000878
```

This problem would only occur on databases with a transaction sequence number larger than 2,147,483,647. The following shows an example sequence that could lead to the bugcheck:

```
SESSION 1:                                SESSION 2:

SQL$
SQL> ATTACH 'FILE FOO';
SQL> SET TRANS READ ONLY;
SQL> SHOW TRANSACTION;

SQL$
SQL> ATTACH 'FILE FOO';
SQL> SET TRANS READ WRITE
      RESER C1 FOR EXCL WRITE;
SQL> SHOW TRANSACTION;
SQL> DELETE FROM C1;
SQL> COMMIT;

SQL> SELECT * FROM C1;
%RDMS-I-BUGCHKDMP
```

This problem has been corrected in Oracle Rdb Release 7.2.4. The invalid access by the read-only transaction is correctly detected and the expected error is generated:

```
SQL> SELECT * FROM C1;
%RDB-F-LOCK_CONFLICT, request failed due to locked resource
-RDMS-F-CANTSNAPE, can't ready storage area $1$DGA1:[DB]FOO.RDB;1 for snapshots
SQL>
```

### 2.1.16 Bugcheck in PSII2INSERTDUPBBC

In prior versions of Oracle Rdb, a problem in establishing correct index node currency during the insertion of a duplicate record dbkey into a sorted ranked index might result in a bugcheck with a "footprint" similar to the following:

```
***** Exception at 0000000000321770 : RDMSHRP72\PSII2INSERTDUPBBC +00002560
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 000000000031A950 : SQLU721\PSII2INSERTBOTTOM + 00000B80
Saved PC = 0000000000301980 : SQLU721\PSII2INSERTTT + 00000400
Saved PC = 0000000000304770 : SQLU721\PSII2INSERTTREE + 00000450
```

As a possible workaround, consider utilizing a sorted index rather than a sorted ranked index.

This problem has been corrected in Oracle Rdb Release 7.2.4.



## 2.2 SQL Errors Fixed

### 2.2.1 Unexpected Failures of TRUNC and ROUND Functions

Bug 6945515

In prior versions of Oracle Rdb, the ROUND and TRUNC functions could fail with an UNSNUMXPR (Unsupported numeric expression) error. The following example shows a query generated by an Oracle RDBMS tool. Here the :1 is a parameter to the dynamic query.

```
SELECT "D_CUST_CODE", "D_ORD_NO", "D_READY_DATE" FROM "DOR_ORDER" WHERE
"D_READY_DATE" >= ADD_MONTHS(LAST_DAY(TRUNC(:1))+1, -1)
ERROR at line 6:
ORA-32800: internal error [No corresponding Oracle message for Rdb error]
%SQL-F-UNSNUMXPR, Unsupported numeric expression
```

The error from Oracle Rdb indicates that it defaulted to a numeric TRUNC function.

This problem has been corrected in Oracle Rdb Release 7.2.4. SQL now does a better job of predicting the data type of input parameters based on secondary arguments to ROUND and TRUNC and also from the function context.

### 2.2.2 Unexpected Restrictions in Date/Time Subtraction

Bug 2946256

In prior versions of Oracle Rdb, the use of TIMESTAMP and DATE VMS in subtraction expressions was restricted and resulted in either DATESUBILL or DATESCANEQ errors. The following examples show these problems.

- DATE VMS could not be used in a subtraction expression without being CAST as a TIMESTAMP. In this example, DVMS is defined as DATE VMS, and TS is defined as TIMESTAMP.

```
select
    (dvms - ts) year(7) to month
from
    dt_table;
%SQL-F-UNSDATXPR, Unsupported date expression
-SQL-F-DATESUBILL, Operands of date/time subtraction are incorrect
```

- TIMESTAMP types with different fractional second precision could not appear in a subtraction expression. In this example, TS1 is defined as TIMESTAMP(1).

```
select
    (ts - ts1) day(7) to second(2)
from
    dt_table;
%SQL-F-UNSDATXPR, Unsupported date expression
-SQL-F-DATESCANEQ, Date/time expressions with different
fractional seconds precision are not comparable
```

- Expressions that should have resulted in a number of days (ORACLE LEVEL1 or ORACLE2 dialects only) were not permitted. In this example, BIRTHDAY is defined as TIMESTAMP(0) which is not

compatible with SYSDATE which is type DATE VMS.

```
set dialect 'oracle level2';
select (sysdate-bday)*(60*24) from foo;
%SQL-F-UNSDATXPR, Unsupported date expression
-SQL-F-DATESUBILL, Operands of date/time subtraction are incorrect
```

These restrictions have been lifted in Oracle Rdb Release 7.2.4. For the purposes of subtraction, DATE VMS is treated as a TIMESTAMP(2) and differences in fractional second precision are no longer considered incompatibilities between these types.

## 2.2.3 Unexpected DEFVALINC Error When Altering a Column's Domain

Bug 8242912

In prior versions of Oracle Rdb, the ALTER TABLE ... ALTER COLUMN command could fail when altering a domain for a column. This occurred because the old DEFAULT for the column was compared to the older (and in this case smaller) DEFAULT instead of the new DEFAULT inherited from the new domain.

The following example shows the unexpected error.

```
SQL> create domain char_12 char(12) default 'short str';
SQL> create domain char_20 char(20) default 'this is longer';
SQL> create table mytab (col1 char_20);
SQL> alter table mytab alter column col1 char_12;
%SQL-W-CHR_TOO_SHO, Character length of column COL1 is too short
%SQL-F-DEFVALINC, You specified a default value for COL1 which
is inconsistent with its data type
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.2.4. In this release, no error is reported although a warning will be issued because of the reduced size of the column.

## 2.2.4 Unexpected FLDNOTCRS Error When Referencing DBKEY or ROWID in Quotes

Bug 8294794

In prior versions of Oracle Rdb, a column reference to DBKEY or ROWID within quotes would result in a FLDNOTCRS error. The following example shows the error.

```
SQL> select a."ROWID" from rdb$database a;
%SQL-F-FLDNOTCRS, Column A.ROWID was not found in the tables in current scope
SQL>
```

When either SET DIALECT or SET QUOTING RULES enable quoting of columns, the references to the pseudo columns DBKEY or ROWID should be permitted.

This problem has been corrected in Oracle Rdb Release 7.2.4. SQL now allows the pseudo columns DBKEY or ROWID to be quoted. They must be in all upper case letters.

## 2.2.5 Unexpected Bugcheck When Assigning Value to FOR Cursor Variables

Bug 8429420

In prior versions of Oracle Rdb, attempts to assign a value to a FOR cursor variable would cause a bugcheck. These variables are read-only copies of the select expressions and may not be used as targets for SELECT ... INTO clause, UPDATE ... RETURNING clause and the INSERT ... RETURNING clause.

```
SQL> create module M
cont>
cont> procedure P;
cont> begin
cont> for :emp as each row of
cont>   select employee_id, last_name from employees
cont> do
cont>   select last_name
cont>     into :emp.last_name
cont>   from employees
cont>   where employee_id = :emp.employee_id;
cont> end for;
cont> end;
cont>
cont> end module;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DISK1:[TEST]SQLBUGCHK.DMP;
  SQL$721  SQL$SQL  SQL$$FLUSH_INPUT_ON_CONTROL_C
                        8000 0000000000000798 0000000000150358
  SQL$721  SQL$SQL  SQL$$FLUSH_INPUT_ON_CONTROL_C
                        8000 0000000000000798 0000000000150358
%SQL-F-BUGCHK, There has been a fatal error. Please contact your Oracle support
representative. SQL$SEMMSG - MERGE_PROC_MSG found
parsym not in signature
```

A similar bugcheck is reported for a multistatement procedure.

```
SQL> begin
cont> for :emp as each row of
cont>   select employee_id, last_name from employees
cont> do
cont>   select last_name
cont>     into :emp.last_name
cont>   from employees
cont>   where employee_id = :emp.employee_id;
cont> end for;
cont> end;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DISK1:[TEST]SQLBUGCHK.DMP;
  SQL$721  SQL$SQL  SQL$$FLUSH_INPUT_ON_CONTROL_C
                        8016 00000000000007C4 0000000000150384
  SQL$721  SQL$SQL  SQL$$FLUSH_INPUT_ON_CONTROL_C
                        8016 00000000000007C4 0000000000150384
%SQL-F-BUGCHK, There has been a fatal error. Please contact your Oracle support
representative. SQL$$BLR_MSG_FIELD_REF - NF and HV in two messages
```

This problem has been corrected in Oracle Rdb Release 7.2.4. SQL now correctly diagnoses such illegal usage, as shown in the following example.

```
SQL> begin
```

```

cont> for :emp as each row of
cont>   select employee_id, last_name from employees
cont> do
cont>   select last_name
cont>   into :emp.last_name
cont>   from employees
cont>   where employee_id = :emp.employee_id;
cont> end for;
cont> end;
%SQL-F-FORFCHVARRO, FOR statement variable "LAST_NAME" is read only

```

## 2.2.6 Unexpected Bugcheck When Executing External Routine

Bug 8231715

In prior versions of Oracle Rdb V7.2 on OpenVMS Integrity systems, a bugcheck might occur when a stored procedure or stored function attempted to execute an external routine to which the user did not have EXECUTE access.

The following example shows a stored function (CALL\_FN) in the module CALLED\_MODULE that executes a call to the external procedure LIB\$SIGNAL. This external procedure denies access to the current user and should result in the error: %RDB-E-NO\_PRIV, privilege denied by database facility.

```

SQL> show privilege on procedure LIB$SIGNAL;
Privileges on Procedure LIB$SIGNAL
  ( IDENTIFIER=[RDB,TESTER],ACCESS=NONE)
SQL> show privilege on module CALLED_MODULE;
Privileges on Module CALLED_MODULE
  ( IDENTIFIER=[RDB,TESTER],ACCESS=EXECUTE+SHOW+ALTER+DROP+REFERENCES)
SQL>
SQL> set flags 'trace';
SQL> begin
cont> declare :a integer;
cont> set :a = CALL_FN (100);
cont> end;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file
DISK1:[TESTER]RDSBUGCHK.DMP_MINI
%RDMS-I-BUGCHKDMP, generating bugcheck dump file
DISK1:[TESTER]RDSBUGCHK.DMP_MINI
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DISK1:[TESTER]RDSBUGCHK.DMP;

```

This problem has been corrected in Oracle Rdb Release 7.2.4. Oracle Rdb now correctly performs the authorization check while running the stored routine.

## 2.2.7 Unexpected Error When Importing a Database With Profile Definitions

In prior versions of Oracle Rdb, the SQL EXPORT DATABASE and IMPORT DATABASE commands did not handle correctly databases that contained objects created using the CREATE PROFILE. These objects were incorrectly imported as roles (which would fail) and assignments of profiles to users was not preserved.

The following example shows the errors that are reported by IMPORT.

```
SQL> import database
cont>     from abc_ex
cont>     filename a_ex
cont>     trace;
IMPORTing Users and Roles
Completed CDD$SYSTEM. DIO = 19, CPU = 0:00:00.01, FAULTS = 7
Completed SQLNET4RDB. DIO = 6, CPU = 0:00:00.00, FAULTS = 0
Completed RDB_EXECUTE. DIO = 14, CPU = 0:00:00.01, FAULTS = 6
Completed RDBUSER2. DIO = 13, CPU = 0:00:00.00, FAULTS = 6
%SQL-F-NOPRFRES, unable to import profile DEVELOPMENT_USER
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-SECNOTINT, database security checking is not internal
Completed DEVELOPMENT_USER. DIO = 6, CPU = 0:00:00.00, FAULTS = 10
%SQL-F-NOPRFRES, unable to import profile QUERY_USER
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-SECNOTINT, database security checking is not internal
Completed QUERY_USER. DIO = 7, CPU = 0:00:00.00, FAULTS = 7
IMPORTing Granted Users and Roles
%SQL-F-NOPGPRES, unable to import granted profile for grantee RDB_EXECUTE
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-SECNOTINT, database security checking is not internal
Completed RDB_EXECUTE. DIO = 5, CPU = 0:00:00.00, FAULTS = 3
%SQL-F-NOPGPRES, unable to import granted profile for grantee RDBUSER2
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-SECNOTINT, database security checking is not internal
Completed RDBUSER2. DIO = 5, CPU = 0:00:00.00, FAULTS = 0
.
.
.
Completed import. DIO = 1287, CPU = 0:00:00.18, FAULTS = 790
```

This problem has been corrected in Oracle Rdb Release 7.2.4. SQL EXPORT now correctly preserves the assigned profile for each user and IMPORT correctly rebuilds the assigned role list for each user.

## 2.2.8 Unexpected Bugcheck When Table Partitions Used in RESERVING Clause

Bug 5861614

In prior releases of Oracle Rdb, a bugcheck could result if two tables had partitions that reside in the same storage area. This occurred if a partition of one table was reserved for EXCLUSIVE WRITE and a partition of the other table was reserved for PROTECTED WRITE or SHARED WRITE, as in the following example.

```
SQL> set transaction
cont>     read write
cont>     reserving
cont>     EMPLOYEES partition (1) for EXCLUSIVE WRITE,
cont>     JOB_HISTORY partition (1) for PROTECTED WRITE;
SQL>
SQL> insert into EMPLOYEES
cont>     (employee_id, last_name, first_name, middle_initial)
cont>     value ('00100', 'Jones', 'Fred', 'E');
1 row inserted
SQL>
SQL> insert into JOB_HISTORY
cont>     (employee_id)
```

```
cont>      select employee_id from EMPLOYEES where employee_id = '00100';
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TESTER]RDSBUGCHK.DMP;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TESTER]SQLBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual address=
0000000000000010, PC=00000000028944C, PS=0000001B
```

When using the PARTITION clause and reserving that partition for EXCLUSIVE WRITE, the underlying storage area is reserved for EXCLUSIVE WRITE also and therefore all other partitions that reference this storage area must also be promoted to EXCLUSIVE WRITE.

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.2.9 Unexpected Failure of GRANT and REVOKE When Using Synonyms

Bug 8485872

In prior versions of Oracle Rdb, none of the commands GRANT, REVOKE, SHOW PROTECTION and SHOW PRIVILEGES supported the use of synonyms. The synonyms can be created explicitly by the CREATE SYNONYM command or implicitly by the RENAME command.

The following example shows the problems that were reported:

```
SQL> create table t1 ( a1 int );
SQL> create synonym s1 for t1;
SQL> grant update on s1 to public;
%RDB-E-NO_RECORD, access by dbkey failed because dbkey is no longer associated
with a record
-RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-NODBK, 0:314:7824 does not point to a data record
SQL> show protection on t1;
Protection on Table T1
  ( IDENTIFIER=[RDB,TESTER],ACCESS=SELECT+INSERT+UPDATE+DELETE+SHOW+CREATE+
  ALTER+DROP+DBCTRL+REFERENCES)
  ( IDENTIFIER=[*,*],ACCESS=NONE)
SQL> show protection on s1;
Protection on Table S1
%SQL-F-RELNOTDEF, Table S1 is not defined in database or schema
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.2.10 Unexpected Correlation Names Appearing as Column Headers

Bug 8648202

In prior releases of Oracle Rdb V7.2, an interactive SQL select which includes a join and also concatenates a zero length string with a column would include just the correlation name in the column header. In prior releases, this would be left blank.

The following example shows the problem.

```
SQL> select '||a.first_name, '||b.last_name
cont> from employees a, employees b
cont> limit to 1 row;
A.          B.
Terry      Ames
1 row selected
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.2.4. SQL now suppresses the column header as in prior versions.

## 2.2.11 Wrong Results When CONCAT Used in an Aggregate Expression with DISTINCT Clause

Bug 8644211

In prior versions of Oracle Rdb V7.2, queries that included CONCAT (or ||) in an aggregate expression with a DISTINCT clause might return the wrong results.

The following example shows such a query.

```
SELECT T.COL4,
       COUNT (DISTINCT ET.COLB4 || ET.COLB6) AS C2
FROM MYTABLE T LEFT OUTER JOIN MYTABLE_DET ET
  ON ET.COLB1 = T.COL1 AND ET.COLB3 = 'E'
WHERE T.COL6 BETWEEN '0' AND '5'
GROUP BY T.COL4;
```

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.2.12 UPDATE ONLY CURSOR Clause Quietly Ignored by SQL Module Language Processor and SQL Precompiler

Bug 7758451

In prior versions of Oracle Rdb, the SQL Module Language processor and SQL precompiler would quietly ignore the UPDATE ONLY CURSOR clause for a DECLARE cursor if it did not detect any update actions for that cursor. That is, no UPDATE ... WHERE CURRENT OF or DELETE ... WHERE CURRENT OF clauses referenced that cursor name.

However, in some cases, the UPDATE ONLY CURSOR clause is used to encourage stricter locking on the table during processing. In many cases, the number of lock operations can be reduced by first locking for update in preparation for the subsequent update, rather than locking first for READ and later upgrading the lock to an UPDATE lock on the row.

This problem has been corrected in Oracle Rdb Release 7.2.4. When such modules are recompiled with this release of Oracle Rdb, the clause UPDATE ONLY CURSOR will no longer be ignored and will be used by Oracle Rdb.

## 2.3 RMU Errors Fixed

### 2.3.1 Application Hangs Using Automatic AIJ Backups

Using automatic AIJ backups, the application hangs waiting for the next AIJ to become available. A dump of the database shows inaccessible AIJ files.

```
$ RMU/DUMP db
...
- 1 journal is inaccessible
  AIJ backup not possible
...
  File is inaccessible
    journal has been made inaccessible by system
    journal is not empty
...
```

The database dump may also show other symptoms of failed or stalled AIJ backups.

This problem has been corrected in Oracle Rdb Release 7.2.4.

### 2.3.2 Failed RMU MOVE or COPY Deletes Wrong Files

Bug 8400242

A failed RMU MOVE or COPY operation could delete the wrong version of a created database file at the target location. For example, this can happen if, during the RMU operation, a target device becomes full.

For this to happen, there have to be other files in the target location with the exact same filenames as the current database.

This problem has been corrected in Oracle Rdb Release 7.2.4.

### 2.3.3 RMU Extract Now Extracts Comments for Constraints

Bug 8238128

In prior versions of Oracle Rdb, the RMU Extract command did not extract the comments applied to a constraint by the COMMENT ON CONSTRAINT statement. The following example shows that an additional COMMENT ON statement is output after the CREATE TABLE statement.

```
set verify;
set language ENGLISH;
set default date format 'SQL92';
set quoting rules 'SQL92';
set date format DATE 001, TIME 001;
attach 'filename ABC';
create table TESTCON (
  T1K
    SMALLINT
    constraint T1K_PRIMARY
```



```

        primary key
        initially deferred deferrable
    comment is
        'Primary',
T2
    SMALLINT);

comment on constraint T1K_PRIMARY is
    'Primary key';

commit work;

```

This problem has been corrected in Oracle Rdb Release 7.2.4. RMU Extract now extracts the constraint comments.

## 2.3.4 Bugcheck in RMU /COLLECT OPTIMIZER\_STATISTICS /STATISTICS=WORKLOAD

Bug 8433938

The following command bugchecks:

```

$ RMU/COLLECT OPTIMIZER_STATISTICS /STATISTICS=WORKLOAD databasename
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=0000000000000000, PC=000000008002CA10, PS=0000001B
%RMU-F-FATALOSI, Fatal error from the Operating System Interface.
%RMU-I-BUGCHKDMP, generating bugcheck dump file DISK:[DIR]RMUBUGCHK.DMP;
%RMU-F-FTL_COL_STAT, Fatal error for COLLECT OPTIMIZER_STATISTICS
operation at 14-APR-2009 17:59:08.77

```

This problem is caused by the code accessing data beyond the last entry of a collision error table.

The workaround would be to delete the offending workload column groups.

```

$ RMU/Delete Optimizer_Statistics/Column=(COMPANY_ID, LEGAL_ENTITY_ID, -
GROUP_ID)/Table= CRIT_GROUP database
$ RMU/Delete Optimizer_Statistics/Column=(COMPANY_ID, LEGAL_ENTITY_ID, -
CPTY_MEDIA_IND, GROUP_ID)/Table= CRIT_GROUP database
$ RMU/Delete Optimizer_Statistics/Column=(COMPANY_ID, LEGAL_ENTITY_ID, -
CPTY_MEDIA_IND, GROUP_ID, RECORD_STATE)/Table= CRIT_GROUP database

```

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.3.5 Increased Cardinality Limit in RMU /COLLECT OPTIMIZER\_STATISTICS /STATISTICS=WORKLOAD

Bug 8460493

In prior versions of Oracle Rdb, Optimizer Workload Collection in RMU ignored any table with cardinality higher than 249.95 million. See the following example.

```

$RMU/COLLECT OPTIMIZER_STATISTICS /STATISTICS=WORKLOAD TESTDB/LOG
Start loading tables... at 23-APR-2009 09:55:49.41
Done loading tables.... at 23-APR-2009 09:55:49.42

```

## Oracle® Rdb for OpenVMS

```
Start collecting workload stats... at 23-APR-2009 09:55:53.47
Maximum memory required (bytes) = 1166796
Table : XTEST exceeds 249.95 million rows; default statistics values used.
Done collecting workload stats... at 23-APR-2009 09:55:53.51
Start calculating stats... at 23-APR-2009 09:55:53.51
Done calculating stats... at 23-APR-2009 09:55:53.51
Start writing stats... at 23-APR-2009 09:55:55.38
```

```
-----
Optimizer Statistics collected for table : FOO
  Workload Column group   : COL1
  Duplicity factor        : 1.0000000
  Null factor             : 0.0000000
Done writing stats... at dd-mmm-yyyy hh:mm:ss.01
```

Notice that the duplicity and null factors do not get computed at all.

The restriction on maximum cardinality at 249.95 million has been increased to approximately  $2^{52}$ .

This problem has been corrected in Oracle Rdb Release 7.2.4.

### 2.3.6 Incorrect LOCK\_TIMEOUT With RMU /BACKUP /ONLINE /LOCK\_TIMEOUT /LOG

Bug 4047148

In prior releases of Oracle Rdb, if both RMU/BACKUP/AFTER\_JOURNAL and RMU/BACKUP specify LOCK\_TIMEOUT then it appears that subsequent timeouts are increased. For example, if /LOCK\_TIMEOUT=60 is specified on an after-image journal backup, it stalls for approximately this time. If the database backup also has a 60 second wait time, then this appears to double and it waits for 2 minutes.

Session 1:

```
SQL> attach 'file mf_personnel';
SQL> update employees set sex='M'
```

Session 2:

```
$ set noon
$ show time
  3-DEC-2008 17:17:09
$ rmu/back/after/lock=60 mf_personnel mf_per.ajj
%RMU-W-DATACMIT, unjournalized changes made; database may not be recoverable
%RMU-F-TIMEOUT, timeout on quiet
-COSI-W-CANCEL, operation canceled
%RMU-F-FTL_BCK, Fatal error for BACKUP operation at  3-DEC-2008 17:18:11.91
$ show time
  3-DEC-2008 17:18:12
$ rmu/back/log/online/lock=60 mf_personnel mf_per.rbf
%RMU-I-QUIETPT, waiting for database quiet point at  3-DEC-2008 17:18:22.27
%RMU-F-TIMEOUT, timeout on quiet
-COSI-W-CANCEL, operation canceled
%RMU-F-FTL_BCK, Fatal error for BACKUP operation at  3-DEC-2008 17:20:22.15
$ show time
  3-DEC-2008 17:20:22
```

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.3.7 RMU /RESTORE /ONLINE /JUST\_CORRUPT Bugcheck

Bug 8568870

In some cases, the "RMU /RESTORE /ONLINE /JUST\_CORRUPT" command can fail with an internal consistency failure bugcheck "footprint" similar to:

```
Exception occurred at RMU72\RMUCLIS$RESTORE + 000038A4
COSI-F-BUGCHECK, internal consistency failure
Called from RMU72\RMU_DISPATCH + 00000F44
Called from RMU72\RMU_STARTUP + 000004CC
Called from RMU72\RMU$MAIN + 00000034
```

This problem can be triggered by having deleted storage areas as in the following example:

```
$ SQL$
CREATE DATABASE FILENAME TESTDB
  NUMBER OF CLUSTER NODES IS 1 NUMBER OF USERS IS 10
  CREATE STORAGE AREA A0 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A1 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A2 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A3 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A4 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A5 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A6 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A7 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A8 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8
  CREATE STORAGE AREA A9 ALLOCATION IS 8 SNAPSHOT ALLOCATION IS 8;
ALTER DATABASE FILENAME TESTDB
  DROP STORAGE AREA A3
  DROP STORAGE AREA A4
  DROP STORAGE AREA A5;
CREATE TABLE T1 (C1 INTEGER, C2 CHAR(5));
CREATE STORAGE MAP M1 FOR T1 STORE IN A1;
CREATE TABLE T8 (C1 INTEGER, C2 CHAR(5));
CREATE STORAGE MAP M8 FOR T8 STORE IN A8;
INSERT INTO T1 VALUES (11, 'ONE') RETURNING DBKEY;
INSERT INTO T8 VALUES (81, 'ONE') RETURNING DBKEY;
COMMIT;
EXIT;
$ RMU/BACKUP/NOLOG TESTDB TESTDB
$ RMU/ALTER TESTDB ! Invalidate checksums
  AREA 3 PAGE 5
  DEPOSIT CHECKSUM = 101010101
  COMMIT
  AREA 10 PAGE 5
  DEPOSIT CHECKSUM = 101010101
  COMMIT
  EXIT
$ RMU/OPEN/WAIT TESTDB
$ SQL$ ! Detect checksum errors and add to CPT
  ATTACH 'FILE TESTDB';
  SELECT * FROM T1;
  SELECT * FROM T8;
  EXIT;
$ RMU/SHOW CORRUPT TESTDB
$ RMU/RESTORE/ONLINE/JUST_CORRUPT/NOLOG TESTDB
$ RMU/CLOSE/WAIT TESTDB
```

As a workaround, omit the "/ONLINE" qualifier; perform the /JUST\_CORRUPT restore operation offline.

This problem has been corrected in Oracle Rdb Release 7.2.4. The online /JUST\_CORRUPT restore now correctly detects and ignores invalid FILID entries.

## 2.3.8 RMU /LOAD /PARALLEL COSI-F-READERR Without VMS BYPASS Privilege

Bug 8468385

The RMU/LOAD/PARALLEL documentation correctly states that the VMS BYPASS process privilege is not required for RMU/LOAD/PARALLEL as long as the Oracle Rdb database Access Control List permits the user to execute the RMU/LOAD command. However, a %COSI-F-READERR error resulted from a code problem that caused the VMS BYPASS process privilege to be required to access a VMS mailbox device used for communicating with the executor processes that are created by RMU/LOAD/PARALLEL. This problem has been fixed and the code now conforms to the documentation that the BYPASS process privilege is not required for RMU/LOAD/PARALLEL.

The following example shows the problem. If the user is not granted the VMS BYPASS privilege, a %COSI-F-READERR is returned when each executor is created by the RMU/LOAD/PARALLEL command when loading the TEST\_LOAD table from the TEST\_LOAD.UNL unload file in the TESTDB.RDB database. No data is loaded and the parallel load terminates with a fatal error status.

```
RYEROX>show process/privileges
Authorized privileges:
  CMKRNL      NETMBX      OPER          PRMGBL      SYSGBL      SYSLCK
  SYSPRV      TMPMBX      WORLD

Process privileges:
  NETMBX          may create network device
  TMPMBX          may create temporary mailbox

$rmu/show privileges testdb.rdb
  ( IDENTIFIER=[ *, * ]ACCESS=READ+WRITE+CONTROL+RMU$ALL)

$rmu/unload testdb.rdb test_load test_load.unl
$rmu/load/debug=trace/commit=10000 -
/parallel=(exec=2,buff=64) testdb.rdb test_load test_load.unl

* Using EXECUTOR_IMAGE:"SYS$COMMON:[SYSLIB]RDMRLE72.EXE"

%COSI-F-READERR, read error

%COSI-F-READERR, read error

%RMU-I-DATRECREAD, 0 data records read from input file.
%RMU-I-DATRECSTO, 0 data records stored.
%RMU-F-FTL_LOAD, Fatal error for LOAD operation at 27-APR-2009 09:36:20.77
```

The only way to avoid this problem in earlier versions of Oracle Rdb which have this problem is to grant the VMS BYPASS process privilege to the account from which the parallel load is being run or to do a non-parallel load of the database table.

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.4 LogMiner Errors Fixed

### 2.4.1 Continuous LogMiner Startup Serialized With AIJ Backups

Bug 7634512

Although it is documented that it is not permitted to perform AIJ backup operations when using the Continuous LogMiner feature until the Continuous LogMiner has fully transitioned to the live AIJ files, doing so could result in unintended data loss in the output stream from the LogMiner.

The impact of this problem has been reduced in Oracle Rdb Release 7.2.4. The Continuous LogMiner now utilizes a global "AIJ Backup" lock to serialize the startup of the Continuous LogMiner with AIJ backup operations.

### 2.4.2 RMU /UNLOAD /AFTER\_JOURNAL Allows /RESTART Without /CONTINUOUS

Previously, the /RESTART qualifier required the /CONTINUOUS qualifier for restarting the LogMiner. This restriction has been relaxed. The /RESTART qualifier is now allowed with or without the CONTINUOUS qualifier.

In addition, when the /RESTART qualifier is specified, it is no longer enforced that the first AIJ backup file supplied be from after a quiet-point AIJ backup. This implies that it is possible that the transaction returned with the specified AERCP (AIJ Extract Restart Control Point) could be an incomplete transaction (where not every modification made by the transaction is returned in the output data stream). Applications that use /RESTART are expected to understand this behavior and to ignore the transaction with the restarted AERCP.

This problem has been corrected in Oracle Rdb Release 7.2.4.

## 2.5 Row Cache Errors Fixed

### 2.5.1 Incremental Backup With Row Cache

Bug 8363084

Previously, it was possible for incremental database backups to not correctly save all modified database rows since the last full backup when the row cache feature was in use. When modified rows were copied from cache back to the database, the "MAX\_SNAP\_TSN" field on the database page was not correctly maintained. This field is used by the incremental database backup feature as an indication of when rows on the page may have been modified and if it is out of date, the incremental backup may not consider the page content as a candidate to be saved.

As a possible workaround for this problem, perform full database backups rather than incremental database backups when using the row cache feature.

This problem has been corrected in Oracle Rdb Release 7.2.4. The "MAX\_SNAP\_TSN" field on the database page is now maintained correctly when modified rows are copied from cache back to the database.

## 2.6 Hot Standby Errors Fixed

### 2.6.1 Hot Standby LRS Database Prefetch Count Limit

Previously, the LRS process would set its APF (asynchronous prefetch) depth to half of its buffer count. For large numbers of buffers, this depth could result in a significant number of outstanding database prefetch IO requests. In some cases, this could result in possible quota exhaustion or storage controller saturation.

This problem has been addressed by providing additional control of the LRS process APF depth along with a lower default limit.

The logical name (must be defined system-wide prior to the startup of an LRS process)

RDM\$BIND\_LRS\_MAX\_APF\_DEPTH can be used to limit the maximum number of default LRS APF IO depth. If not specified, the default value is 500. The minimum value is 2 and the maximum value is 524288. If the database specifies a higher APF depth value, that value will be utilized.

---

# **Chapter 3**

## **Software Errors Fixed in Oracle Rdb Release 7.2.3.5**

This chapter describes software errors that are fixed by Oracle Rdb Release 7.2.3.5.



## 3.1 Software Errors Fixed That Apply to All Interfaces

### 3.1.1 Patch Required When Using VMS V8.3 or V8.2-1 on Itanium

During qualification testing of Oracle Rdb, a possible problem related to OpenVMS image activation on some Itanium systems was discovered. This problem could, in some cases, result in unexpected messages similar to the following:

```
-RDB-E-UNAVAILABLE, Rdb/Dispatch is not available on your system
-RDB-I-TEXT, %LIB-E-ACTIMAGE, error activating image
DUA0:[SYS1.SYSCOMMON.][SYSLIB]RDB$COSIP.EXE,
-SYSTEM-F-NOPRIV, insufficient privilege or object protection violation
```

To avoid this problem, Oracle recommends that customers install one of the following architecture-specific patch kits (or subsequent replacement if superseded) for Itanium systems:

- VMS821I\_SYS-V0400 (I64)
- VMS83I\_SYS-V0200 (I64)

### 3.1.2 Wrong Result From Outer Zigzag Match

Bug 6310672

The match join query with outer zigzag skip returns the wrong result due to the filtered predicate referencing the middle segment of the index, as seen in the example below.

```
SELECT * FROM T1;
      COL1      COL2      COL3      RID
      -----      -----      -----      ---
           1           9278           10605      OR1
           46           9270           46002      OR2
           46           9270           46997      OR3
           46           9278           46002      OR4
4 rows selected

SELECT * FROM T2;
      COL1      COL2      COL3      RID
      -----      -----      -----      ---
           46           46002           146002      IR1
           46           46011           146002      IR2
2 rows selected

set flags 'strategy, detail(2)';
SELECT T1.COL1, T1.COL2, T1.COL3, T2.COL3, T1.RID, T2.RID
FROM T1 t1, T2 t2
WHERE
  T1.COL1 = T2.COL1 AND
  T1.COL3 = T2.COL2 AND
  T1.COL2 = 9278 AND
  T2.COL3 = 146002;
```

```

Tables:
  0 = T1
  1 = T2
Conjunct: (0.COL1 = 1.COL1) AND (0.COL3 = 1.COL2)
Match Q1
Outer loop      (zig-zag)
Match_Keys:0.COL1, 0.COL3
Index_Keys:COL1, COL2, COL3, RID
  Conjunct: 0.COL2 = 9278
  Index only retrieval of relation 0:T1
    Index name  T1_NDX [0:0]
Inner loop      (zig-zag)
Match_Keys:0.COL1, 0.COL3
Index_Keys:COL1, COL2, COL3, RID
  Conjunct: 1.COL3 = 146002
  Index only retrieval of relation 1:T2
    Index name  T2_NDX [0:0]
0 rows selected

```

The query works if 'ZIGZAG\_OUTER' is disabled, as in the following example.

```

set flags 'NOZIGZAG_OUTER';
...execute the above query....
Tables:
  0 = T1
  1 = T2
Conjunct: (0.COL1 = 1.COL1) AND (0.COL3 = 1.COL2)
Match Q1
Outer loop
Match_Keys:0.COL1, 0.COL3
  Conjunct: 0.COL2 = 9278
  Index only retrieval of relation 0:T1
    Index name  T1_NDX [0:0]
Inner loop      (zig-zag)
Match_Keys:0.COL1, 0.COL3
Index_Keys:COL1, COL2, COL3, RID
  Conjunct: 1.COL3 = 146002
  Index only retrieval of relation 1:T2
    Index name  T2_NDX [0:0]
    T1.COL1      T1.COL2      T1.COL3      T2.COL3      T1.RID      T2.RID
    46           9278         46002       146002      OR4         IR1
1 row selected

```

This problem occurs when two tables are joined with zigzag match strategy using the index that defines two columns as leading and trailing segments with the middle segment as a filter predicate.

In the above example query, the first and third columns of the outer index T1\_NDX are joined by zigzag match to the first and second columns of the inner index T2\_NDX.

```

create index T1_NDX on table T1 (COL1, COL2, COL3);
create index T2_NDX on table T2 (COL1, COL2, COL3);

```

The middle segment column of the outer index T1\_NDX is referenced by the filter key predicate 'T1.COL2 = 9278'.

As of today, Oracle Rdb disallows the middle and trailing segment columns to be included in the generation of ikeys.

This problem has been corrected in Oracle Rdb Release 7.2.3.5.

### 3.1.3 Cannot Disable After–Image Journaling when Fast–Commit is Enabled

Bug 8286207

In certain circumstances, attempts to disable after–image journaling (AIJ) in Oracle Rdb would fail with the following error message when the Fast–Commit feature was enabled:

```
$ RMU/SET AFTER/DISABLE mf_personnel
%RMU-F-MUSTRECDB, database must be closed or recovered
```

When Fast–Commit is enabled, certain checks must be executed to ensure that committed transactions are synchronized between the database and AIJ file before journaling can be disabled.

If the database was closed cleanly (no users remaining), then after–image journaling could be disabled. However, if the database was closed abnormally while users were still attached, a check was made to determine if any of those users had an active checkpoint. If so, journaling could not be disabled.

It was recently discovered that this second level check was not being performed correctly and would report that an abnormally–terminated user had an active checkpoint when they, in fact, did not.

This problem has been corrected in Oracle Rdb Release 7.2.3.5.

### 3.1.4 Alignment Faults Related to Remote Database Access

Bug 8355675

When accessing a remote database, both the client and server processes may experience alignment faults due to access of an internal unaligned network buffer.

This problem has been corrected in Oracle Rdb Release 7.2.3.5.

### 3.1.5 RDMS–E–NOSOL\_FOUND Error From Full OJ Query

In prior releases of Oracle Rdb, it was possible to have queries fail with a NOSOL\_FOUND as shown in the following example.

```
SELECT V.VOYAGE_NUM, S.SHIP_NAME
      FROM
          VOYAGE V FULL OUTER JOIN SHIP S ON
              V.VOYAGE_NUM = S.SHIP_NAME
      WHERE
          V.VOYAGE_NUM = 4903 AND
          S.SHIP_NAME = 'AMIGO'
      ;
~S: Full OJ query with cross strategy was not possible
%RDB-E-IMP_EXC, facility-specific limit exceeded
-RDMS-E-NOSOL_FOUND, No possible solution has been found by Rdb optimizer
```

However, such queries may work if one of the predicates is commented out, as in the following example.

```
set flags 'strategy,detail(2)';
SELECT V.VOYAGE_NUM, S.SHIP_NAME
      FROM
            VOYAGE V FULL OUTER JOIN SHIP S ON
            V.VOYAGE_NUM = S.SHIP_NAME
      WHERE
            ! works if one of the predicates is commented out
            !V.VOYAGE_NUM = 0 AND
            S.SHIP_NAME = 'AMIGO'
      ;
```

Tables:

0 = VOYAGE

1 = SHIP

Conjunct: 1.SHIP\_NAME = 'AMIGO'

Match (Full Outer Join) Q2

Outer loop

Match\_Key:CAST (0.VOYAGE\_NUM AS CHAR(15))

Sort: CAST (0.VOYAGE\_NUM AS CHAR(15))(a)

Index only retrieval of relation 0:VOYAGE

Index name VOYAGE\_NUM\_INDEX [0:0]

Inner loop

Match\_Key:CAST (0.VOYAGE\_NUM AS CHAR(15))

Conjunct: 1.SHIP\_NAME = 'AMIGO'

Index only retrieval of relation 1:SHIP

Index name SHIP\_NAME\_INDEX [1:1]

Keys: 1.SHIP\_NAME = 'AMIGO'

V.VOYAGE\_NUM S.SHIP\_NAME

NULL AMIGO

1 row selected

```
SELECT V.VOYAGE_NUM, S.SHIP_NAME
      FROM
            VOYAGE V FULL OUTER JOIN SHIP S ON
            V.VOYAGE_NUM = S.SHIP_NAME
      WHERE
            ! works if one of the predicates is commented out
            V.VOYAGE_NUM = 4903
            !and S.SHIP_NAME = 'AMIGO'
      ;
```

Tables:

0 = VOYAGE

1 = SHIP

Conjunct: 0.VOYAGE\_NUM = 4903

Match (Full Outer Join) Inner TTBL Q2

Outer loop

Match\_Key:CAST (0.VOYAGE\_NUM AS CHAR(15))

Index only retrieval of relation 1:SHIP

Index name SHIP\_NAME\_INDEX [0:0]

Inner loop

Match\_Key:CAST (0.VOYAGE\_NUM AS CHAR(15))

Temporary relation

Sort: CAST (0.VOYAGE\_NUM AS CHAR(15))(a)

Conjunct: 0.VOYAGE\_NUM = 4903

Index only retrieval of relation 0:VOYAGE

Index name VOYAGE\_NUM\_INDEX [1:1]

Keys: 0.VOYAGE\_NUM = 4903

V.VOYAGE\_NUM S.SHIP\_NAME

4903 NULL

1 row selected

## Oracle® Rdb for OpenVMS

If all the predicates are removed, the query of full outer join still compiles and executes using the match strategy.

```
SELECT V.VOYAGE_NUM, S.SHIP_nAME
      FROM
          VOYAGE V FULL OUTER JOIN SHIP S ON
              V.VOYAGE_NUM = S.SHIP_NAME
      ;
```

Tables:

0 = VOYAGE

1 = SHIP

Match (Full Outer Join) Q2

Outer loop

Match\_Key:CAST (0.VOYAGE\_NUM AS CHAR(15))

Sort: CAST (0.VOYAGE\_NUM AS CHAR(15))(a)

Index only retrieval of relation 0:VOYAGE

Index name VOYAGE\_NUM\_INDEX [0:0]

Inner loop

Match\_Key:CAST (0.VOYAGE\_NUM AS CHAR(15))

Index only retrieval of relation 1:SHIP

Index name SHIP\_NAME\_INDEX [0:0]

V.VOYAGE_NUM	S.SHIP_NAME
4902	NULL
4903	NULL
4904	NULL
4905	NULL
4906	NULL
4907	NULL
4908	NULL
4909	NULL
4910	NULL
4911	NULL
NULL	AMIGO
NULL	ANNACONDA
NULL	BLUEGILL
NULL	CLEARVIEW
NULL	DAFFODIL
NULL	LADY JANE
NULL	LISA T.
NULL	MATILDA
NULL	MIDDIE
NULL	OCEAN STAR
NULL	PATHFINDER
NULL	PUDDLE JUMPER
NULL	SANDRA C.
NULL	SILENT PASSAGE
NULL	TRANS DELTA
NULL	ULTRA
NULL	VOLORIA II
NULL	WALTER B. FINN

28 rows selected

When a filter predicate is added to both outer and inner legs, the query gets an error message saying that the cross strategy is not a possible solution for a full outer join. This is an erroneous message as the query should not be using cross strategy. The match strategy should be allowed in this case.

This problem has been corrected in Oracle Rdb Release 7.2.3.5.

## 3.2 SQL Errors Fixed

### 3.2.1 Wrong Result from CONCAT Query Under ORACLE Dialects

Bug 7595255

In prior versions of Oracle Rdb V7.2, when the dialect was set to ORACLE LEVEL1 or ORACLE LEVEL2, the CONCAT function or concatenation (||) operator would not return the correct result for the case where the concatenation resulted in a zero length string. Oracle dialect requires this result to be interpreted as NULL.

This problem has been corrected in Oracle Rdb Release 7.2.3.5. Oracle Rdb now correctly processes such results.

### 3.2.2 Unexpected Bugcheck During RENAME Operation

Bug 8245741

In prior releases of Oracle Rdb, it was possible for the RENAME statement (or RENAME TO clause of an ALTER statement) to fail with a bugcheck. Often repeating the operation would succeed. The bugcheck would have a footprint similar to the following.

- SYSTEM-F-ACCVIO, access violation, virtual address=000000001F60000
- Exception occurred at symbol not found
- Called from RDMSHRP72\RDMS\$\$FIND\_SYNONYM + 000001C0
- Called from RDMSHRP72\RDMS\$\$CHANGE\_INDEX + 000009C0
- Called from RDMSHRP72\RDMS\$\$RELEASE\_DDL\_VM\_HNDLR + 00003270

This problem has been corrected in Oracle Rdb Release 7.2.3.5. This problem occurred on rare occasions when the DDL buffer was followed by a non-readable page. Oracle Rdb has been modified to avoid this condition.

### 3.2.3 Unexpected DDL\_CONFLICT Error Generated by ALTER TABLE ... ENABLE CONSTRAINT

Bug 8301379

In prior versions of Oracle Rdb, attempts to enable constraint such as shown in the example below, failed with an unexpected DDL\_CONFLICT error. In fact, ENABLE CONSTRAINT, ENABLE PRIMARY KEY and ENABLE UNIQUE should all be supported when the table is reserved for SHARED DATA DEFINITION.

```
SQL> set transaction read write
cont>     reserving employees for shared data definition;
SQL> alter table employees enable primary key;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDB-E-LOCK_CONFLICT, request failed due to locked resource
-RDMS-E-DDL_CONFLICT, relation EMPLOYEES specified reserved for data definition
SQL> rollback;
```

## Oracle® Rdb for OpenVMS

This problem has been corrected in Oracle Rdb Release 7.2.3.5. The query executed by ENABLE CONSTRAINT was not identifying itself as only reading the data and so was rejected.

## 3.3 RMU Errors Fixed

### 3.3.1 Unexpected RMU/REPAIR Error if Index RDB\$REL\_REL\_ID\_NDX is Sorted Ranked

In prior releases of Oracle Rdb, the RMU/REPAIR command assumed that the database system index RDB\$REL\_REL\_ID\_NDX on the system table RDB\$RELATIONS would only be of type SORTED but it can also be of type SORTED RANKED. This caused RMU/REPAIR and RMU/REPAIR/SPAM to abort with an error if the system index RDB\$REL\_REL\_ID\_NDX was defined as SORTED RANKED.

RMU/REPAIR accesses the system index RDB\$REL\_REL\_ID\_NDX defined for the system table RDB\$RELATIONS at the start of RMU/REPAIR to gather information which is later used as a basis for the Oracle Rdb database repair of SPAM pages.

The following examples show the error that was output by RMU/REPAIR if Oracle Rdb database SPAM pages needed to be reconstructed.

```
$ RMU/REPAIR RDB_DATABASE
%RMU-I-FULBACREQ, A full backup of this database should be performed after
RMU REPAIR
%RMU-F-NORELIDACC, The RDB$REL_REL_ID_NDX index could not be processed
%RMU-F-FTL_REP, Fatal error for REPAIR operation at 31-JAN-2009 01:02:51.50
$ RMU/REPAIR/SPAM RDB_DATABASE
%RMU-I-FULBACREQ, A full backup of this database should be performed after
RMU REPAIR
%RMU-F-NORELIDACC, The RDB$REL_REL_ID_NDX index could not be processed
%RMU-F-FTL_REP, Fatal error for REPAIR operation at 31-JAN-2009 01:02:51.54
```

There is no way to avoid this problem except to either specify RMU/REPAIR/NOSPAMS or to recreate the database using the SYSTEM INDEX clause to specify TYPE IS SORTED.

This problem has been corrected in Oracle Rdb Release 7.2.3.5. RMU/REPAIR has now been modified to access the RDB\$REL\_REL\_ID\_NDX as either a Sorted or Sorted Ranked index.



## 3.4 LogMiner Errors Fixed

### 3.4.1 Continuous LogMiner Shutdown Via SYS\$FORCEX Causes DBR

Previously, if a process running the Continuous LogMiner was terminated via the SYS\$FORCEX system service, a database recovery process would be initiated unexpectedly.

This problem has been corrected in Oracle Rdb Release 7.2.3.5. The RMU /UNLOAD /AFTER\_JOURNAL /CONTINUOUS command now establishes an exit handler intended to attempt to disconnect from the database in an orderly fashion to avoid a database recovery process.

## 3.5 RMU Show Statistics Errors Fixed

### 3.5.1 RMU/SHOW STATISTICS Row Cache Overview Invalid Percentage

Bug 8366854

The RMU/SHOW STATISTICS "Row Cache Overview" was not correctly processing all quadword statistics values. This could lead to integer truncation which would lead to invalid results when calculating percentages prior to display.

This problem has been corrected in Oracle Rdb Release 7.2.3.5.

---

# **Chapter 4**

## **Software Errors Fixed in Oracle Rdb Release 7.2.3.4**

This chapter describes software errors that are fixed by Oracle Rdb Release 7.2.3.4.

## 4.1 Software Errors Fixed That Apply to All Interfaces

### 4.1.1 Single Byte String Comparison With Some Relational Operators Returns Different Result on Itanium

Bug 8225666

Using the relational operators for greater-than, less-than, greater-than-or-equal or less-than-or-equal on single-byte strings on I64 systems would incorrectly perform signed comparisons of the bytes rather than the expected unsigned comparisons. The following example demonstrates the incorrect results on Itanium:

```
$ SQL$
SQL> CREATE DATA FILE FOO;
SQL> CREATE TABLE T1 (F1 CHAR(1));
SQL> INSERT INTO T1 VALUES ('0');
1 row inserted
SQL> SELECT * FROM T1 WHERE F1 > X'FF';
  F1
  0
1 row selected
SQL> SELECT * FROM T1 WHERE F1 < X'FF';
0 rows selected
SQL> SELECT * FROM T1 WHERE F1 > X'00';
  F1
  0
1 row selected
SQL> SELECT * FROM T1 WHERE F1 < X'00';
0 rows selected
SQL> COMMIT;
```

This problem has been corrected in Oracle Rdb Release 7.2.3.4. Correct unsigned relational comparisons are now being performed on single byte strings.

### 4.1.2 Wrong Results Queries When Using DISTINCT

Bug 7829199

In prior versions of Oracle Rdb V7.2 on Integrity systems, it was possible to see wrong results when a query used DISTINCT and other sort operations such as GROUP BY, ORDER BY, DISTINCT, or UNION DISTINCT clauses. This happened when these sort operations were active at the same time such as in nested cursors or nested FOR loops.

A workaround for this problem is to define the logical name RDMS\$BIND\_MAX\_QSORT\_COUNT to 0. This disables the use of the Quicksort routine and so avoids this problem.

This problem has been corrected in Oracle Rdb Release 7.2.3.4. The Quicksort routine for Oracle Rdb on Integrity systems has been corrected to reset the compare routine prior to use.

This problem does not affect OpenVMS Alpha systems.

### 4.1.3 Potential System Crash Using LARGE MEMORY IS ENABLED or SHARED MEMORY IS SYSTEM on Itanium

When using the SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED feature on some Itanium systems, it was possible for Oracle Rdb to cause an OpenVMS system crash with a bugcheck type of "MACHINECHK, Machine check while in kernel mode". The MACHINECHK can be triggered by an invalid PFN (page frame number) value being loaded into a PTE (page table entry).

This problem only effects those Itanium systems with physical memory accessible to the system of PFNs greater than 268,435,455 (Hex 0FFFFFFF). The SDA command CLUE CONFIG can be used to display the system memory configuration.

The workaround to this problem, for such systems with PFNs greater than 268,435,455, is to discontinue use of the LARGE MEMORY IS ENABLED and SHARED MEMORY IS SYSTEM features if enabled, by using one or both, of the SQL commands:

- ALTER DATABASE ... GLOBAL BUFFERS ...LARGE MEMORY IS DISABLED
- ALTER DATABASE ... SHARED MEMORY IS PROCESS RESIDENT

This problem has been corrected in Oracle Rdb Release 7.2.3.4. The Oracle Rdb VLM feature correctly processes 40-bit PFN values.

### 4.1.4 Bugcheck from Nested GROUP BY Query with OR Predicate

Bug 7669128

The following nested GROUP BY query with OR predicate bugchecks.

```
SELECT * FROM
  (SELECT
    S.DELIV_MONTH, S.DESTINATION_CODE, SUM(S.QUANTITY)
  FROM
    (SELECT
      S1.DELIV_MONTH, S1.DESTINATION_CODE, SUM(S1.QUANTITY)
    FROM SALES S1
    GROUP BY S1.DELIV_MONTH, S1.DESTINATION_CODE) AS
    S (DELIV_MONTH, DESTINATION_CODE, QUANTITY),
    FUEL_VALUE F
  WHERE S.DELIV_MONTH = F.DELIV_MONTH AND
        S.DESTINATION_CODE = F.DESTINATION_CODE
    GROUP BY S.DELIV_MONTH, S.DESTINATION_CODE) AS
  V1 (DELIV_MONTH, DESTINATION_CODE, QUANTITY)
WHERE
  DESTINATION_CODE = 401 OR DESTINATION_CODE = 506
GROUP BY DELIV_MONTH, DESTINATION_CODE;
%DEBUG-I-DYNMODSET, setting module RDMS$PREEXE
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=0000000000000000, PC=00000000007F2954, PS=0000001B
%RDB-E-BAD_REQ_HANDLE, invalid request handle
```

The query works if the OR predicate is omitted, as in the following example.

```

SELECT * FROM
  (SELECT
    S.DELIV_MONTH, S.DESTINATION_CODE, SUM(S.QUANTITY)
  FROM
    (SELECT
      S1.DELIV_MONTH, S1.DESTINATION_CODE, SUM(S1.QUANTITY)
    FROM SALES S1
    GROUP BY S1.DELIV_MONTH, S1.DESTINATION_CODE) AS
    S (DELIV_MONTH, DESTINATION_CODE, QUANTITY),
    FUEL_VALUE F
  WHERE S.DELIV_MONTH = F.DELIV_MONTH AND
        S.DESTINATION_CODE = F.DESTINATION_CODE
    GROUP BY S.DELIV_MONTH, S.DESTINATION_CODE) AS
  V1 (DELIV_MONTH, DESTINATION_CODE, QUANTITY)
WHERE
  DESTINATION_CODE = 401
GROUP BY DELIV_MONTH, DESTINATION_CODE;
Tables:
  0 = S320_631_AL_EX_AB_SALES
  1 = S320_631_EX_AB_FUEL_VALUE
Reduce: 0.DELIV_MONTH, 0.DESTINATION_CODE
Merge of 1 entries  Q2
  Merge block entry 1  Q3
Aggregate: 0:SUM (Q6-CAGG"=AGG_TOTAL"0.QUANTITY) Q4
Sort: 0.DELIV_MONTH(a), 0.DESTINATION_CODE(a)
Cross block of 2 entries  Q4
  Cross block entry 1
    Merge of 1 entries  Q4
      Merge block entry 1  Q5
        Aggregate: 1:SUM (0.QUANTITY) Q6
        Sort: 0.DELIV_MONTH(a), 0.DESTINATION_CODE(a)
        Conjoinct: 0.DESTINATION_CODE = 401
        Get      Retrieval sequentially of relation 0:S320_631_AL_EX_AB_SALES
    Cross block entry 2
      Conjoinct: (0.DELIV_MONTH = 1.DELIV_MONTH) AND (0.DESTINATION_CODE =
        1.DESTINATION_CODE)
      Get      Retrieval sequentially of relation 1:S320_631_EX_AB_FUEL_VALUE
Deliv   Dest
Month   Code
-----
  0809   401
1 row selected

```

The problem occurs when the OR predicate is applied to the nested GROUP BY query.

This problem has been corrected in Oracle Rdb Release 7.2.3.4.

## 4.1.5 Sort Query Returns Incorrect Data

Bug 7645453

In prior releases of Oracle Rdb, when using an ORDER BY in a query that selects a column which is larger than 32,767 octets, incorrect data values could be returned for some columns.

The problem only occurred on OpenVMS Alpha.

This problem has been corrected in Oracle Rdb Release 7.2.3.4.

## 4.1.6 Wrong Result from Query with UNION and OR Predicates

Bug 7649851

In prior releases of Oracle Rdb, it was possible to get the wrong result from a cross strategy join query with UNION legs at the outer cross block using complex conjuncts involving OR predicates.

The query works correctly if one of the predicates from the OR predicate is commented out, as can be seen in the following two examples:

```
and (df.cpty_id like :cpty_id or :cpty_id = ' ')
```

changed to:

```
and (df.cpty_id like :cpty_id ) --or :cpty_id = ' ')
```

The following simplified query reproduces the problem (should return 198 rows, but returns 0 rows instead).

```
select
  df.ent_id,
  df.c_id,
  df.y_id,
  df.tid
from
  tabl t1,
  (select
    c_id
    ,com_id
    ,ent_id
    ,y_id
    ,vdate
    ,bdate
    ,date_ind
    ,unit_id
    ,rate
    ,tid
    ,rstate
  from DF06 U1
 union all
  select
    c_id
    ,com_id
    ,ent_id
    ,y_id
    ,vdate
    ,bdate
    ,date_ind
    ,unit_id
    ,rate
    ,tid
    ,rstate
  from DF08 U2
 ) as df
where
  df.com_id = 'ABCDE'
 and t1.unit_id = df.unit_id
```

## Oracle® Rdb for OpenVMS

```

and t1.pg_id = -100
and df.bdate >= '00000000'
and df.bdate <= '99999999'
and df.rstate = 'V'
and df.date_ind > ' '
and (df.ent_id = 'FOO' or 'FOO' = ' ')
--
-- The following conjuncts are missing in the strategy
--
and (df.cpty_id = 'FOOBAR' or 'FOOBAR' = ' ')
and (df.vdate >= '20071211' and df.vdate <= '20071211')
and (df.rate >= -99999999 and df.rate <= 99999999)
;
~S: Outline "BUG4_OUTLINE" used
Tables:
  0 = TAB1
  1 = DF06
  2 = DF08
Cross block of 2 entries  Q1
Cross block entry 1
  Index only retrieval of relation 0:TAB1
  Index name  TAB1_IDX [1:1]
  Keys: 0.PG_ID = -100
Cross block entry 2
  Conjunct: 0.UNIT_ID = <mapped field:28>
Merge of 1 entries  Q1
  Merge block entry 1  Q2
  Merge of 2 entries  Q2
  Merge block entry 1  Q3
  Leaf#01 FFirst 1:DF06 Card=1000
  Bool: (1.COM_ID = 'ABCDE') AND (1.BDATE >= '00000000')
        AND (1.BDATE <= '99999999') AND (1.RSTATE = 'V')
        AND (1.DATE_IND > ' ') AND ((1.ENT_ID = 'FOO') OR (
        'FOO' = ' ')) AND ((1.CPTY_ID = 'FOOBAR') OR ('FOOBAR' = ' '))
        AND (1.VDATE >= '20071211' AND (1.VDATE <=)
        '20071211') AND (1.RATE >= -99999999 AND (1.RATE <=)
        99999999)
  BgrNdx1 DF_2006_IDX [1:1] Fan=294
  Keys: <mapped field:22> = 'ABCDE'
  Bool: (1.BDATE >= '00000000' AND (1.BDATE <=)
        '99999999') AND (1.RSTATE = 'V') AND ((1.ENT_ID
        = 'FOO') OR ('FOO' = ' ')) AND ((1.CPTY_ID = 'FOOBAR') OR (
        'FOOBAR' = ' ')) AND (1.VDATE >= '20071211') AND (
        1.VDATE <= '20071211')
  Merge block entry 2  Q4
  Conjunct: (2.COM_ID = 'ABCDE') AND (2.BDATE >= '00000000')
        AND (2.BDATE <= '99999999') AND (2.RSTATE = 'V'
        )
        <== See Note:
  Conjunct: 2.VDATE >= '20071211'
  Leaf#02 FFirst 2:DF08 Card=1198
  Bool: 2.VDATE <= '20071211'
  BgrNdx1 DF_2008_IDX [1:1] Fan=294
  Keys: <mapped field:22> = 'ABCDE'
  Bool: (2.BDATE >= '00000000' AND (2.BDATE <=)
        '99999999') AND (2.RSTATE = 'V') AND ((<mapped field:23>
        = 'FOO') OR ('FOO' = ' ')) AND ((<mapped field:21> = 'FOOBAR'
        ) OR ('FOOBAR' = ' ')) AND (2.VDATE >= '20071211') AND (
        2.VDATE <= '20071211')
0 rows selected

```

Note that the following conjuncts in the strategy under the "Merge block entry 2":



## Oracle® Rdb for OpenVMS

```
Merge block entry 2 Q4
Conjunct: (2.COM_ID = 'ABCDE') AND (2.BDATE >= '00000000')
          AND (2.BDATE <= '99999999') AND (2.RSTATE = 'V'
          )
Conjunct: 2.VDATE >= '20071211'
Leaf#02 FFirst 2:DF08 Card=1198
```

are missing the following predicates:

```
and df.date_ind > ' '
and (df.ent_id = 'FOO' or 'FOO' = ' ')
and (df.cpty_id = 'FOOBAR' or 'FOOBAR' = ' ')
and (df.vdate >= '20071211' and df.vdate <= '20071211')
and (df.rate >= -99999999 and df.rate <= 99999999)
```

This problem occurs when the OR predicates are incorrectly mapped as conjuncts into the merge leg of UNION legs which are joined as the inner block of a cross join, as seen below.

```
Cross block of 2 entries
Cross block entry 1
...etc...
Cross block entry 2
Merge of 1 entries
Merge block entry 1
Merge of 2 entries
Merge block entry 1
=>Correct Conjuncts (...)
...etc...

Merge block entry 2
=>Missing conjuncts (...)
...etc...
```

The query works if the merge leg of UNION legs is placed at the outer block of the cross join, as in the following example.

```
Cross block of 2 entries
Cross block entry 1
Merge of 1 entries
Merge block entry 1
Merge of 2 entries
Merge block entry 1
=>Correct Conjuncts (...)
...etc...
Merge block entry 2
=>Correct Conjuncts (...)
...etc...
Cross block entry 2
...etc...
```

This problem has been corrected in Oracle Rdb Release 7.2.3.4.

## 4.2 SQL Errors Fixed

### 4.2.1 Unexpected Data Type Restrictions for Some SQL Operations

In prior releases of Oracle Rdb V7.2, the DECLARE variable statement, the CAST operator and routine parameter declarations did not handle correctly the following data types.

- LONG and LONG RAW types

These types were added to Oracle Rdb V7.2 as synonyms for LIST OF BYTE VARYING AS TEXT and LIST OF BYTE VARYING AS BINARY to provide compatibility with Oracle RDBMS table definitions.

```
SQL> declare :c long ;
declare :c long ;
      ^
%SQL-F-LOOK_FOR, Syntax error, looking for VARCHAR, found ; instead
SQL> declare :c long raw;
declare :c long raw;
      ^
%SQL-F-LOOK_FOR, Syntax error, looking for VARCHAR, found RAW instead
SQL>
```

- LONG VARCHAR type

This type is a synonym for VARCHAR(16383). However, the clause CHARACTER SET available for CHARACTER and CHARACTER VARYING (VARCHAR) types was not supported.

```
SQL> select cast(rdb$flags as long varchar character set isolatin9) from -
rdb$database;
select cast(rdb$flags as long varchar character set isolatin9) from
rdb$database
      ^
%SQL-F-LOOK_FOR, Syntax error, looking for ), found CHARACTER instead
SQL>
```

- CHARACTER VARYING type

This type was incorrectly interpreted as CHARACTER.

These problems have been corrected in Oracle Rdb Release 7.2.3.4. In addition, the SHOW VARIABLE statement has been enhanced to display more attributes of these data types.

### 4.2.2 Unexpected Bugcheck from DECLARE Cursor Through Dynamic SQL

Bug 7688956

In some cases, when a cursor is declared through Dynamic SQL and the SELECT expression passed is not legal SQL syntax, a bugcheck is generated. The signature of the dump is similar to that shown below.

- Alpha OpenVMS 8.3
- Oracle Rdb Server 7.2.3.1.1
- Got a SQLBUGCHK.DMP

- SQL-F-BUGCHK, There has been a fatal error. SQL\$SEMCUR - 9
- Exception occurred at SQL\$SHR72\SQL\$\$WALK\_DYN\_SELECT + 000002DC
- Called from SQL\$SHR72\SQL\$\$WALK\_DATATYPE + 00002068
- Called from SQL\$SHR72\SQL\$\$WALK\_DATATYPE + 00002AC0
- Called from SQL\$SHR72\SQL\$\$SEMANTICS + 00000110

The workaround for this problem is to correct the SQL select statement syntax and ensure that it conforms to the syntax defined by the Oracle Rdb SQL Reference Manual.

This problem has been corrected in Oracle Rdb Release 7.2.3.4. SQL will now issue an error instead of creating a bugcheck dump.

## 4.2.3 Unexpected INVQUAL Error When Using DCL SET PROCESS/PARSE\_STYLE Qualifier

Bug 7327258

In prior releases of Oracle Rdb, DCL procedures that used the SET PROCESS qualifier PARSE\_STYLE=EXTENDED would cause SQL Module Language and SQL Precompiler to not recognize command qualifiers typed in lower case. This is shown in the following example.

```
$ set process/parse_style=traditional
$ sqlmod/noinitialize_handles/connect SQL_CONNECTIONS_C.SQLMOD
$ set process/parse_style=extended
$ sqlmod/noinitialize_handles/connect SQL_CONNECTIONS_C.SQLMOD
%SQL-F-INVQUAL, Invalid qualifier noinitialize_handles specified on command
line
```

This problem has been corrected in Oracle Rdb Release 7.2.3.4.

## 4.2.4 Improved Diagnostics for INVVALIS Error

Bug 6026938

In prior releases of Oracle Rdb, an INSERT statement may fail with an INVVALIS error with little indication of what caused the error. Typically this error is issued when no explicit column list is specified and the table contains COMPUTED BY, AUTOMATIC AS or IDENTITY columns. Since these are read-only columns, they are not considered valid target columns for the default column list created by SQL. Another reason could be that the table has been altered to add or drop columns.

In either case, Rdb now displays the expected number of values as well as the actual number of values passed to the INSERT statement. The following example shows additional information.

```
SQL> create table T0 (a integer identity, b char);
SQL> insert into T0 as the_zeroth values (1, 'a');
%SQL-F-INVVALIS, The value list must have as many items as the column list
-SQL-E-INSEXPVAL, Table "T0" has 1 updatable column but was passed 2 values
```

This problem has been corrected in Oracle Rdb Release 7.2.3.4.

# Chapter 5

## Software Errors Fixed in Oracle Rdb Release 7.2.3.3

The Oracle Rdb 7.2.3.3 kit was obsoleted and replaced by Oracle Rdb 7.2.3.4. The release notes for all fixes since Rdb 7.2.3.2 shipped are contained in the "Software Errors Fixed in Oracle Rdb Release 7.2.3.4" chapter.

---

# **Chapter 6**

## **Software Errors Fixed in Oracle Rdb Release 7.2.3.2**

This chapter describes software errors that are fixed by Oracle Rdb Release 7.2.3.2.

## 6.1 Software Errors Fixed That Apply to All Interfaces

### 6.1.1 Result Data Type From NEGATE was Incorrect

Exact numeric type (TINYINT, SMALLINT and INTEGER) have an interesting property. The most negative value cannot be negated and still be represented by the same number of bits. For example, a TINYINT can hold the value  $-128$ , but the negated value ( $128$ ) requires a SMALLINT. Similarly, a SMALLINT type can hold  $-32768$  but the negated value ( $32768$ ) requires an INTEGER, and so on.

Prior releases of Oracle Rdb failed to allocate a sufficiently large data type for the result of the negate operation which could lead to integer overflow errors.

This problem has been corrected in Oracle Rdb Release 7.2.3.2. Negate, and related operations such as the ABS function, will now result in the correct type.

### 6.1.2 Floating Multiply May Produce Lower Precision Result

Bug 7383036

Oracle Rdb produces a floating point result whenever a divide operation (/) is used in a SQL or RDO query. It has been reported that in some cases only involving scaled numeric data that this floating result, when multiplied by a value and then converted back to the source data type, may contain a significant rounding error on Alpha hardware. This was obvious when comparing the results from OpenVMS systems running on VAX or Integrity hardware.

This release of Oracle Rdb has restructured the generated code for Alpha systems so that this operation maintains a larger magnitude value during the multiply (delaying the scaling of the integer value for a short time).

This change affects TINYINT, SMALLINT, INTEGER and BIGINT with scale value greater than 0. In most cases, there should be little difference in the results but, in some cases, the change of operation order will be noticeable because it preserves extra bits of accuracy.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

### 6.1.3 Zero–Length String Compare on Itanium Using Run–Time Native Compiler Could Cause Bugcheck

Bug 7457310

When using the optional run–time native code generator feature (enabled by default starting with Oracle Rdb Release 7.2.3), certain zero–length string compares could result in incorrect code being generated leading to a bugcheck at runtime due to an ACCVIO from the generated code.

The following example demonstrates one possible cause of this problem:

```
SQL> CREATE DATABASE FILENAME TESTDB;
SQL> CREATE TABLE T1 (C1 CHAR(7));
SQL> INSERT INTO T1 (C1) VALUES ('1234567');
1 row inserted
SQL> INSERT INTO T1 (C1) VALUES ('7654321');
1 row inserted
SQL> COMMIT;
SQL> SELECT DISTINCT C1, ' ' FROM T1;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DUA0:[USER]RDSBUGCHK.DMP;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DUA0:[USER]SQLBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=000000001A54000, PC=000000000641120, PS=0000001B
```

In this example, the bugcheck dump "footprint" would be similar to:

```
***** Exception at 00000000816D2040 : symbol not found
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
virtual address=000000001A54000, PC=00000000816D2040, PS=0000000B
Saved PC = 00000000809BFE30 : RDMSHRP721\RDMS$$EXE_NEXT + 00005E10
Saved PC = 00000000809C1820 : RDMSHRP721\RDMS$$EXE_NEXT + 00007800
Saved PC = 0000000080AA8A20 : RDMSHRP721\RDMS$$C_EXE_NEXT + 00000080
Saved PC = 0000000080373D20 : RDMSHRP721\RDMS_EXE_INTERP + 0000FC00
```

As a workaround for this problem, the run-time native compiler could be disabled by defining the logical name RDMS\$BIND\_CODE\_OPTIMIZATION to the value "0" or setting the debug flag CODE\_OPTIMIZATION to zero.

This problem has been corrected in Oracle Rdb Release 7.2.3.2. The run-time native compiler now generates correct code for zero-length string comparisons.

## 6.1.4 Query Slows Down Using Seq Scan After Reaching ThreLim

Bug 7384420

The simple select query slows down dramatically when the cardinality of the table hits 1651910465 and the strategy switches to "Fin Seq" after the BgrNdx1 reaches "ThreLim".

```
select * from T1 where a1=5;
Tables:
  0 = T1
Leaf#01 FFirst 0:T1 Card=1651910465
  Bool: 0.A1 = 5
  BgrNdx1 U1 [1:1] Fan=14
  Keys: 0.A1 = 5
~Estim U1 Sorted: Split lev=3, Seps=1 Est=222
~E#0003.01(1) Estim  Index/Estimate 1/222
~E#0003.01(1) BgrNdx1 ThreLim  DBKeys=0  Fetches=2+0  RecsOut=0
~E#0003.01(1) FgrNdx  FFirst  DBKeys=0  Fetches=0+0  RecsOut=0*`ABA
      A1          A2          A3
      5           1           1
      5           2           1
...etc...
      5           2047        1
~E#0001.01(1) Fin  Seq  DBKeys=1659910464  Fetches=0+3277523  RecsOut=2048
      5           2048        1
2048 rows selected
```

## Oracle® Rdb for OpenVMS

show stat

```
process statistics at 10-OCT-2008 11:04:46.45
elapsed time = 0 00:52:23.97 CPU time = 0 00:50:34.64
page fault count = 41 pages in working set = 19984
buffered I/O count = 2133 direct I/O count = 3277555
open file count = 13 file quota remaining = 1987
locks held = 159 locks remaining = 31841
CPU utilization = 96.5% AST quota remaining = 995
```

If one row is deleted from the table, the query runs fast (see following example).

```
delete from T1 where A1 = 1671910465;
commit;
select * from T1 where a1=5;
Tables:
  0 = T1
Leaf#01 FFirst 0:T1 Card=1651910465
  Bool: 0.A1 = 5
  BgrNdx1 U1 [1:1] Fan=14
  Keys: 0.A1 = 5
~Estim U1 Sorted: Split lev=3, Seps=1 Est=222
~E#0004.01(1) Estim Index/Estimate 1/222
      A1          A2          A3
      5           1           1
      ...etc...
      5           1023         1
~E#0001.01(1) BgrNdx1 EofBuf DBKeys=1024 Fetches=2+1 RecsOut=1024
~E#0001.01(1) FgrNdx FFirst DBKeys=1024 Fetches=0+3 RecsOut=1024`ABA
~E#0001.01(1) BgrNdx1 EofData DBKeys=2048* Fetches=0+0 RecsOut=1024 #Bufs=5
      5           1024         1
      ...etc...
      5           2047         1
~E#0002.01(1) Fin TTbl DBKeys=2048 Fetches=0+0 RecsOut=2048
      5           2048         1
2048 rows selected
```

The problem is caused by the overflow of the index threshold limit in BgrNdx1, where the background index reaches the threshold limit (ThreLim), abandons the foreground FgrNdx (FFirst), and finally resolves to the sequential tactic "Fin Seq".

The cardinality value of 1671910465 causes an overflow in the threshold limit of the index internally.

If the cardinality value is reduced by just one, the BgrNdx1 successfully scans the background index, switches to FgrNdx (FFirst) after it hits the end of buffer (EofBuf) at 1024 dbkeys, and finally delivers the rows from the background index.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

## 6.1.5 String Compares of More Than 255 Bytes Could Return False Matches

Bug 7560946

When using the optional run-time native code generator feature (enabled by default starting with Oracle Rdb Release 7.2.3), certain compares of fixed-length strings greater than 255 bytes could result in incorrect code



## Oracle® Rdb for OpenVMS

being generated leading to possible false matches if the leading bytes of the string were equal.

As a workaround for this problem, the run-time native compiler could be disabled by defining the logical name `RDMS$BIND_CODE_OPTIMIZATION` to the value "0" or setting the debug flag `CODE_OPTIMIZATION` to zero.

This problem has been corrected in Oracle Rdb Release 7.2.3.2. The run-time native compiler now generates correct code for string comparisons of any length.

## 6.2 SQL Errors Fixed

### 6.2.1 Incorrect COMPRESSION Attributes Applied to Vertical Row Partitioned Table

Bug 7278508

In prior releases of Oracle Rdb, the CREATE STORAGE MAP statement that used the COLUMNS clause to create a vertical row partitioned table would not apply the correct compression to each partition.

The following example shows that the CREATE STORAGE MAP specifies that only the last partition be compressed but the created map has all partitions with compression enabled.

```
SQL> create storage map t_m for t
cont>     store columns (i,tx)
cont>         disable compression
cont>         in a1
cont>     store columns (i2,tx2)
cont>         disable compression
cont>         in a2
cont>     store columns (tx3)
cont>         enable compression
cont>         in a3
cont> ;
SQL> show storage map (partition) t_m;
      T_M
For Table:          T
Partitioning is:    UPDATABLE
Store clause:       STORE columns (i,tx)
                   disable compression
                   in a1
                   store columns (i2,tx2)
                   disable compression
                   in a2
                   store columns (tx3)
                   enable compression
                   in a3

Partition information for storage map:

Vertical Partition: VRP_P001
Compression is:     ENABLED
Partition: (1) SYS_P00061
Storage Area: A1

Vertical Partition: VRP_P002
Compression is:     ENABLED
Partition: (1) SYS_P00062
Storage Area: A2

Vertical Partition: VRP_P003
Compression is:     ENABLED
Partition: (1) SYS_P00063
Storage Area: A3
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.2.3.2. Existing tables can still be accessed. After this release has been installed, the table can be unloaded, the storage map re-created and the data reloaded to achieve the desired compression characteristics.

## 6.2.2 Unexpected Bugcheck During ALTER INDEX ... BUILD PARTITION on Sorted Ranked Index

Bug 7360420

In prior releases of Oracle Rdb, it was possible that ALTER INDEX could bugcheck with a footprint similar to the one shown below. This occurred while building the bitmap for the SORTED RANKED index.

- PSIBUILD2BUILDFROMBOTTOM + 00000A2C
- PSII2CREATETREE + 00000244
- RDMS\$\$KOD\_CREATE\_TREE + 000001E4
- RDMS\$\$KOD\_CREATE\_INDEX + 000005DC
- RDMS\$\$CREATE\_INDEX\_INFO + 00004CFC
- RDMS\$\$RELEASE\_DDL\_VM\_HNDLR + 00000A24
- BLI\$CALLG + 000000BC
- KOD\$SETSTK\_AND\_CONTINUE + 0000019C

This problem can be avoided by using either CREATE INDEX or ALTER INDEX ... BUILD ALL PARTITIONS.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

## 6.2.3 Change in Handling of Dynamic Cursor Names

In prior versions of Oracle Rdb, cursor names passed to the DECLARE CURSOR statement (extended dynamic) were implicitly upper-cased. However, we have observed considerable CPU overhead for this action. Interfaces such as Oracle SQL/Services, OCI Services for Rdb, Native JDBC Driver, and so on, all use a consistent naming scheme for cursors and do not benefit from the upper-casing of the name. Therefore, with this release of Oracle Rdb, we no longer uppercase the cursor name passed to the dynamic SQL statement.

Applications will typically create and use a consistent case for cursor names and should see no change in behavior from this performance change to Dynamic SQL.

## 6.2.4 Unexpected DISABLED Indices After Database Created by IMPORT DATABASE Statement

Bug 7578428

The ALTER INDEX ... BUILD ALL PARTITIONS statement implicitly marks the index as MAINTENANCE IS ENABLED IMMEDIATE. Unfortunately, if the index was originally marked as MAINTENANCE IS DISABLED, a subsequent EXPORT DATABASE and IMPORT DATABASE would recreate the index again marked as MAINTENANCE IS DISABLED.

## Oracle® Rdb for OpenVMS

The following example shows this sequence. Note that the index after the import is marked as ***Index is no longer maintained.***

```
SQL> create table my_tab (my_col1 integer, my_col2 integer);
SQL> insert into my_tab value (1,1);
1 row inserted
SQL> create index my_ndx on my_tab (my_col2);
SQL> alter index my_ndx maintenance is disabled;
SQL> alter index my_ndx truncate all partitions;
%RDB-W-META_WARN, metadata successfully updated with the reported warning
-RDMS-W-IDXBLDAPEND, index in build pending state - maintenance is disabled
SQL> alter index my_ndx build all partitions;
SQL> show index *
User indexes in database with filename testdb
Indexes on table MY_TAB:
MY_NDX                with column MY_COL2
  Duplicates are allowed
  Type is Sorted
  Key suffix compression is DISABLED
SQL> commit;
SQL> disconnect all;
SQL>
SQL> export database filename testdb into save_testdb;
SQL> drop data filename testdb;
SQL> import database from save_testdb filename testdb;
SQL>
SQL> show index *
User indexes in database with filename testdb
Indexes on table MY_TAB:
MY_NDX                with column MY_COL2
  Duplicates are allowed
  Type is Sorted
  Key suffix compression is DISABLED
  Index is no longer maintained
SQL>
```

This problem is caused by stale metadata (in the Rdb\$EXTENSION\_PARAMETERS) column for the index.

To avoid this problem, the index maintenance can be changed as follows:

1. Name the partitions for the index to make it easier to specify these names in the various ALTER INDEX commands.

```
SQL> create index MY_NDX2 on MY_TAB (my_col2)
cont> store in RDB$SYSTEM (partition MY_NDX2);
```

For an existing index, an ALTER INDEX ... RENAME PARTITION ... statement can be used to override the Rdb-generated names for each partition.

2. Replace the BUILD ALL PARTITIONS action with a BUILD PARTITION for each partition.

```
SQL> alter index MY_NDX2 build partition MY_NDX2;
%RDB-W-META_WARN, metadata successfully updated with the reported warning
-RDMS-W-IDXBLDAPEND, index in build pending state - maintenance is disabled
```

3. Finally, execute the ALTER INDEX command to set MAINTENANCE as ENABLED IMMEDIATE. This command correctly removes the stale metadata.

```
SQL> alter index MY_NDX2 maintenance is enabled immediate;
```

## Oracle® Rdb for OpenVMS

This problem has been corrected in Oracle Rdb Release 7.2.3.2. ALTER INDEX now removes the stale setting from the index allowing EXPORT and IMPORT to process the index correctly.

## 6.3 RMU Errors Fixed

### 6.3.1 Use of External Media Manager or Tape Librarian Causes a Hang or an ACCVIO

An RMU command can exit with an ACCVIO or hang, especially when the command is using a media manager or tape librarian. In some cases, the per thread stack might have been exhausted and be the cause of such a hang or ACCVIO. But there could be other reasons for this situation as well.

This modification adds a no-access virtual memory page to either side of a per thread stack. This causes an ACCVIO if a thread tries to use stack space beyond the stack limits.

Here is how to find out whether or not the ACCVIO was caused by a stack size too small. Repeat the command with the following logicals defined to see the current stack limits:

```
$ CREATE/NAME_TABLE COSI_ENV_USER
$ DEFINE/TABLE=COSI_ENV_USER RMU$DEBUG_FLAGS -
    "RMU_DEBUG_TIME,RMU_DEBUG_ANY"
$ RMU...
15:59:52.80: BF thread DIO: X00C86010 length: X0138 STKLIM: X00C8A000:X00C92000
15:59:52.80: DBF thread DIO: X00C94160 length: X0140 STKLIM: X00C98000:X00CA0000
15:59:52.80: DBF thread DIO: X00CA22B0 length: X0140 STKLIM: X00CA6000:X00CAE000
%SYSTEM-F-ACCVIO, access violation, reason mask=01, -
    virtual address=0000000000C89FB0, PC=...
```

The example shows the virtual address of the ACCVIO 00C89FB0 is just below the low stack limit 00C8A000 and within the guard page. This indicates a possible stack size problem and can be avoided by increasing the stack size per thread.

To add extra stack space, define the following logical:

```
$ DEFINE RDM$BIND_LIBRARIAN_ADDITIONAL_STACK 10240
...adds another 10KB stack space per RMU thread.
```

This logical can be used even if no media manager or tape librarian is being used.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

### 6.3.2 RMU/BACKUP Sometimes Dismounts Last Volume in a Multi-tape Drive Case

When running an RMU Backup using multiple tape drives, the last volume is sometimes dismounted.

In a multi-tape drive case, all volumes except the last one are so-called continuation volumes and cannot be appended to. Hence these volumes are dismounted at the end of the backup. The last volume can still be appended to and is therefore not to be dismounted.

This was caused by a false setting of a last volume flag. This problem has been corrected in Oracle Rdb Release 7.2.3.2.

### 6.3.3 Parallel Backup Hangs

Bug 6628917

A parallel backup can hang depending on the number and size of storage areas and the number of tape drives being used.

This is caused by a feature in RMU Backup that waits for the first worker process to save the system area before other worker processes can proceed. In some cases, the system area was not assigned to executor process WORKER\_001 and therefore it caused a hang.

As a workaround, the plan file can be edited and the line with the system area can be moved in front of the areas assigned to WORKER\_001.

Example of PLAN file listing the system area RDB\$SYSTEM first in the storage area list for WORKER\_001:

```

Executor Parameters :
  Executor Name = WORKER_001
  Executor Type = Worker
  Start Storage Area List
    RDB$SYSTEM,
    A0010,
    ...

```

With this change, the system area is always assigned to the top of the list for WORKER\_001.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

### 6.3.4 Users Were Audited Who Held Identifiers Not Enabled For Auditing

Bug 6531721

If you specify the command

```
$ RMU/SET AUDIT=IDENTIFIER=(X) database
```

identifier X is added to a list of enabled user identifiers whose access to an Oracle Rdb database can be audited. If a user does not hold an identifier that matches this list, then Oracle Rdb will not send audit entries to VMS for the database user when auditing is started. The following command shows the identifiers currently enabled for audit.

```
$ RMU/SHOW AUDIT/IDENTIFIERS database
```

Because of a problem checking user identifiers against the list of identifiers currently enabled for auditing, users of Rdb databases who did not hold one of these identifiers were audited. This problem has been fixed and now only users holding enabled identifiers will be audited when they access Rdb databases.

The following example of this problem shows that although auditing is enabled only for "USERA", if, when auditing is started, another user accesses the database who holds an identifier which has not been enabled for auditing, auditing entries for that user have been created when they should not have been.

```

$ RMU/SET AUDIT/ENABLE=IDENTIFIER=(USERA) MF_PERSONNEL
$ RMU/SET AUDIT/ENABLE=DACCESS=DATABASE/PRIVILEGES=(ALL) MF_PERSONNEL
$ RMU/SET AUDIT/ENABLE=(DACCESS,PROTECTION,RMU) MF_PERSONNEL
$ RMU/SET AUDIT/EVERY MF_PERSONNEL
$ RMU/SET AUDIT/FLUSH MR_PERSONNEL
$ RMU/SET AUDIT /TYPE=AUDIT /START MF_PERSONNEL
$ RMU/SET AUDIT/ENABLE=DACCESS=TABLE=(EMPLOYEES) -
  /PRIVILEGES=(ALTER, CREATETAB, DBADM, DBCTRL, DELETE, DROP, INSERT, -
  SECURITY, UPDATE, FAILURE) MF_PERSONNEL
$ RMU/SHOW AUDIT/IDENTIFIERS MF_PERSONNEL
Enabled identifiers:
  (IDENTIFIER=[RDB,USERA])

$ SQL
ATT 'FILENAME MF_PERSONNEL';
SELECT DISTINCT AUDIT$TIME_STAMP,AUDIT$OPERATION,AUDIT$USER_NAME,
AUDIT$OBJECT_NAME FROM AUDIT_RECORDS -
WHERE AUDIT$USER_NAME = 'USERB';
  AUDIT$TIME_STAMP
    AUDIT$OPERATION
    AUDIT$OBJECT_NAME
22-SEP-2008 11:48:03.56
  Attach Database
    DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1
  >>
  >>
  >>
22-SEP-2008 11:48:20.45
  Insert Record
    EMPLOYEES
  >>
  >>
  >>
22-SEP-2008 11:48:51.71
  Delete Record
    EMPLOYEES
  >>
  >>
  >>
3 rows selected

```

There is no workaround to avoid this problem except to only display audit entries in SQL for named users.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

## 6.3.5 RMU-W-AIJBCKFAIL – RMU/RECOVER Reports Incomplete AIJ Backup

Using /FORMAT=NEW\_TAPE to backup and restore a journal file may report a warning at the end of the recover operation.

```

$ RMU /RESTORE /NOLOG /NOCDD /AFTER MFP.RBF
$ RMU /RECOVER /NOLOG MFP.ABF /FORMAT=NEW_TAPE
...

```



## Oracle® Rdb for OpenVMS

%RMU-W-AIJBCKFAIL, the AIJ backup that created the AIJ file did not complete  
%RMU-I-AIJALLDONE, after-image journal roll-forward operations completed  
%RMU-I-AIJFNLSEQ, to start another AIJ file recovery, the sequence number  
needed will be 0

This would happen when an RMU/BACKUP/AFTER\_JOURNAL command would exactly fill an internal buffer when saving the last AIJ record. An internal block counter was not correctly incremented and hence the final AIJ\_END record was tagged with a duplicate block number. The problem remained undetected during the save operation. All AIJ records were restored during a recover operation. However, the reported next sequence number was incorrect which could prevent a further recovery operation.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

## 6.4 LogMiner Errors Fixed

### 6.4.1 LogMiner Incorrectly Returns DBK Multiple Times Within Transaction

Bug 7532856

Starting with Oracle Rdb Release 7.2.3, the `RMU /UNLOAD /AFTER_JOURNAL` command could, in rare cases, return multiple versions of a row with the same database key within one transaction (this case is not allowed; only the final content of a row is to be returned within a transaction). The cause of the problem was related to an incorrect internal sort comparison not correctly handling duplicate database key removal.

This problem has been corrected in Oracle Rdb Release 7.2.3.2.

---

# **Chapter 7**

## **Software Errors Fixed in Oracle Rdb Release 7.2.3.1**

This chapter describes software errors that are fixed by Oracle Rdb Release 7.2.3.1.

# 7.1 Software Errors Fixed That Apply to All Interfaces

## 7.1.1 Page Transfer Via Memory Lost Database Update

Bug 7243961

In very rare cases, when using the "Page Transfer Via Memory" feature, a database recovery (DBR) process may incorrectly avoid re-applying a database update. The conditions leading to this situation are complex and rare.

As a workaround, users of the "Page Transfer Via Memory" feature may wish to consider disabling it.

This problem has been corrected in Oracle Rdb Release 7.2.3.1. The database recovery (DBR) process starts applying updates from the correct point in the AIJ file.

## 7.1.2 Rollback Performance Improvement with Row Cache

Performance of ROLLBACK operations has been improved when the Row Cache feature is enabled for a database. Previously, all working set entries in all caches for the user were evaluated even if the transaction did not modify any cache or perform any cache latch operations.

This problem has been corrected in Oracle Rdb Release 7.2.3.1. Working set entries are only checked if a transaction has modified any cache or performed any cache latching operations.

## 7.1.3 Wrong Result from Query with Shared NULL Conjunct in OR Predicate

Bug 7254564

The following query returns the wrong result (100 rows instead of 0 rows).

```
set flags 'strategy,detail';
select employee_id, first_name, last_name
from employees e1
where
  (employee_id is not null and
   (employee_id is null or
    not exists (select * from employees e2
                where e2.employee_id = e1.employee_id
                   and e2.last_name = e1.last_name)));
```

Tables:

0 = EMPLOYEES

1 = EMPLOYEES

Cross block of 2 entries

Cross block entry 1

Leaf#01 FFirst 0:EMPLOYEES Card=100

Bool: NOT MISSING (0.EMPLOYEE\_ID)

BgrNdx1 EMP\_EMPLOYEE\_ID [0:1] Fan=17

## Oracle® Rdb for OpenVMS

```
Keys: NOT MISSING (0.EMPLOYEE_ID)
Cross block entry 2
Aggregate-F1: 0:COUNT-ANY (<subselect>)
Conjunct: 1.LAST_NAME = 0.LAST_NAME
Get      Retrieval by index of relation 1:EMPLOYEES
Index name EMPLOYEES_HASH [1:1]      Direct lookup
Keys: 1.EMPLOYEE_ID = 0.EMPLOYEE_ID
Bool: NOT MISSING (1.EMPLOYEE_ID)
EMPLOYEE_ID  FIRST_NAME  LAST_NAME
00164        Alvin        Toliver
...etc...
100 rows selected
```

In Rdb Release 7.2–200, the same query returns 0 rows.

Tables:

```
0 = EMPLOYEES
1 = EMPLOYEES
Cross block of 2 entries  Q1
Cross block entry 1
Leaf#01 FFirst 0:EMPLOYEES Card=100
Bool: NOT MISSING (0.EMPLOYEE_ID)
BgrNdx1 EMP_EMPLOYEE_ID [0:1] Fan=17
Keys: NOT MISSING (0.EMPLOYEE_ID)
Cross block entry 2
Conjunct: MISSING (0.EMPLOYEE_ID) OR (<agg0> = 0)
Aggregate-F1: 0:COUNT-ANY (<subselect>) Q2
Conjunct: 1.LAST_NAME = 0.LAST_NAME
Get      Retrieval by index of relation 1:EMPLOYEES
Index name EMPLOYEES_HASH [1:1]      Direct lookup
Keys: 1.EMPLOYEE_ID = 0.EMPLOYEE_ID
Bool: NOT MISSING (1.EMPLOYEE_ID)
0 rows selected
```

This problem is caused by the NULL conjunct in one of the operands in the OR predicate, shared by the NOT NULL conjunct at the outer leg of a cross strategy.

The key parts of this query which contributed to the error are:

1. The WHERE clause contains a NULL conjunct as one of the operands in an OR predicate which is an operand to an AND predicate with the other operand being a NOT NULL conjunct on the same column.
2. The other operand of the OR predicate is a NOT EXISTS clause checking the existence of the joined operation between the same two tables.

This problem has been corrected in Oracle Rdb Release 7.2.3.1.

## 7.2 SQL Errors Fixed

### 7.2.1 Unexpected SQL-F-DATTYPUNK Error from Dynamic SQL Applications

Bug 7249833

In Oracle Rdb Release 7.2.3, a problem was introduced in the handling of parameters to Dynamic SQL states for DECODE and the Simple CASE expression. Dynamic fails to assign a data type in some cases and reported a SQL-F-DATTYPUNK error. The following example is run using a Dynamic SQL testing tool and shows the reported problem.

```
$          RUN TEST$TOOLS:TESTER
attach 'filename PERSONNEL';
set dialect 'oracle level2';
select E.LAST_NAME, E.FIRST_NAME
from EMPLOYEES E
where EMPLOYEE_ID = DECODE(:b1, NULL ,E.EMPLOYEE_ID, :b2);
Error -1:
%SQL-F-DATTYPUNK, Data type unknown. Expression cannot use only host variables
select E.LAST_NAME, E.FIRST_NAME
from EMPLOYEES E
where EMPLOYEE_ID = (case :b1 when NULL then E.EMPLOYEE_ID else :b2 end);
Error -1:
%SQL-F-DATTYPUNK, Data type unknown. Expression cannot use only host variables
```

This problem has been corrected in Oracle Rdb Release 7.2.3.1.

### 7.2.2 Updated Support for SQL Pseudo Types for C Language

Bug 7248505

In prior versions of Oracle Rdb, SQL date/time types were supported by the SQL precompiler for C using a struct of two long fields (10,11). This caused informational errors from the C compiler when using /NOMEMBER\_ALIGNMENT. In this release, Oracle Rdb has been changed to no longer generate this struct and instead use an int64 variable for SQL date/time values declared using SQL\_TIMESTAMP, SQL\_TIME, SQL\_DATE\_ANSI, SQL\_DATE\_VMS, and SQL\_INTERVAL pseudo type definitions.

The following example shows the error previously generated by the C compiler.

```
$ sql$pre/cc=decc b /nomember_alignment/warn=enable=alignment

    long 10,11;
    .....^
%CC-I-MISALGNDSTRCT, This member requires longword alignment for efficient
access, but is contained in a struct containing byte alignment. Consider
using #pragma nomember_alignment longword.
....
```

This problem has been corrected in Oracle Rdb Release 7.2.3.1.

## Note

*Although the internal layout was not documented nor was the struct supported for customer use by Oracle, it is possible that existing applications reference the fields of the struct. If this is the case, then it is possible to redefine the struct for use in applications. Include this definition at the top of the C module that requires the field references.*

```
#pragma member_alignment save
#pragma member_alignment
typedef struct {
    long l0,l1;
} SQL_DATE_ANSI_X;
#define SQL_DATE_ANSI SQL_DATE_ANSI_X
#pragma member_alignment restore
```

The macro will be used to replace the definition generated by the SQL precompiler.

## 7.2.3 Unexpected SQL-F-NUMXPRESX Error While Processing ABS Function

In prior releases of Oracle Rdb V7.2, the ABS function used with a parameter marker in Dynamic SQL could return the SQL-F-NUMXPRESX error. The following example shows such an example:

```
select * from rdb$database where rdb$flags = ABS(?);
Error -1:
%SQL-F-NUMXPRESX, A numeric expression was expected
```

When SQL cannot determine the type of the ABS function source, it should default to DOUBLE PRECISION as shown in this example.

```
select * from rdb$database where rdb$flags = ABS(?);
[1 fields]
0//G Float: 0.1e0
```

This problem has been corrected in Oracle Rdb Release 7.2.3.1.

## 7.2.4 Unexpected Truncation of Application Parameter Data in Dynamic SQL

Bug 1523789

In prior releases of Oracle Rdb, Dynamic SQL would set the data type of a parameter marker based on context. Often this data type might have been resolved from a string literal and in many cases this string literal may have been smaller than the source application parameter. This resulted in truncation of the input and might have led to incorrect query results. Consider this simple example:

```
select * from PORTS where city = ? and ? <> 'Boston';
```

In this example, the second parameter was described in the SQLDA as CHAR(6) inheriting the length of the string literal, but the input source from the application was really a CHAR(20). This resulted in application data being truncated when passed to the SQL query.

This problem has been corrected in Oracle Rdb Release 7.2.3.1. In this release, any parameter which derives its data type from a string literal will be promoted to VARCHAR(2000) to ensure sufficient space is available for application parameter data.

## 7.2.5 Unexpected Prompt When Adding a Comment to an Index

In prior versions of Oracle Rdb, the ALTER INDEX statement tried to prevent unintended changes to the STORE clause by confirming this action with the prompt "This index was previously specified with a STORE clause. Continue? [N]". However, this prompt was incorrectly issued when only the COMMENT IS clause was used.

The following example shows the problem.

```
SQL> ALTER INDEX foo COMMENT 'SUPPORTS CONSTRAINT bar';  
This index was previously specified with a STORE clause. Continue? [N]N  
%SQL-F-CHGINDMAPSTP, Terminating operation at user's request
```

The workaround is to use the COMMENT ON INDEX statement or answer Y to the prompt.

This problem has been corrected in Oracle Rdb Release 7.2.3.1.



## 7.3 RMU Errors Fixed

### 7.3.1 RMU/BACKUP Exits with ACCVIO in ENCRYPTSHR

An RMU/BACKUP/ENCRYPT can fail with an ACCVIO with the following pattern in the RMUBUGCHK dump:

```
***** Exception at 0000000000F5C730 : Image ENCRYPT$ALG$AES + 00032730
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=0000000000BE6000, PC=0000000000F5C730, PS=0000001B
Saved PC = 0000000000E9AAEC : Image ENCRYP$SHR + 00082AEC
Saved PC = 0000000000393808 : RMU72\RMUIO$CRYPT_INIT_BCK + 000002F8
```

The problem was caused by a buffer overrun in OpenVMS ENCRYPTSHR code. The problem has been reported to OpenVMS Engineering.

This problem has been corrected in Oracle Rdb Release 7.2.3.1. The RMU code has been modified to protect against the buffer overrun.

## 7.4 RMU Show Statistics Errors Fixed

### 7.4.1 RMU/SHOW STATISTICS Does Not Show BLOB Areas On Logical Area Information Screen

Bug 7242685

When selecting BLOB logical areas on the "Logical Area Overview" display, if system logical areas have not been enabled, BLOB logical areas will not be displayed. If the tools main item "Display all logical areas" is used, then the BLOB logical areas will be displayed.

This problem has been corrected in Oracle Rdb Release 7.2.3.1. When selecting BLOB logical areas, they will be displayed regardless of the "Display all logical areas" setting.

---

# **Chapter 8**

## **Software Errors Fixed in Oracle Rdb Release 7.2.3.0**

This chapter describes software errors that are fixed by Oracle Rdb Release 7.2.3.0.

# 8.1 Software Errors Fixed That Apply to All Interfaces

## 8.1.1 Inability to Cause Free Space Scan for First Insert After Attach

Bug 6807755

Oracle Rdb Release 7.2 introduced a mechanism that allows database users on the same cluster node to share information regarding the availability of free space. This optimization helps reduce cases of unexpected storage area growth when one process is inserting into a table while another process on the same node is deleting from the same table.

In a cluster environment, however, there is no cross-node sharing regarding the availability of free space within a logical area. In prior releases of Oracle Rdb, a periodic detach and attach could be used to force a process to initially scan for the next available free space when inserting a record. However, this technique would not always work with Oracle Rdb Release 7.2 due to the shared insert location.

This problem has been corrected in Oracle Rdb Release 7.2.3. In a cluster environment, a process will ignore the shared information regarding the availability of free space for the process's first insert into a logical area.

## 8.1.2 Insert ACCVIO Bugcheck on IA64

Bug 6820142

An ACCVIO bugcheck could occur during an insert into a table with many constraints. This was only an issue on Itanium and also only when running in "Interpreter Mode", which was the default for versions prior to 7.2.3.

```
***** Exception at FFFFFFFF80004090 : symbol not found
%SYSTEM-F-ACCVIO, access violation, reason mask=04, virtual address=
FFFFFFFFFFFFFF8870, PC=FFFFFFFF80004090, PS=0000000B
Saved PC = 0000000080328710 : RDMSHRP731\RDMS_EXEINT_CREATE_IKEY + 000000F0
Saved PC = 00000000803282C0 : RDMSHRP731\RDMS_EXE_INTERP + 000446F0
```

A possible workaround is to use the Run-Time Routine Native Compiler which is enabled by default starting in this release, Release 7.2.3.

Please see the "Enhancements and Changes Provided in Oracle Rdb Release 7.2.3" Chapter. The first entry is entitled "Optional Run-Time Routine Native Compiler on I64" and gives more details about how to use this feature.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.1.3 Processes With Multiple Database Attaches Stall "Waiting for RTUPB list" on IA64

Bug 6707976

In rare cases, a process with multiple streams (meaning attached to databases multiple times) could get into a condition where one stream needed to execute and was able to run but another stream was hibernating. A change in implementation during the IA64 port resulted in the hibernation happening in the stream's context rather than in the internal Rdb thread scheduler (as it would have been on Alpha). This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.1.4 Multi-Threaded Processes Hang in HIB Status

Bug 6707976

In some cases, multi-threaded processes may appear to hang in a hibernating state when using Oracle Rdb. This condition can be due to both Oracle Rdb and the OpenVMS user-mode process thread scheduler using hiber/wake sequences and Oracle Rdb not having always correctly "re-armed" the process wake-pending state by issuing a \$WAKE.

This problem has been corrected in Oracle Rdb Release 7.2.3. Oracle Rdb now issues a \$WAKE to the process after a hibernate stall has completed.

---

Application and Oracle Rdb Both Using SYSSHIBER

*Oracle Rdb's use of the \$WAKE system service will 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, applications that use HIBER/WAKE facilities must use a code sequence that avoids continuing without a check for the operation (such as a delay or a timer firing) being complete.*

---

## 8.1.5 Tape Savesets With More Than 999,999 Blocks Report Wrong Block Number

Writing savesets to tape with more than 999,999 blocks reports a wrong block number using the DCL DIRECTORY command with /SIZE.

This is caused by using only 6 bytes in the ANSI HDR1 and EOF1 labels to store the block count in ASCII. A recent change in OpenVMS (BACKUP, MTAACP) extended the use of the implementation specific field in these labels to use 4 more bytes to store the high order digits for block counts larger than 999,999.

This problem has been corrected in Oracle Rdb Release 7.2.3. RMU has been changed to also use the extra 4 bytes for larger tape savesets.

## 8.1.6 Possible DBR Failure Due to Address Space Being Exhausted

Bug 7033216

In rare cases, it is possible for the Database Recovery (DBR) process to fail with a "footprint" similar to the following:

```
***** Exception at 00000000002EAFE0 : RDMDBR72\KUTREC$RV_CREATE + 000001E0
%COSI-F-VASFULL, virtual address space full
-SYSTEM-W-REGISFULL, specified region is full
Saved PC = 00000000002ECE60 : RDMDBR72\KUTREC$RV_GET + 000000B0
Saved PC = 00000000002EC980 : RDMDBR72\KUTREC$RV_FIND + 000004D0
Saved PC = 0000000000EFA60 : RDMDBR72\DBR$PROCESS_AIJBUF + 00001000
Saved PC = 0000000000107BC0 : RDMDBR72\DBR$REDO + 00001C50
Saved PC = 000000000010A930 : RDMDBR72\DBR$REDO_LOOP + 000001B0
Saved PC = 0000000000F8B20 : RDMDBR72\DBR$RECOVER_USER + 00001390
Saved PC = 0000000000F70D0 : RDMDBR72\DBR$RECOVER + 00000A50
```

This problem can be caused by the Database Recovery process failing to release recovery buffers in unusual circumstances. This problem may be more likely when using the row cache feature. It is also possible that other Database Recovery (DBR) process "out of memory" symptoms may be caused by the same problem.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.1.7 Control Number and Size of AIJ Initialization I/O Buffers

Bug 7040647

When an AIJ backup operation completes, the after image journal file(s) are initialized with a pattern of -1 (hex FF) bytes. This initialization is designed to be fast and thus may effectively utilize the I/O subsystem by performing large, asynchronous I/Os. This speed can, however, come at the cost of a high load on I/O components during the initialization. In rare cases, this load could slow down other I/Os on the system.

In order to allow control over the relative I/O load from the AIJ initialization operation, two logical names may be used.

- **RDM\$BIND\_AIJ\_INITIALIZE\_IO\_COUNT** – Specifies the number of asynchronous I/O operations that will be queued at once to the AIJ file. The default value if the logical name is not defined is 2, the minimum value is 1 and the maximum value is 128.
- **RDM\$BIND\_AIJ\_INITIALIZE\_IO\_SIZE** – Specifies the number of 512-byte disk blocks to be written per I/O. The default value if the logical name is not defined is 256 for fixed length circular journals and 32 for single extensible journals. The minimum value is 4 and the maximum value is 256.

Reducing the value of either logical might increase the amount of time needed to initialize the AIJ file after a backup. However, it may also reduce queue load on the I/O subsystem.

In prior versions of Oracle Rdb, these logical names would be translated only at database open time. They are

now also translated when processes attach to the database. This problem has been corrected in Oracle Rdb Release 7.2.3.

The values of the number of asynchronous I/O operations and number of blocks to be written per I/O for AIJ initialization may also be controlled with the RMU SHOW STATISTICS utility "AIJ Dashboard" screen using the dashboard items "Init IO Count" and "Init IO Size".

## 8.1.8 Some Sort Operations May Bugcheck With More Than 134 Million Rows

Bug 7113254

In some cases, a sort operation with more than 134 million rows can cause a bugcheck. For example, when building an index on a huge table, as in the following example.

```
SQL> SET TRANSACTION READ WRITE RESERVING MyTable FOR EXCLUSIVE WRITE;  
SQL> CREATE INDEX MyIndex ON MyTable (  
_SQL> Field1,Field2,...,Fieldn ) TYPE IS SORTED STORE IN MyIndexTable );  
%RDMS-I-BUGCHKDMP, generating bugcheck dump file ...
```

The problem was caused by an erroneous check for the maximum number of rows included in a sort sequence.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.1.9 DBR Hangs on Page Lock After Node Crash

Bug 7029177

Starting with Oracle Rdb V7.1–2, under certain rare circumstances, it was possible for an Oracle Rdb Database Recovery (DBR) process to hang during database recovery after a node failure. Because the DBR process in this case was hung, all users of the database were blocked.

For this problem to occur, the database must be accessed on multiple nodes, including the one that fails. During cluster recovery, an Rdb monitor on a surviving node will create a DBR process to recover the users on the failed node. It is possible that the DBR may become blocked waiting on a page lock held by another user who, in turn, is waiting on a MEMBIT lock for the recovery to complete. Due to use of the NODLCKWT and NODLCKBLK flags on some of the locks involved, it may not be possible for the OpenVMS lock manager to detect possible deadlocks in this situation.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.1.10 Alter Storage Map ... Reorganize Pages Corrupts SORTED Indices

Bug 7022352

An ALTER STORAGE MAP ... REORGANIZE AREAS or ALTER STORAGE MAP ... REORGANIZE PAGES could cause some SORTED indices to become corrupt.

A RMU/VERIFY might show an index corruption similar to:

```
%RMU-W-BADIDXREL, Index TP_INDEX either points to a non-existent record or
                    has multiple pointers to a record in table TP.
                    The logical dbkey in the index is 3128:11:0.
%RMU-W-DATNOTIDX, Row in table TP is not in any indexes.
                    Logical dbkey is 3135:198:7.
```

A possible workaround would be to use SORTED RANKED indices.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.1.11 RDMALS Bugcheck at AIJUTL\$FORMAT\_ARBS + 00000DA4

Bug 6523851

On rare occasions, the AIJ Log Server (ALS) may fail with an exception at AIJUTL\$FORMAT\_ARBS. The bugcheck Stack Dump Summary will be similar to:

```
***** Exception at 00000000000DB154 : RDMALS72\AIJUTL$FORMAT_ARBS + 00000DA4
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 00000000000C58C8 : RDMALS72\ALS$FLUSH_ONE_CACHE + 000015C8
Saved PC = 00000000000C3C64 : RDMALS72\ALS$FLUSH + 00000664
Saved PC = 00000000000C27F4 : RDMALS72\ALS$MAIN + 00000F94
Saved PC = FFFFFFFF8013075C : symbol not found
```

The actual offset may be different depending on the platform and Rdb version being run.

This has been diagnosed to be a timing issue while the ALS is preparing updates to the After Image Journal (AIJ) file. This problem does not lead to data corruption or missing AIJ entries.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.1.12 Bugcheck Due to Command Line Too Long

Bug 7196946

The SQL Precompiler and SQL Module Language compiler support a maximum command line length of 255 characters. This length limit was not properly enforced and excessively long command lines could cause memory corruption leading to bugchecks or unexpected results.

This problem has been corrected in Oracle Rdb Release 7.2.3. The SQL Precompiler and SQL Module Language compiler now return an error message for command lines that are greater than 255 characters, as in the following example.

```
$ SQL$PRE72 -
  /CC/SQLOPT=(CONN,C_STR=NOBLANK,NOEXTERNAL_GLOBALS, -
  NODECLARE_MESSAGE_VECTOR,NOTRANSACTION)/NOMEMB -
  /NAME=SHORT/NEST=NONE/NOOPT/CHECK=NOFP_MODE/FLOAT=IEEE_FLOAT -
  /IEEE_MODE=FAST/WARNING=(DISABLE=BADBOUNDCHK)/DEBUG -
  /OBJECT=PND_XXXXX/INCL=PND_XXXXX/LIST=PND_XXXXX -
```



## Oracle® Rdb for OpenVMS

```
EMPTY__MODULES_0.SC ENV_SC.SQL_CTX
%SQL-F-CMDLN_TOO_LONG, The command line is too long.
Please shorten the command
```

## 8.2 SQL Errors Fixed

### 8.2.1 ALTER PROFILE Generates a SQL Bugcheck Dump

In prior releases of Oracle Rdb, the ALTER PROFILE command would generate a SQL bugcheck when an attempt was made to change or add a DEFAULT TRANSACTION. The following example shows the problem.

```
SQL> create profile default_profile
cont> ;
SQL> alter profile default_profile
cont> default transaction read only;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TESTING]SQLBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=04,
virtual address=0000000000000000, PC=000000000024D484, PS=0000001B
```

To avoid this problem, use DROP PROFILE to remove the profile and recreate it. This might require using the CASCADE option and then later altering each user to replace the PROFILE.

This problem has been corrected in Oracle Rdb Release 7.2.3.

### 8.2.2 Incorrect File Specification Reported for RDB\$STORAGE\_AREAS Information Table

Bugs 6763712 and 6006502

Starting with Oracle Rdb V7.2, the domain RDB\$FILE\_NAME was changed to be a COMPUTED BY column that returns the file specification for the database root. For the table RDB\$DATABASE, this allowed space saving in the root file. Unfortunately, this domain was inadvertently used by several information tables. This led to unexpected results for those file specifications from these tables.

The following example shows that when selecting the names of the snapshot files from the table RDB\$STORAGE\_AREAS, only the root file specification is returned.

```
SQL> SELECT RDB$AREA_FILE AS SNAP_NAME EDIT USING 'T(50)',
cont> RDB$INITIAL_ALLOCATION AS INITIAL EDIT USING 'Z(9)',
cont> RDB$CURRENT_ALLOCATION AS CURRENT EDIT USING 'Z(9)'
cont> FROM RDB$STORAGE_AREAS
cont> WHERE (RDB$CURRENT_ALLOCATION = RDB$INITIAL_ALLOCATION) AND
cont> (RDB$AREA_NAME = ' ') limit to 3 rows;
SNAP_NAME INITIAL CURRENT
DISK14:[TESTING.DATABASE]SAMPLE.RDB;1 5000 5000
DISK14:[TESTING.DATABASE]SAMPLE.RDB;1 500 500
DISK14:[TESTING.DATABASE]SAMPLE.RDB;1 1000 1000
3 rows selected
```

This problem has been corrected in Oracle Rdb Release 7.2.3. The INFO\_TABLES.SQL script now uses a new domain RDB\$FILE\_SPECIFICATION. If these information tables exist in a V7.2 database, they must be dropped and replaced using the revised SQL\$SAMPLE:INFO\_TABLES.SQL script.

#### *Example*

## Oracle® Rdb for OpenVMS

This example shows the execution of the SQL\$SAMPLE:INFO\_TABLES.SQL script to replace the information tables.

```
SQL> attach 'filename MF_PERSONNEL';
SQL> @info_tables
```

Copyright (c) 1997, 2008, Oracle Corporation. All Rights Reserved.

Please ignore any FIELD\_EXISTS error for the domain RDB\$FILE\_SPECIFICATION %SQL-F-FIELD\_EXISTS, Domain RDB\$FILE\_SPECIFICATION already exists in this database or schema

Creating RDB\$CACHE

Creating RDB\$DATABASE\_JOURNAL

Creating RDB\$DATABASE\_ROOT

Creating RDB\$DATABASE\_USERS

%RDB-E-NO\_META\_UPDATE, metadata update failed

-RDMS-F-VIEWEXI, relation RDB\$DATABASE\_USERS is referenced by view V\$X

-RDMS-F-RELNOTDEL, relation RDB\$DATABASE\_USERS has not been deleted

%SQL-F-REL\_EXISTS, Table RDB\$DATABASE\_USERS already exists in this database or schema

Creating RDB\$STORAGE\_AREAS

Creating RDB\$JOURNALS

Creating RDB\$LOGICAL\_AREAS

Creating RDB\$CHARACTER\_SETS

Creating RDB\$NLS\_CHARACTER\_SETS

Type COMMIT if there were no unexpected errors, otherwise ROLLBACK

```
SQL> rollback;
```

### Usage Notes

- The INFO\_TABLES script attempts to define a special domain RDB\$FILE\_SPECIFICATION. If this domain already exists in the database, then it is reasonable to see this message and it can be ignored.
- If one of the existing information tables is referenced by a view or stored procedure, then SQL will report the dependency on that object. Assuming that the sources for that view exist for recreation in the database, the database administrator can manually drop the table (RDB\$DATABASE\_USERS in this example) using the CASCADE option. Once INFO\_TABLES.SQL is re-executed and committed then the views can be recreated. Stored procedures are not dropped by the CASCADE option so will execute correctly when the information table is recreated.

## 8.2.3 Unexpected Bugcheck When Processing Zero Length Strings

Bug 6834122

In prior versions of Oracle Rdb, generation of zero length string literals might cause a bugcheck, as shown in the following example.

```
SQL> create table test_char(fld1 char(10));
SQL> insert into test_char values ('|||');
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER1:[TESTING]SQLBUGCHK.DMP;
%SQL-F-BUGCHK, There has been a fatal error. Please contact your Oracle
support representative. SQL$BLRXPR - 30
```

This problem has been corrected in Oracle Rdb Release 7.2.3. Oracle Rdb now correctly handles this case. As a workaround, this query will succeed if the dialect is changed to either ORACLE LEVEL1 or ORACLE LEVEL2.

## 8.2.4 Unexpected Bugcheck from RENAME TABLE Statement

In prior releases of Oracle Rdb, the RENAME TABLE command would generate a bugcheck if the target table was a GLOBAL or LOCAL TEMPORARY table. The command was attempting to modify the Rdb AIP structure which is not used by temporary tables.

The following shows an example of this problem.

```
SQL> create global temporary table SAMPLE
cont> (a integer)
cont> on commit preserve rows;
SQL> rename table SAMPLE to SAMPLING;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TESTING]RDSBUGCHK.DMP;
```

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.2.5 Unexpected RDMS-F-NODBK Reported by DROP USER Statement

Bug 6959409

In prior releases of Oracle Rdb, the DROP USER statement could report a RDMS-F-NODBK when processing a database defined with SECURITY CHECKING IS INTERNAL. In this type of database, the DROP USER also implicitly performs REVOKE ENTRY on each access control list (ACL) defined for all objects in the database. During this action, it is possible that the database key for the access control list will not be saved correctly.

The following example shows the type of error reported.

```
SQL> drop user rdbtest restrict;
%RDB-E-NO_RECORD, access by dbkey failed because dbkey is no longer associated
with a record
-RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-NODBK, 0:67108864:0 does not point to a data record
SQL>
```

It is possible that defining RDM\$BIND\_BUFFERS to a sufficiently large value may avoid this problem. This

problem has been corrected in Oracle Rdb Release 7.2.3. Oracle Rdb has been corrected to save the reference to the LIST OF BYTE VARYING access control list.

## 8.2.6 Unexpected COSI-F-VASFULL Reported by Long Running SQL Application

Bug 6952372

In some rare cases, SQL precompiler or SQL Module Language applications may leak memory when fetching numeric values into host variables defined as character types (for instance char in C, or PIC X in COBOL).

The following example shows an extreme case where a long running application aborted because of insufficient virtual address space.

```
$ R RDB_V_TEST3
%RDB-E-ARITH_EXCEPT, truncation of a numeric value at runtime
-RDB-E-ARITH_EXCEPT, truncation of a numeric value at runtime
-COSI-F-VASFULL, virtual address space full
-SYSTEM-F-ILLPAGCNT, illegal page count parameter
%SYSTEM-E-ABORT, abort
$ exit
```

This problem has been corrected in Oracle Rdb Release 7.2.3. This problem was caused by an incorrectly initialized descriptor that contained an invalid character set identification.

Oracle recommends that any affected application be rebuilt once this release is installed. This rebuild will require recompile with either the SQL precompiler or the SQL module language compilers.

In addition to this correction, some performance work was implemented in this area that will benefit all SQL precompiler or SQL Module Language applications.

## 8.2.7 Diagnostics Lost Following Procedure Call in Dynamic SQL

Bug 6379149

In prior versions of Oracle Rdb, it was possible that an exception raised by a stored procedure or a multi-statement procedure could be lost when executed from Dynamic SQL. The procedure must issue a COMMIT or ROLLBACK prior to raising the exception.

The following example shows that the signalled error (in this case from an explicit SIGNAL statement) is not reported through Dynamic SQL. The reported case was missing an exception reporting a failed constraint validation. Note: the output in this example is from a Dynamic SQL testing tool.

```
-> ATTACH 'filename sql$database';
-> SET DIALECT 'oracle level2';
-> SET HOLD CURSORS 'all';
-> SET TRANSACTION READ WRITE;
-> ! OCI ALWAYS DECLARES CURSORS AS WITH HOLD PRESERVE ALL, SEE SET HOLD
CURSORS ABOVE;
-> DYN P SELECT RDB$RELATION_NAME FROM RDB$RELATIONS;
```

```

Embedded DYN PREP executed (Implicit).
-> DYN O;
Embedded DYN OPEN executed (Implicit).
-> DYN F;
  0/RDB$RELATION_NAME: CANDIDATES
Embedded DYN FETCH executed (Implicit).
-> \
Entering block mode...
-> BEGIN
-> COMMIT;
-> SIGNAL 'RR002' ('unexpected error');
-> END;
-> \
Leaving block mode...
-> COMMIT;
-> EXIT;

```

For this erroneous behavior to be observed, the following must be true:

- A cursor must be open at the time that the procedure (stored or multi-statement) is executed.
- The procedure must execute a COMMIT or ROLLBACK statement.
- The cursor must be a WITH HOLD cursor otherwise the COMMIT or ROLLBACK will implicitly close any open cursors. This can be achieved using the WITH HOLD clause on DECLARE CURSOR or by executing the command SET HOLD CURSORS 'all'. Cursors open through OCI Services for Rdb are implicitly WITH HOLD cursors.

The returned diagnostics were lost when SQL checked that all non-HOLD cursors were closed following a transaction state change made by the procedure call. This problem has been corrected in Oracle Rdb Release 7.2.3. The returned diagnostics are now preserved and reported.

## 8.2.8 Unexpected Bugcheck from SET TRANSACTION Statement

Bug 7010980

In rare cases, Oracle Rdb would generate a bugcheck dump during a SET TRANSACTION statement that included a PARTITION list as part of the RESERVING clause.

The following example shows this problem.

```

SQL> set transaction
cont>   read only
cont>   wait 0
cont>   reserving
cont>     table1 for shared read,
cont>     table2 for shared read,
cont>     table3 for shared read,
cont>     table70 for shared read,
cont>     table83 for shared read,
cont>     table101 for shared read,
cont>     table367 partition(15) for shared read;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TESTING]RDSBUGCHK.DMP;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TESTING]SQLBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
virtual address=00000000225E7DC, PC=0000000001F48DC, PS=0000001B

```

## Oracle® Rdb for OpenVMS

This problem only occurs when the database contains many tables, indices and partitions and one of the tables referenced in the SET TRANSACTION statement has a logical area much larger than that referenced in the PARTITION clause.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.3 RMU Errors Fixed

### 8.3.1 RMU Extract Reports BAD\_CODE Error for TRANSLATE Function

Bug 6752325

In prior releases of Oracle Rdb, RMU Extract would generate a BAD\_CODE error when trying to extract the TRANSLATE function. The following example shows the reported error.

```
%RMU-F-BLRINV, internal error - BLR string 83 for SAMPLE_VIEW. is invalid
-COSI-E-BAD_CODE, corruption in the query string
%RMU-F-FTL_RMU, Fatal error for RMU operation at 16-JAN-2008 06:52:12.81
```

This problem has been corrected in Oracle Rdb Release 7.2.3.

### 8.3.2 Restore of Selected Area Fails With ACCVIO

A restore using a journal file to restore one or more areas from a multi-volume or multi-file backup fails with an ACCVIO.

```
$ RMU/RESTORE/AREA/JOURNAL=JF/REWIND/VOLUME=3 -
    TAPE1:MFP/MASTER,TAPE2:/MASTER,TAPE3:/MASTER DEPARTMENTS
...
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
    virtual address=0000000048, PC=...
```

Using the journal file, RMU realized the storage area DEPARTMENTS was located on the first tape volume. But remaining code still tried to access the other tapes without a proper parameter block.

This problem has been corrected in Oracle Rdb Release 7.2.3. The code now does not try to access an input volume or input file that is not needed.

### 8.3.3 RMU Load Quietly Truncated String Data During Insert

Bug 6732438

In prior versions of Oracle Rdb, the RMU Load command would quietly truncate data loaded into CHAR and VARCHAR columns. This loss of data might be significant but was never reported by Oracle Rdb.

This problem has been corrected in Oracle Rdb Release 7.2.3. RMU Load now defaults to SQL dialect SQL99 which implicitly checks for and reports truncations during INSERT operations.

This behavior is controlled by a new /DIALECT qualifier.

- /NODIALECT, /DIALECT=SQL89 or /DIALECT=NONE will revert to prior behavior and no truncation error will be reported.



- /DIALECT=SQL99 (the default) will enable reporting of truncation errors. Note that truncation occurs if non-space characters are discarded during the insert.

```
$ rmu/load abc f2 f1
%RMU-I-LOADERR, Error loading row 1.
%RDB-E-TRUN_STORE, string truncated during assignment to a column
%RMU-I-DATRECREAD, 1 data records read from input file.
%RMU-I-DATRECSTO, 0 data records stored.
%RMU-F-FTL_LOAD, Fatal error for LOAD operation at 13-FEB-2008 15:39:44.40
$
```

### 8.3.4 RMU /UNLOAD /AFTER\_JOURNAL Sort Performance

The RMU /UNLOAD /AFTER\_JOURNAL command sorts all records for each transaction in order to remove duplicate modifications to the same record within a transaction. Previous versions of the RMU /UNLOAD /AFTER\_JOURNAL command used the "SORT32" package and an internal sort algorithm to perform this sorting. However, for certain numbers of records being sorted, the internal sort algorithm could consume an unexpectedly large amount of CPU time.

The RMU /UNLOAD /AFTER\_JOURNAL command has been improved with a significantly faster internal sort algorithm. This should result in significant performance improvements for certain cases of extracting data from the after image journal file.

### 8.3.5 Possible RMU/MOVE\_AREA/ONLINE Database Corruption

Bug 6664950

RMU/MOVE/ONLINE could fail to correctly move Oracle Rdb database storage areas if other users were modifying the same areas. The Move operation could complete without any errors. Or, the Move operation would fail to delete the original area file(s) that had been moved because another database user incorrectly still had access to them.

```
%RMU-F-CANTDELETE, error deleting "DISK:[DIRECTORY]AREA.RDA;1"
%COSI-E-FLK, file currently locked by another user
-RMS-E-FLK, file currently locked by another user
%RMU-F-CANTDELETE, error deleting "DISK:[DIRECTORY]AREA.SNP;1"
%COSI-E-FLK, file currently locked by another user
-RMS-E-FLK, file currently locked by another user
```

An RMU/VERIFY of the database done after the move would show one of the following errors:

The first type of verify error reports that rows have been inserted in the storage area moved that were not in an index.

```
RMU/VERIFY/ALL DATABASE.RDB
%RMU-W-DATNOTIDX, Row in table ACCOUNT is not in any indexes.
Logical dbkey is 68:270:2.
%RMU-W-DATNOTIDX, Row in table ACCOUNT is not in any indexes.
Logical dbkey is 68:343:1.
```

## Oracle® Rdb for OpenVMS

The second type of verify error reports that the maximum page number in the database root for the storage area snapshot file moved is invalid.

```
RMU/VERIFY/ALL DATABASE.RDB
%RMU-E-BADMAXPNO, unable to read last page (1478) in file
DISK:[DIRECTORY.ONE]AREA.SNP;1
%RMU-F-ABORTVER, fatal error encountered; aborting verification
```

If a previous version of Rdb is installed which does not include the solution to this problem, it is highly recommended not to use online moves while the storage area being moved is being modified.

The following example of this problem shows an online move of the area HASHES of the database MOVIE from DISK:[DIRECTORY] to DISK:[DIRECTORY.ONE]. At the end of the move, the %RMU-F-CANTDELETE error was output when RMU attempted to delete the original source storage area files that had been moved. This occurred because another user still had access to these files because of this problem.

```
$ rmu/move/online/log/directory=[.one] movie.rdb hashes
%RMU-I-QUIETPT, waiting for database quiet point at 26-NOV-2007
10:44:23.86
%RMU-I-RELQUIETPT, Database quiet point lock has been released at
26-NOV-2007 10:44:33.39.
%RMU-I-MOVTEXT_04, Starting move of storage area
DISK:[DIRECTORY.ONE]HASHES.RDA;1 at 26-NOV-2007 10:44:35.68
%RMU-I-MOVTEXT_01, Completed move of storage area
DISK:[DIRECTORY.ONE]HASHES.RDA;1 at 26-NOV-2007 10:44:40.44
%RMU-I-MOVTEXT_05, Moved snapshot area file
DISK:[DIRECTORY.ONE]HASHES.SNP;1
%RMU-I-MOVTEXT_15, Live area files moved
%RMU-I-MOVTEXT_10, Root entry modified for moved area file
DISK:[DIRECTORY.ONE]HASHES.RDA;1
%RMU-I-MOVTEXT_11, Root entry modified for moved snapshot file
DISK:[DIRECTORY.ONE]HASHES.SNP;1
%RMU-I-MOVTEXT_06, Database root updated for all moved areas
%RMU-I-MOVTEXT_16, On line users notified of completion of move
%RMU-F-CANTDELETE, error deleting "DISK:[DIRECTORY]HASHES.RDA;1"
%COSI-E-FLK, file currently locked by another user
-RMS-E-FLK, file currently locked by another user
%RMU-F-CANTDELETE, error deleting "DISK:[DIRECTORY]HASHES.SNP;1"
%COSI-E-FLK, file currently locked by another user
-RMS-E-FLK, file currently locked by another user
%RMU-I-MOVTEXT_07, Obsolete files deleted
%RMU-W-DOFULLBCK, full database backup should be done to ensure future
recovery
%RMU-I-COMPLETED, MOVE_AREA operation completed at 26-NOV-2007
10:44:41.80
```

A workaround to avoid this problem is to not specify /ONLINE for the move command. If /ONLINE is specified, always verify the database after the move. If the verify shows the index diagnostics, the indexes should be rebuilt. If the verify shows the root diagnostics, the database can be restored and the move repeated without specifying /ONLINE.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.3.6 RMU /BACKUP Performance Enhanced On Some Types Of Tape Drives

Performance testing of the RMU/BACKUP utility using LTO4 tape drives indicated that some magnetic tape devices experienced unexpectedly long delays while repositioning after writing tape marks. This delay was found to have been the result of not using the "fast" method of skipping over tape marks. In some situations, this delay was found to be several minutes long.

This problem has been corrected in Oracle Rdb Release 7.2.3. Performance of the RMU/BACKUP utility has been improved for magnetic tape devices that support the SCSI space-by-file-marks command and the SCSI read position command.

As a workaround, the DCL command "SET MAGTAPE/FAST\_SKIP=ALWAYS dev" can be used to set the tape drive to use the "fast" method of skipping over tape marks by default.

## 8.3.7 RMU Analyze Reports Incorrect Compression Ratio for Some Tables

Bug 6992051

Starting with Oracle Rdb V7.1, the run-length compression algorithm was changed such that very long strings are compressed using a more compact encoding scheme. Unfortunately, the RMU/ANALYZE command did not use the new algorithm and so in some cases miscalculated the compression ratio for some table logical areas.

The following example shows a fragment of output from the RMU Analyze command.

```
Logical area: COL1 for storage area : BIGAREA
Larea id: 57, Record type: 30, Record length: 3608, Compressed

Data records: 1, bytes used: 68 (0%)
  average length: 68, compression ratio: 0.45
```

In this example, the stored rows were 3603 bytes long (record length minus the row overhead) and therefore the compression ratio should have been about 0.02.

This problem has been corrected in Oracle Rdb Release 7.2.3. RMU Analyze now uses the correct algorithm to compute the expanded row length and to compute a more accurate compression ratio.

## 8.3.8 Restore May Miss Storage Area Without Warning

A restore of a full multi-volume or multi-file backup may complete without a warning even though not all savesets have been specified in the restore command. Hence storage areas saved in the missing saveset(s) are not restored.

A RMU/VERIFY of the restored database does not always detect the error.

```
$ RMU/BACKUP/DISK=WRITER_THREADS:2 MF_PERSONNEL MF_PERSONNEL.RBF, [ ]
$ SQL DROP DATABASE FILE MF_PERSONNEL;
```

```

$ !
$ ! Restore command with missing 2nd save set
$ !
$ RMU/RESTORE/NOCCD/DISK MF_PERSONNEL.RBF
%RMU-I-AIJRSTAVL, 0 after-image journals available for use
%RMU-I-AIJISOFF, after-image journaling has been disabled
%RMU-W-USERECCOM, Use the RMU Recover command. The journals are not available.
$ RMU/VERIFY/ALL/NOLOG MF_PERSONNEL
%RMU-W-BADPROID, DKA100:[TEST]EMPIDS_LOW.RDA;1 file contains a bad identifier
Expected "RDMSDATA", found ""
%RMU-W-BADDBPRO, DKA100:[TEST]EMPIDS_LOW.RDA;1 file does not belong to
database DKA100:[TEST]MF_PERSONNEL.RDB;1
found references to database
...
$ !
$ ! Correct restore command
$ !
$ RMU/BACKUP/DISK=READER_THREADS:2 MF_PERSONNEL MF_PERSONNEL.RBF,[]

```

A check has been added to RMU/RESTORE which issues a warning message after the completion of the restore if not all areas have been restored.

```
%RMU-W-RESINCOMP, Not all storage areas have been restored
```

Explanation:

One or more storage areas could not be restored.

User Action:

Use RMU/VERIFY to find any missing or incomplete storage area. Repeat the restore command and include the save sets with the missing storage areas.

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.3.9 Incorrect View Syntax Generated by RMU Extract Item=VIEW

Bug 6905847

Prior releases of Oracle Rdb would not correctly extract views generated by third-party database tools. The following example shows the incorrect syntax in the SELECT value expression list.

```

$ rmu/extract-
  /item=view-
  /option=(noheader,mat:V2%,filename_only) abc
set verify;
set language ENGLISH;
set default date format 'SQL92';
set quoting rules 'SQL92';
set date format DATE 001, TIME 001;
attach 'filename ABC';
create view "V2"
  (A,
   B,
   C,
   D,
   E) as
(select

```

```

      ,
      ,
      C1.C,
      C1.D,
      C1.E
from ABCD C1);

```

```
commit work;
```

This problem has been corrected in Oracle Rdb Release 7.2.3. RMU Extract now handles the variations in these view definitions and generates a valid SQL view definition.

### 8.3.10 Revised Format for the RMU/UNLOAD/AFTER\_IMAGE CONTROL File

In prior releases of Oracle Rdb, the command RMU/UNLOAD/AFTER\_IMAGE permitted the unloaded after-image table data to be loaded into an Oracle RDBMS database using the SQL\*Loader tool. This was enabled by generating a control file for SQL\*Loader using the CONTROL qualifier and outputting the data in a portable TEXT format. There were several problems with this support that have been corrected in this release.

- The LogMiner timestamp information was described as Oracle DATE type which does not support fractional seconds precision. Now the fields RDB\$LM\_START\_TAD and RDB\$LM\_COMMIT\_TAD are defined as TIMESTAMP types with fractional seconds precision.
- The syntax generated for the control file was incorrect in several places. This included specifying 12 hour date/time masks instead of 24 hour masks for DATE, incorrect type specifiers for numeric values, and incorrect offset values.
- For some data types, the unloaded TEXT format was too narrow allowing possible truncation of the data and for some too wide, wasting space in the unloaded data file. This has been corrected for both LogMiner and RMU/UNLOAD for FORMAT=TEXT.

This problem has been corrected in Oracle Rdb Release 7.2.3. At this time only partial support for TIME and INTERVAL is available in the RMU/UNLOAD/AFTER\_IMAGE CONTROL file.

### 8.3.11 RMU Created an Incorrect Restore Options File for Single File Databases

Bug 7033861

The Oracle Rdb RMU commands RMU/DUMP/RESTORE\_OPTIONS, RMU/DUMP/BACKUP/RESTORE\_OPTIONS and RMU/BACKUP/RESTORE\_OPTIONS created an incorrect options file to be used by the RMU/RESTORE/OPTIONS command. The file specification of the database in the first comment line was missing. The "/FILE" qualifier, illegal for single file databases which do not have an \*.RDA storage area file, was included and was not followed by a file name if the database had not yet been backed up. The /SNAPSHOT qualifier "FILE=" parameter was also included when it should not have been and was also not followed by a file specification if the database had not yet been backed up. The /SNAPSHOT qualifier "FILE=" parameter in an options file for a single file database cannot be used to change either the location or name of the snapshot file.

## Oracle® Rdb for OpenVMS

For an Oracle Rdb single file database, the single storage area is located in the \*.RDB file along with the database root. There is only one snapshot file which has the same name as the database \*.RDB file. For example, for the single file personnel database, the single storage area is located in the PERSONNEL.RDB file and the single snapshot file is named PERSONNEL.SNP. The option file used by the restore command can be used to change other parameters of the single storage area but not the \*.RDB and \*.SNP file names and locations.

The following example of this problem shows that the REST.OPT options file created by the RMU/DUMP command had incorrect syntax which caused an error when the RMU/RESTORE/OPTION=REST.OPT command was executed.

```
$ RMU/DUMP/OUT=NL:/RESTORE_OPT=REST.OPT PERSONNEL.RDB
$ TYPE REST.OPT
! Options file for database
! Created dd-mmm-yyyy hh:mm:ss.xxxx
! Created by DUMP command

RDB$SYSTEM -
    /file= -
    /extension=ENABLED -
    /read_write -
    /spams -
    /snapshot=(allocation=100, -
                file=)
$ RMU/RESTORE/NOLOG/OPTION=REST.OPT/NOCCD PERS.RBF
%RMU-F-CONFLSWIT, conflicting qualifiers /FILE and single file
database organization
%RMU-F-FTL_RSTR, Fatal error for RESTORE operation at dd-mmm-yyyy hh:mm:ss.xxxx
```

A workaround to avoid this problem is to edit the options file to eliminate the "/file=" qualifier and the "file=" parameter that follows the "/snapshot" qualifier and add the database file specification to the end of the first comment line. Here is the correct restore options file format for a single file database that the RMU commands named above will output now that this problem has been corrected.

```
$ RMU/DUMP/OUT=NL:/RESTORE_OPT=REST.OPT PESONNEL.RDB
$ TYPE REST.OPT
! Options file for database SERDB_USER1:[HOCHULI.BUG7033861]PERSONNEL.RDB;1
! Created 25-JUN-2008 14:51:31.90
! Created by DUMP command

RDB$SYSTEM -
    /extension=ENABLED -
    /read_write -
    /spams -
    /snapshot=(allocation=100)
$ RMU/RESTORE/NOLOG/OPTION=REST.OPT/NOCCD PERS.RBF
%RMU-I-AIJRSTAVL, 0 after-image journals available for use
%RMU-I-AIJISOFF, after-image journaling has been disabled
%RMU-W-USERECCOM, Use the RMU Recover command. The journals are not
available.
```

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 8.4 RMU Show Statistics Errors Fixed

### 8.4.1 RMU/SHOW STATISTICS Checkpoint Information Screen Sort By Oldest Transaction Did Not Work Correctly

Bug 6735923

In previous Rdb 7.2 releases, the RMU /SHOW STATISTICS "Checkpoint Information" screen option to "Sort by oldest transaction" option did not work properly. Users were unable to order the display by transaction start time.

This problem has been corrected in Oracle Rdb Release 7.2.3. The "Sort by oldest transaction" option now works as expected.

### 8.4.2 Date and Time Included In Lock Deadlock Log File

Bug 4069186

Previously, only the time of a deadlock was included in the "Lock Deadlock Log" file. In some cases, the lack of the date field made analysis difficult.

This problem has been corrected in Oracle Rdb Release 7.2.3. The "Lock Deadlock Log" file now always includes a date field as demonstrated in this example output.

```
Oracle Rdb V7.2-200 Performance Monitor Lock Deadlock Log
Database DPA10:[DB]AARCH.RDB;1
Lock Deadlock Log created 20-MAR-2008 12:18:41.18
2040045C:1 20-MAR-2008 18:11:20.3145724 - waiting for page 18:8 (PW) [6 missed]
20400DA3:1 21-MAR-2008 08:00:01.9816693 - waiting for page 53:17111 (PW)
```

### 8.4.3 Incorrect RMU/SHOW STATISTICS Layout File RMU\$SHOW\_STATISTICS72.CDO

Bug 7156903

The file RMU\$SHOW\_STATISTICS72.CDO that shipped with prior versions of Oracle Rdb V7.2 did not correctly describe the layout of the RMU/SHOW STATISTICS data file.

This problem has been corrected in Oracle Rdb Release 7.2.3. A revised RMU\$SHOW\_STATISTICS72.CDO file is provided by this release. In addition two companion files are provided in SYSS\$LIBRARY: RMU\$SHOW\_STATISTICS72.SQL, which creates a SQL table for use by RMU/LOAD; and RMU\$SHOW\_STATISTICS72.H, a C header file to simplify processing of the data file by C applications.

---

# **Chapter 9**

## **Enhancements And Changes Provided in Oracle Rdb Release 7.2.4**



# 9.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.4

## 9.1.1 Date/Time Arithmetic Enhancements

Bug 6219485

This release of Oracle Rdb lifts many restrictions on the DATE VMS data type. SQL now treats DATE VMS type as a TIMESTAMP(2) for the purposes of the add and subtract with intervals or when generating intervals.

- A year–month interval can be added to or subtracted from a DATE VMS value and result in a DATE VMS value.
- A day–time interval can be added to or subtracted from a DATE VMS value and result in a DATE VMS value.
- A DATE VMS value can be subtracted from another DATE VMS value to produce either a year–month interval or a day–time interval.
- A DATE ANSI value can be subtracted from a DATE VMS value to produce a year–month interval.
- A DATE VMS value can be subtracted from a DATE ANSI value to produce either a year–month interval or a day–time interval.
- A DATE VMS value can be subtracted from a TIMESTAMP value to produce either a year–month interval or a day–time interval.
- A TIMESTAMP value can be subtracted from a DATE VMS value to produce either a year–month interval or a day–time interval.

Also in this release the following rules have been relaxed.

- Relax rule that TIME values could only be subtracted if they had the same fractional seconds precision.
- Relax rules that TIMESTAMP values could only be subtracted if they had the same fractional seconds precision.
- Relax assignment rules and let Rdb handle truncating time portion when TIMESTAMP is assigned to a DATE ANSI column, variable or parameter.
- Add rule that merge of DATE VMS and TIMESTAMP(n) will always be TIMESTAMP(n), where n is inherited from the TIEMSTAMP expression. Merge rules are used by UNION, INTERSECT, EXCEPT, MINUS operators and CASE expression processing.


These enhancements have been made in Oracle Rdb Release 7.2.4.

## 9.1.2 New DEFAULT PROFILE Feature

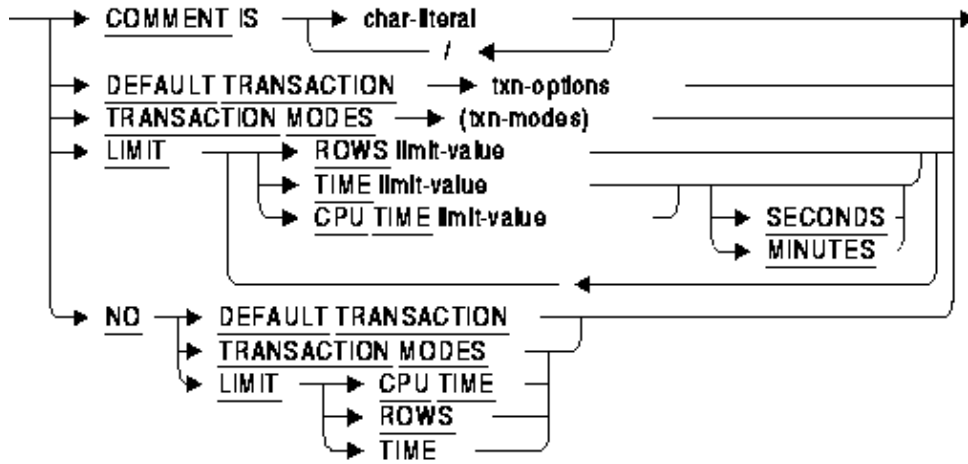
This release of Oracle Rdb enhances the PROFILE support with a new DEFAULT profile. When a user attaches to the database using ATTACH, CONNECT or SET SESSION AUTHORIZATION, they will either load their assigned profile definition or inherit the default profile (if defined).

*Syntax*



The CREATE, ALTER and DROP PROFILE syntax is changed as shown. The existing profile-options diagram remains the same.



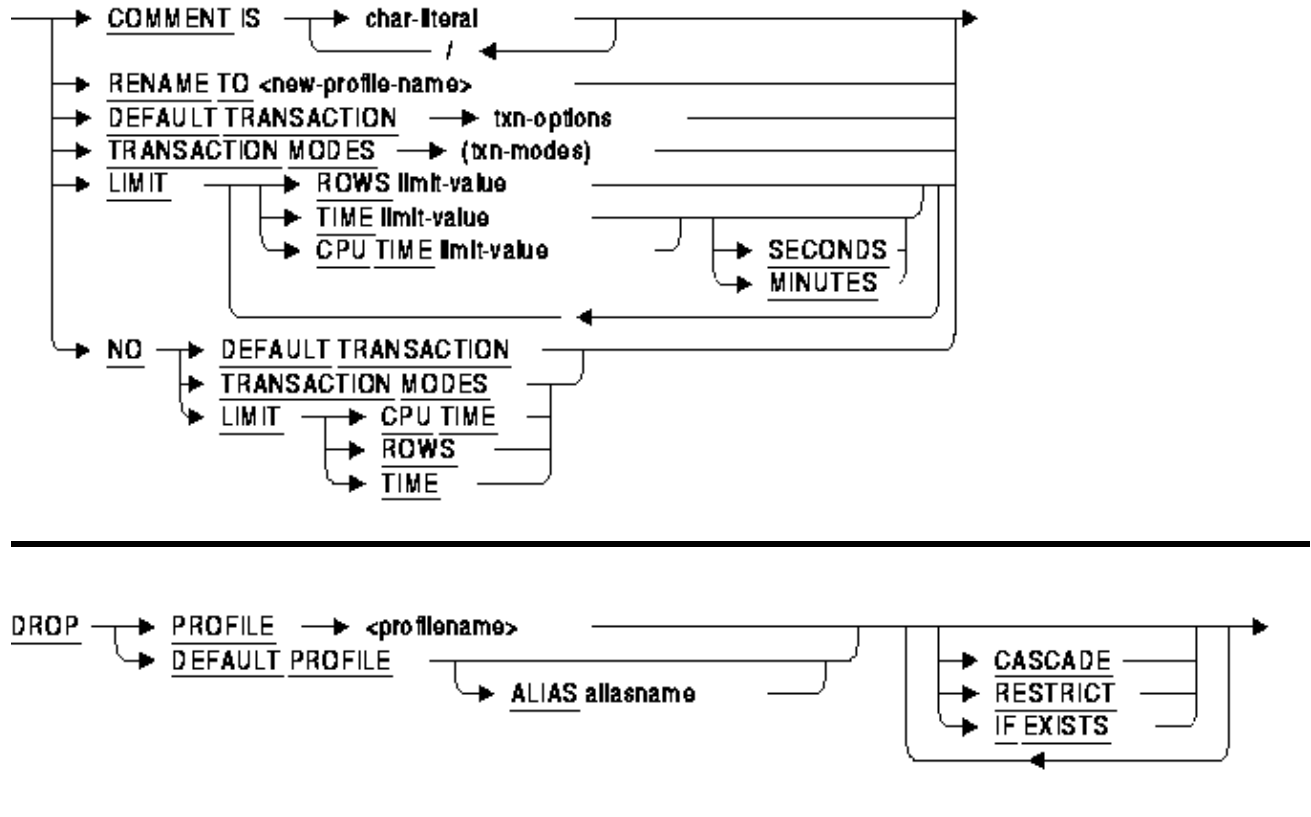
profile-options =



limit-value =



profile-options =



**Arguments**

- **ALIAS aliasname**  
When attached to multiple databases, the aliasname is required to direct the CREATE, ALTER or DROP command to the appropriate database.
- **DEFAULT PROFILE**  
Creates the special profile RDB\$DEFAULT\_PROFILE. This profile will be used by any user who is not assigned a profile using the PROFILE clause of CREATE or ALTER PROFILE.

**Usage Notes**

- It is possible to restrict the transaction modes to READ ONLY using the default profile. Use caution in this case because it is possible that no user will have READ WRITE access to undo such a definition. In this case, you can define the logical name RDBMS\$SET\_FLAGS to the value PROFILE\_OVERRIDE to allow a suitably privileged user to start a transaction without using the transaction mode restrictions in the default profile. Such a user must have database SECURITY privilege, possibly inherited from the OpenVMS SECURITY process privilege.

## 9.1.3 RMU /DUMP/BACKUP/OPTIONS=ROOT /HEADER\_ONLY Displays the Header Information Only

Bug 8235615

A new feature has been added to RMU/DUMP/BACKUP/OPTIONS=ROOT command to process only the header information when the /HEADER\_ONLY qualifier is used.

In prior releases of Oracle Rdb, the user had to wait until the entire backup file (.RBF) was processed. If the backup file was stored on tape and spanned multiple tapes then all the tapes had to be mounted and processed. When using /HEADER\_ONLY, RMU now ceases processing of the backup file once the header has been displayed.

```
$ RMU/DUMP /BACKUP/OPTIONS=ROOT /HEADER_ONLY DBNODE$DMA200:GLORY.RBF
*-----*
* Oracle Rdb V7.2-4                               26-FEB-2009 07:21:58.69
*
* Dump of Database Backup Header
*   Backup filename: GLORY.RBF
*   Backup file database version: 7.2
*
*-----*
```

### Database Parameters:

```
Root filename is "USER1:[BUG.8235615.FIX]GLORY.RDB;1"
Created at 25-APR-2007 16:43:05.52
Oracle Rdb structure level is 72.1
Maximum user count is 23
Maximum node count is 3
Database open mode is AUTOMATIC
Database close mode is AUTOMATIC
Database will be mapped in process space
All transaction modes are allowed
Prestarted transactions are enabled
Snapshot mode is NON-DEFERRED
Statistics are enabled
Operator notification is disabled
Logical area count is 512
Storage Areas...
  - Active storage area count is 4
  - Reserved storage area count is 8
Row Caches...
  - Active row cache count is 0
  - Reserved row cache count is 1
  - Checkpoint information
    No time interval is specified
    Default source is updated rows
    Default target is backing file
    Default backing file directory is database directory
    RUJ Global Buffers are disabled
  - WARNING: Maximum node count is 3 instead of 1
  - WARNING: After-image journaling is disabled
  - WARNING: Fast commit is disabled
Buffers...
  - Default user buffer count is 2000
  - Default recovery buffer count is 2000 (stored as 20)
  - Global buffers are disabled
    Global buffer count is 115
```

## Oracle® Rdb for OpenVMS

Maximum global buffer count per user is 5  
Large memory is disabled

- Buffer size is 12 blocks
  - Maximum pages per buffer is 6
- Asynchronous pre-fetch is enabled
  - Maximum pre-fetch depth is 8 buffers
- Detected asynchronous pre-fetch is enabled
  - Maximum pre-fetch depth is 4 buffers
  - Pre-fetch threshold is 4 buffers
- Asynchronous batch-write is enabled
  - Clean buffer count is 5
  - Maximum batch size is 10 buffers
- Optimized page transfer is disabled

Locking...

- Adjustable record locking is enabled
  - Fanout factor 1 is 10 (10 pages)
  - Fanout factor 2 is 10 (100 pages)
  - Fanout factor 3 is 10 (1000 pages)
- Carry-over lock optimization is enabled
- Lock tree partitioning is disabled

RUJ Journaling...

- No default recovery-unit journal directory

AIJ Journaling...

- After-image journaling is disabled
- Database is configured for 7 journals
- Reserved journal count is 7
- Available journal count is 0
- LogMiner is disabled
- 7 journals can be created while database is active
- Shutdown time is 60 minutes
- Backup operation is manual
- Default backup filename edits are not used
- Log server startup is MANUAL
- Journal overwrite is disabled
- AIJ cache on "electronic disk" is disabled
- Default journal allocation is 512 blocks
- Default journal extension is 512 blocks
- Default journal initialization is 512 blocks
- Current roll-forward sequence number is 0
- Current backup sequence number is 0

Fast Commit...

- Fast commit is disabled
- No checkpointing AIJ interval is specified
- No checkpointing time interval is specified
- No checkpointing transaction interval is specified
- Commit to AIJ optimization is disabled
- Transaction interval is 256

Hot Standby...

- WARNING: After-image journaling is disabled
- WARNING: Fast commit is disabled
- WARNING: Log server startup is MANUAL
- Informational: Operator notification is disabled
- Database is not currently being replicated

Security Auditing...

- Security auditing is disabled
- Security alarm is disabled
- No audit journal filename is specified
- No alarm name is specified
- Synchronous audit record flushing is disabled
- Audit every access

Database Backup...

- Fast incremental backup is enabled

## Oracle® Rdb for OpenVMS

```
- Last full database backup was on 26-FEB-2009 07:16:56.09
- Full database backup TSN is 0:128
- Database was restored on 26-FEB-2009 06:52:11.82
Derived Data...
- Global section size
  With global buffers disabled is 277187 bytes (1MB)
  With global buffers enabled is 1036584 bytes (1MB)
  With Large memory global buffers enabled...
    Database TROOT section is 330024 bytes (1MB)
    Large memory global buffers section is 706560 bytes (1MB)
- Row Cache RUJ buffers section size is 6041088 bytes (6MB)
```

Database root file ACL

```
( IDENTIFIER=[RDB,SFRANN],ACCESS=READ+WRITE+CONTROL+RMU$ALTER+RMU$ANALYZE+
  RMU$BACKUP+RMU$CONVERT+RMU$COPY+RMU$DUMP+RMU$LOAD+
  RMU$MOVE+RMU$OPEN+RMU$RESTORE+RMU$SECURITY+RMU$SHOW+RMU$UNLOAD+RMU$VERIFY)
```

### 9.1.4 GET ENVIRONMENT Now Supports SQLCODE and SQLSTATE Capture

This release of Oracle Rdb allows the GET ENVIRONMENT (SESSION) command to return the SQLCODE and SQLSTATE of the last executed statement. The keyword SQLCODE returns an INTEGER value and SQLSTATE returns a CHAR(5) value. The execution of the GET ENVIRONMENT statement will clear these values so both should be fetched in the same statement.

The following example shows a raised error and the use of GET ENVIRONMENT to capture the SQLCODE and SQLSTATE.

```
SQL> declare :st char(5);
SQL> declare :sc integer = -1;
SQL>
SQL> begin set :sc = :sc / 0; end;
%RDB-E-ARITH_EXCEPT, truncation of a numeric value at runtime
-COSI-F-ARITH, arithmetic exception
-COSI-F-FLTDIV, floating point division exception
SQL>
SQL> get environment (session) :st = SQLSTATE, :sc = SQLCODE;
SQL>
SQL> print :st, :sc;
  ST          SC
22003         -304
SQL>
```

Once the SQLCODE or SQLSTATE value has been saved to a declared variable, it can be used to conditionally execute a COMMIT or a ROLLBACK.

```
begin
if :sc < 0 then
  rollback;
else
  commit;
end if;
end;
```

## 9.1.5 Timestamp Added to Messages For RMU LOAD and UNLOAD

In order to help judge progress of RMU LOAD and UNLOAD operations, a timestamp has been added to the RMU-I-DATRECSTO and RMU-I-DATRECUNL progress messages. The following example shows these timestamps.

```
$RMU /UNLOAD MFP C1 C1
%RMU-I-DATRECUNL, 200000 data records unloaded 5-JUN-2009 07:58:17.23.
$RMU /LOAD /COMMIT=75000 /LOG_COMMIT MFP C1 C1
%RMU-I-DATRECSTO, 75000 data records stored 5-JUN-2009 07:58:43.09.
%RMU-I-DATRECSTO, 150000 data records stored 5-JUN-2009 07:58:43.12.
%RMU-I-DATRECSTO, 200000 data records stored 5-JUN-2009 07:58:43.14.
```

## 9.1.6 New SET SQLDA Statement

Bugs 1088554, 4179408, 5414051, and 7022262

This release of Oracle Rdb introduces a new SET SQLDA statement.

The SQLDA Statement allows a programmer using Dynamic SQL to alter the way the SQLDA (and SQLDA2) and Dynamic SQL statements are processed by Oracle Rdb.

### *Environment*

You can use the SET SQLDA statement:

- In dynamic SQL as a statement to be dynamically executed

### *Syntax*

---

```
SET SQLDA    ┌───▶ literal            ┐
              └───▶ host-variable     ┘
```


---

**sqlda\_options =**


```
┌───▶ sqlda_option            ┐
                                  └───▶
```

---


**sqlda\_option =**



**enable-option =**



**dialect-name =**



### Arguments

- literal  
host-variable  
The parameter passed to the statement must be a literal or a host variable containing one or more SQLDA options (see sqlda\_options syntax diagram for details). If more than one option is specified, they must be separated by commas.
- sqlda\_options  
One or more keyword clauses. If more than one clause is specified, they must be separated by commas.
- ENABLE  
The ENABLE clause activates one of the following behaviors for Dynamic SQL.
  - ◆ INSERT RETURNING – The default behavior of INSERT ... RETURNING when executed by dynamic SQL is to place parameters from the RETURNING INTO clause into the INPUT SQLDA. This behavior is maintained for backward compatibility. This option allows the programmer to force different (and correct) behavior for the non-compound use of this statement.

---

Note



***If the INSERT RETURNING statement is included in a compound statement, the parameters are handled correctly.***

---

- ◆ NAMED MARKERS – as well as traditional parameters markers (?). Dynamic SQL will now accept named, host variable style parameter markers. See the Usage Notes for further details and examples.
- ◆ ROWID TYPE – returns DBKEY values as a special type (SQLDA\_ROWID, 455) to make processing of the DBKEY values easier. For instance, in prior releases, the SQLDA name field (SQLNAME) for DBKEY entries in the SQLDA was the only way to distinguish these values from other CHAR or VARCHAR columns – it would be either DBKEY or ROWID. If a query renamed the DBKEY column, then the application had no information in the SQLDA to indicate that the CHAR or VARCHAR value was binary data. In all respects, the SQLDA\_ROWID type appears as a fixed length string of octets (possibly containing octets of zero which the C language would treat as a NULL terminator for a string).
- DISABLE  
The DISABLE clause deactivates one of the specified behaviors for Dynamic SQL. See ENABLE clause for a list of options.
- ORACLE LEVEL1  
ORACLE LEVEL2  
Either of these options will set the SQLDA to supply enhanced semantics. These options are currently reserved for use of the OCI Services for Rdb product that is part of the Oracle Rdb SQL/Services component. This setting also implicitly enables NAMED MARKERS.
- PADDING n CHARACTERS  
This option directs SQL to configure the SQLDA with larger CHARACTER VARYING strings than would normally be seen. The value of n is an unsigned numeric literal that specifies the number of characters that are added to the estimated length. Any CHARACTER (CHAR) types are converted to CHARACTER VARYING (VARCHAR). This rule is applied to comparison operators <, <=, >, >=, =, <>, and string functions (STARTING WITH, CONTAINING).
- NOPADDING  
This option sets the number of padding characters to 0. This also implies that derived CHARACTER (CHAR) types are not converted to CHARACTER VARYING (VARCHAR) when PADDING CHARACTERS is used.

---

Note

***Oracle recommends that applications always check for SQLDA\_CHAR and SQLDA\_VARCHAR so that the correctly formatted data is made available to SQL.***

---

This is the default setting.

- SQL99
  - SQL92
  - MIA
  - SQL89
  - SQLV40
- Any of these options will revert to the default semantic for the SQLDA which includes disabling NAMED MARKERS.

### ***Usage Notes***

- The ORACLE LEVEL1 and ORACLE LEVEL2 settings are reserved for use by Oracle Corporation. Current behavior of this setting may change with any given release based on requirements of the OCI Services for Rdb component. This setting changes the usage of various SQLDA and SQLDA2 fields.
- Keywords may not be abbreviated and the clauses must be fully specified.
- The SET DIALECT command will implicitly enable NAMED MARKERS if the dialect is changed to either ORACLE LEVEL1 or ORACLE LEVEL2.
- The SET DIALECT command will implicitly disable NAMED MARKERS if the dialect is changed to any dialect other than ORACLE LEVEL1 or ORACLE LEVEL2.
- When NAMED MARKERS are enabled, the contents of the SQLDA and SQLDA2 will reflect one entry for each name. When traditional parameter markers are used, a SQLDA (or SQLDA2) entry will exist for each marker (?) encountered. This change in behavior can simplify the query encoding as well lead to more efficient strategy creation.

### *Examples*

#### *Using the NAMED MARKERS Feature*

This example shows that enabling the NAMED MARKERS feature will allow SQL to prompt for one value and the displayed Rdb strategy shows that only one variable is used.

```
-> SET SQLDA 'ENABLE NAMED MARKERS';
-> SELECT LAST_NAME FROM EMPLOYEES WHERE FIRST_NAME = :F_NAME AND LAST_NAME <>
:F_NAME;
in:  [0] typ=449 len=46
out: [0] typ=453 len=14
[SQLDA - reading 1 fields]
-> Alvin
Tables:
  0 = EMPLOYEES
Conjunct: (0.FIRST_NAME = <var0> AND (0.LAST_NAME <> <var0>))
Get      Retrieval sequentially of relation 0:EMPLOYEES
  0/FIRST_NAME/Varchar(42/46): Alvin
[SQLDA - displaying 1 fields]
  0/LAST_NAME: Toliver
[SQLDA - displaying 1 fields]
  0/LAST_NAME: Dement
```

#### *Using the PADDING Feature*

The following example shows that the derived type for the named parameter MI is a SQLDA\_CHAR (453) of length 1. The input data ('AA') is truncated on assignment and the incorrect results are returned. By adding a small padding, the type is changed to SQLDA\_VARCHAR (449) of length 3 and a correct comparison is performed.

```
-> ATTACH 'filename sql$database';
-> SET SQLDA 'enable named markers, nopadding';
-> SELECT LAST_NAME FROM EMPLOYEES WHERE MIDDLE_INITIAL = :MI;
in:  [0] typ=453 len=1
out: [0] typ=449 len=18
[SQLDA - reading 1 fields]
-> AA
[SQLDA - displaying 1 fields]
  0/LAST_NAME: Toliver
[SQLDA - displaying 1 fields]
  0/LAST_NAME: Lengyel
[SQLDA - displaying 1 fields]
```

```

0/LAST_NAME: Robinson
[SQLDA - displaying 1 fields]
0/LAST_NAME: Ames
-> SET SQLDA 'padding 2 characters';
-> SELECT LAST_NAME FROM EMPLOYEES WHERE MIDDLE_INITIAL = :MI;
in: [0] typ=449 len=7
out: [0] typ=449 len=18
[SQLDA - reading 1 fields]
-> AA
-> EXIT;
Enter statement:

```

Note that the VARCHAR requires an extra 4 bytes for the length information in the SQLDA2 used by the Dynamic SQL testing program.

## 9.1.7 RMU /SHOW VERSION Displays System Architecture and Version

The RMU /SHOW VERSION command has been enhanced to include information about the system architecture and OpenVMS version as shown in the following example:

```

$ RMU /SHOW VERSION
Executing RMU for Oracle Rdb V7.2-400 on OpenVMS IA64 V8.3-1H1
$

```

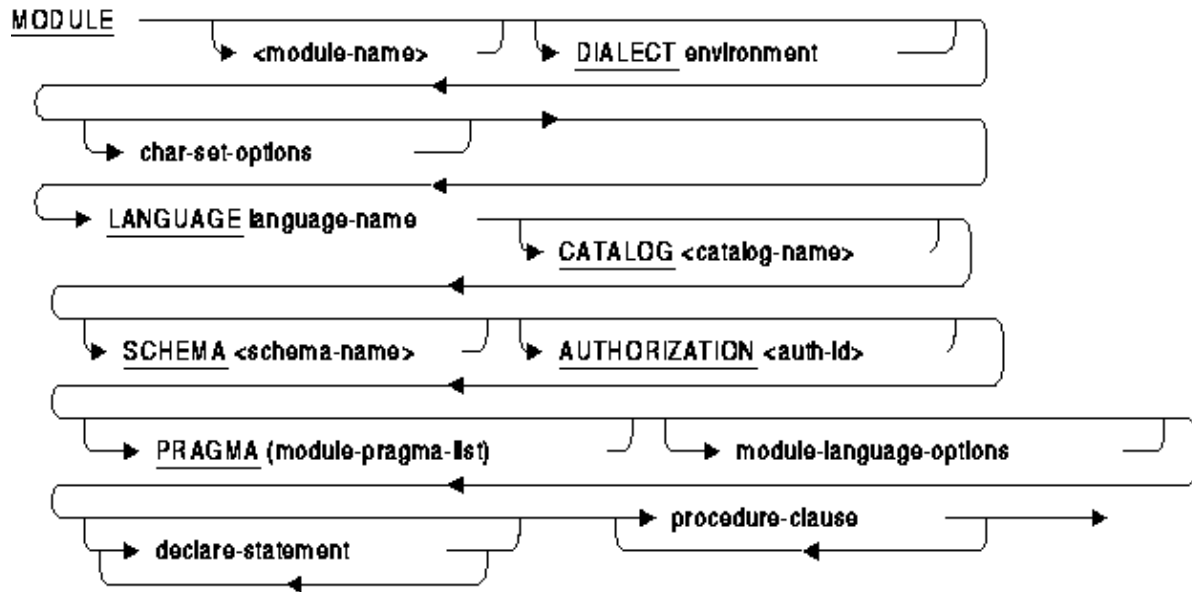
## 9.1.8 New IDENT Option for SQL Module Language PRAGMA Clause

Many OpenVMS compilers allow the programmer to specify the object file IDENT string. This allows tracking of the correct version in linker map files (.map) and within object libraries using the LIBRARIAN command. This release of Oracle Rdb supports setting the IDENT string for the SQL module language.


The IDENT string is specified as part of the PRAGMA clause in the module header.

### *Syntax*

---



module-pragma-list =



**Usage Notes**

- By using the PRAGMA clause with IDENT you can record an identification string in the object module generated by the SQL Module language. This IDENT string is recorded by the OpenVMS LINKER in the image itself and can be viewed in the generated MAP file, examined using ANALYZE/OBJECT, and by the LIBRARIAN command when the object module is stored in an object library.
- OpenVMS limits the IDENT string to a 15 octet string. If the string is longer than this (even with trailing spaces) then an error will be reported by the SQL Module Language compiler.
- If the IDENT clause is omitted, then the default version string will default to 'V1.0' as is the practice with many OpenVMS compilers. Prior versions of Oracle Rdb on Integrity systems would only provide the string '0'.

```

Module name:                "MODSQL$TEST"
Module version:             "0"
Creation date/time:         "15-JUN-2009 20:13"
Language name:              "Oracle Rdb SQL V7.2-351"
    
```

**Examples**

The following example shows the use of a PRAGMA clause in a module header to specify the module ident string.

**Example 9–1 PRAGMA Clause in the Module Header**

---

```

MODULE          MODSQL$TEST
DIAGLECT        SQL99
LANGUAGE        C
AUTHORIZATION   SAMPLE_USER
PRAGMA          (IDENT 'V1.2-300')
ALIAS           RDB$DBHANDLE
PARAMETER       COLONS
    
```

---

The DCL command ANALYZE/OBJECT can be used to examine the ident string in the object file.

**Example 9–2 Examining the IDENT in the Object Module**

---

```

$ sql$mod TEST
$ analyze/object TEST/interactive
This is an OpenVMS IA64 (Elf format) object file

Module Identification Information, in note section 2.

    Module name:                "MODSQL$TEST"
    Module version:             "V1.2-300"
    Creation date/time:        "15-JUN-2009 20:04"
    Language name:             "Oracle Rdb SQL V7.2-401"
Press RETURN to continue, or enter a period (.) for next file:
<Ctrl/Z>
$
    
```

Here is similar output from an OpenVMS Alpha system.

```

$ sql$mod TEST
$ analyze/object TEST
.
.
.
This is an OpenVMS Alpha object file

1.  MODULE HEADER (EOBJ$C_EMH), 71 bytes

    structure level: 2
    maximum record size: 4088
    module name: "MODSQL$TEST"
    module version: "V1.0"
    creation   date/time: 16-JUN-2009 11:02
.
.
.
    
```

This example shows the use of the LIBRARIAN to display the ident strings for object modules in a project object library.

```

$ librarian/list/full project.olb
Directory of ALPHA OBJECT library DISK1:[TESTER]PROJECT.OLB;1 on 16-JUN-2009
11:07:23
    
```

## Oracle® Rdb for OpenVMS

Creation date:	16-JUN-2009 11:07:11	Creator:	Librarian A09-30
Revision date:	16-JUN-2009 11:07:11	Library format:	3.0
Number of modules:	1	Max. key length:	128
Other entries:	5	Preallocated index blocks:	213
Recoverable deleted blocks:	0	Total index blocks used:	2
Max. Number history records:	20	Library history records:	0

MODSQL\$TEST Ident V1.2-300 Inserted 16-JUN-2009 11:07:11 5 symbols

---

### 9.1.9 New Keyword for SQL Module Language /PRAGMA Qualifier

This release of Oracle Rdb adds a new option to the /PRAGMA qualifier. The keyword IDENT can be used to pass a text string to the SQL Module Language compiler to be written to the Object Module Header.

The following example demonstrates the use of the qualifier to establish the generation of the compiler module.

```
$ SQL$MOD TEST/PRAGMA=IDENT="v1.2-32"
```

#### *Usage Notes*

- If the PRAGMA (IDENT ...) clause is used as part of the MODULE header then that value will override any value used on the command line.
- The ANALYZE/OBJECT and LIBRARY commands can be used to display this IDENT string and the value will be displayed in LINKER map files.
- OpenVMS limits the IDENT string to a 15 octet string. If the string is longer than this (even with trailing spaces) then an error will be reported by the SQL Module Language compiler.


### 9.1.10 New IDENT Option for SQL Precompiler DECLARE MODULE Statement

Many OpenVMS compilers allow the programmer to specify the object file IDENT string. This allows tracking of the correct version in linker map files (.map) and within object libraries using the LIBRARIAN command. This release of Oracle Rdb supports setting the IDENT string for the SQL precompiler module header.

The IDENT string is specified as part of the PRAGMA clause in the module header.

#### *Syntax*

---



`module-pragma-llst =`



### Usage Notes

- By using the PRAGMA clause with IDENT, you can record an identification string in the object module generated by the SQL Precompiler. This IDENT string is recorded by the OpenVMS LINKER in the image itself and can be viewed in the generated MAP file, examined using ANALYZE/OBJECT, and by the LIBRARIAN command when the object module is stored in an object library.
- OpenVMS limits the IDENT string to a 15 octet string. If the string is longer than this (even with trailing spaces), then an error will be reported by the SQL Precompiler.
- If the IDENT clause is omitted, then the default version string will default to 'V1.0' as is the practice with many OpenVMS compilers. Prior versions of Oracle Rdb on Integrity systems would only provide the string '0'.

```

Module name:                "MODSQL$TEST"
Module version:             "0"
Creation date/time:        "15-JUN-2009 20:13"
Language name:             "Oracle Rdb SQL V7.2-351"
    
```

### Examples

The following example shows the use of a PRAGMA clause in a module header to specify the module ident string.

#### Example 9–3 PRAGMA clause in the DECLARE MODULE statement

```

$ CREATE TEST_HDR.SQL
DECLARE
  MODULE          MODSQL$TEST
  DIALECT         SQL99
  AUTHORIZATION   SAMPLE_USER
  PRAGMA         ( IDENT 'V1.2-300' )
    
```

;

The DCL command ANALYZE/OBJECT can be used to examine the ident string in the object file. Note that the SQL Precompiler generates two object files which are concatenated. Therefore, the ANALYZE will show two MODULE HEADER records, one for the host language (C for example) and one from the SQL Precompiler.

#### *Example 9–4 Examining the IDENT in the object module*

```
$ sql$pre/CC TEST TEST_HDR
$ analyze/object TEST/output=SYS$OUTPUT
.
.
.
This is an OpenVMS Alpha object file

175.  MODULE HEADER (EOBJ$C_EMH), 75 bytes
      structure level: 2
      maximum record size: 4088
      module name: "MODSQL$TEST"
      module version: "V1.2-300"
      creation   date/time: 16-JUL-2009 16:50
.
.
.
```

## 9.1.11 New Keyword for SQL Precompiler PRAGMA Option

This release of Oracle Rdb adds a new keyword to the SQLOPTIONS qualifiers PRAGMA option. The keyword IDENT can be used to pass a text string to the SQL Precompiler to be written to the Object Module Header.

The following example demonstrates the use of the qualifier to establish the generation of the compiler module.

```
$ SQL$PRE/CC TEST/SQLOPTION=(PRAGMA=IDENT="v1.2-32")
```

#### *Usage Notes*

- If the PRAGMA (IDENT ...) clause is used as part of the DECLARE MODULE statement, then that value will override any value used on the command line.
- The ANALYZE/OBJECT and LIBRARY commands can be used to display this IDENT string and the value will be displayed in LINKER map files.
- OpenVMS limits the IDENT string to a 15 octet string. If the string is longer than this (even with trailing spaces) then an error will be reported by the SQL precompiler.

## 9.1.12 RDB\_STATS\_DATABASE Example Program

Accessing performance information in a tabular fashion for Oracle Rdb databases can often be beneficial. In



particular, stored RMU /SHOW STATISTICS rate information in a database can be utilized to do trend analysis and historical review of performance indicators.

RDB\_STATS\_DATABASE is a sample program that reads an RMU /SHOW STATISTICS binary file and converts all statistic values for each sample into a current rate per second. The statistics values are written to a database table named RMU\$STATISTICS. If the RMU\$STATISTICS table does not exist in the database, it will be created.

To use the RDB\_STATS\_DATABASE program, create a foreign command symbol with a value of "\$SQL\$SAMPLE:RDB\_STATS\_DATABASExx.EXE" (where xx is the version of Rdb) and pass an output database and an input binary statistics file name. The following example command sequence demonstrates one possible way that statistics can be gathered for one hour and then formatted.

```
$ RDB_STATS_DATABASE := $SQL$SAMPLE:RDB_STATS_DATABASE72
$ RMU /SHOW STATISTICS MFP -
    /NOINTERACTIVE -
    /OUTPUT = 2008-11-16-00-56.STATS -
    /UNTIL = "16-NOV-2008 11:00:00" -
    /TIME = 15
$ RDB_STATS_DATABASE MYDB.RDB 2008-11-16-00-56.STATS
```

This program can be used to capture either "static" data (from a perviously collected binary file) or "real time" data where records are written to the database as they are produced from RMU /SHOW STATISTICS, as in the following example (note that the RDB\_STATS\_DATABASE example program should be modified when used in this fashion to commit after every record):

```
$ RDB_STATS_DATABASE := $SQL$SAMPLE:RDB_STATS_DATABASE72
$ CREATE /MAILBOX RDB_STATS_DATABASE$MAILBOX -
    /PERMANENT -
    /LOG -
    /BUFFER_SIZE = 65535 -
    /MESSAGE_SIZE = 10000
$ SPAWN RMU /SHOW STATISTICS MFP -
    /NOINTERACTIVE -
    /OUTPUT = RDB_STATS_DATABASE$MAILBOX:
    /UNTIL = "16-NOV-2009 11:00:00" -
    /TIME = 60
$ RDB_STATS_DATABASE MYDB.RDB RDB_STATS_DATABASE$MAILBOX:
```

This example is intended solely to be used as a template for writing your own program. No support for this example template program is expressed or implied.

Oracle Corporation assumes no responsibility for the functionality, correctness or use of this example program. Oracle Corporation reserves the right to change the format and contents of the Oracle Rdb RMU SHOW STATISTICS binary output file at any time without prior notice.

The RDB\_STATS\_DATABASE example program is comprised of the following source modules found in SQL\$SAMPLE:

- RDB\_STATS\_DATABASExx.C
- RDB\_STATS\_DATABASE\_SQL1\_xx.SQLMOD
- RDB\_STATS\_DATABASE\_SQL2\_xx.SQLMOD

Compile and link the RDB\_STATS\_DATABASE example program as follows:

```

$ CC /FLOAT=IEEE RDB_STATS_DATABASExx.C
$ SQL$MOD /FLOAT=IEEE RDB_STATS_DATABASE_SQL1_xx.SQLMOD
$ SQL$MOD /FLOAT=IEEE RDB_STATS_DATABASE_SQL2_xx.SQLMOD
$ LINK RDB_STATS_DATABASExx+-
      RDB_STATS_DATABASE_SQL1_xx+-
      RDB_STATS_DATABASE_SQL2_xx+-
      SQL$USER /LIBRARY

```

## 9.1.13 RCS Time-Based Cache Sweeping

Previously, the Record Cache Server (RCS) process would perform modified row cache "sweep" operations only when a cache was full (also known as "clogged") with modified rows. Now a database may be configured to perform timed cache sweeps. This feature is intended to help perform "lazy" updates of modified rows to the database from caches without performing a full cache checkpoint operation.

The timer for the periodic cache sweeps is specified with the "SWEEP INTERVAL is numeric-literal seconds" clause of the ALTER DATABASE ... ROW CACHE IS ENABLED statement, as in the following example:

```

ALTER DATABASE FILENAME MF_PERSONNEL
  ROW CACHE IS ENABLED (SWEEP INTERVAL IS 300 SECONDS);

```

The number of slots per cache to sweep is specified with the ALTER CACHE statement. Legal values for "SWEEP INTERVAL" are from 0 seconds (to disable periodic timed sweeps) to 3600 seconds (1 hour).

The RMU /SET ROW\_CACHE command accepts a /[NO]SWEEP\_INTERVAL=n qualifier as an alternate method to specify the periodic cache sweep timer. /NOSWEEP\_INTERVAL disables periodic timed sweeps and /SWEEP\_INTERVAL=n can be used to set the timer for the periodic cache sweeps. Legal values for /SWEEP\_INTERVAL=n are from 1 second to 3600 seconds (1 hour).

The Record Cache Server (RCS) process log file contains information about periodic row cache "sweep" operations and can be a useful analysis tool.

---

### Default Value

*The intended default value for the SWEEP INTERVAL in a database is zero seconds (meaning disabled). It is, however, possible for a database that had originally been created with Oracle Rdb Release 7.0 to have a non-zero value. Customers using the row cache feature are advised to explicitly set the SWEEP INTERVAL parameter to either zero (to disable periodic timed sweeps) or the desired sweep interval period on all databases after upgrading to Release 7.2.4.*

---

### Use RMU /SET ROW\_CACHE /NOSWEEP\_INTERVAL

*In Oracle Rdb Release 7.2.4, the "SWEEP INTERVAL is 0 seconds" clause of the ALTER DATABASE ... ROW CACHE IS ENABLED statement may not disable the periodic row cache "sweep" operations. Oracle recommends using the RMU /SET ROW\_CACHE /NOSWEEP\_INTERVAL command as an alternative. This problem will be corrected in a future release.*

---



---

### Incomplete Display Support

*Oracle Rdb Release 7.2.4 does not include complete support for showing a database's periodic row cache "sweep" operation timer value. As a workaround prior to the next Oracle Rdb release, the RMU/DUMP/HEADER/OPTIONS=DEBUG command can be used to display the database parameter RCS\_SWEEP\_INTERVAL as in the following example:*

```
$ RMU/DUMP/HEADER/OPTIONS=DEBUG/OUTPUT=X.X MF_PERSONNEL
$ SEARCH X.X RCS_SWEEP_INTERVAL
RCACHE_CNT = 11.          RCACHE_VBN = 153.          RCS_SWEEP_INTERVAL = 123.
```

---

## 9.1.14 RMU Command TSN Keyword and Qualifier Value

RMU commands that accept a TSN keyword or qualifier value now accept input formats as follows:

- A decimal string representing a quadword TSN value
- A hexadecimal string starting with "%X" representing a quadword TSN value
- A two-part decimal string separated by a colon representing a quadword TSN value as high and low longwords

Following are some example uses of input TSN values:

```
$ RMU /DUMP /AFTER_JOURNAL J1.AIJ /FIRST=TSN=54321
$ RMU /DUMP /AFTER_JOURNAL J1.AIJ /FIRST=TSN=123456234253245
$ RMU /DUMP /AFTER_JOURNAL J1.AIJ /FIRST=TSN=%X7655
$ RMU /DUMP /AFTER_JOURNAL J1.AIJ /FIRST=TSN=%X000000715F856AB
$ RMU /DUMP /AFTER_JOURNAL J1.AIJ /FIRST=TSN=0:871251
$ RMU /DUMP /AFTER_JOURNAL J1.AIJ /FIRST=TSN=3:53487
$ RMU /DUMP /AFTER_JOURNAL J1.AIJ /FIRST=TSN=21:653156
```

## 9.1.15 New Support for RENAME and CREATE SYNONYM Commands

With this release of Oracle Rdb, the RENAME and CREATE SYNONYM commands support INDEX and STORAGE MAP database objects.

- RENAME INDEX changes the name of the index in all system tables.
  - A synonym is created using the old index name to reference the new name of the index. This synonym will be used by any query outline that previously referenced the index using the old name. Note that only a single synonym name may exist. Therefore, if you have indices with the same name as another object, then the RENAME INDEX command may fail if creating the synonym detects a duplicate name.
  - The command ALTER INDEX ... RENAME TO ... is synonymous with the RENAME INDEX command.
- RENAME STORAGE MAP changes the name of the storage map in all system tables.

If the storage map has a companion function in the RDB\$STORAGE\_MAPS system module, then that function will also be renamed. A synonym is created using the old function name to reference the new name of the function. This synonym will be used by any other routine, computed by column, automatic column, and so on that referenced the old storage mapping function.

The command ALTER STORAGE MAP ... RENAME TO ... is synonymous with the RENAME STORAGE MAP command.

- CREATE SYNONYM ... FOR INDEX ... is now supported. Synonyms for indices can be created, altered and dropped.
- CREATE SYNONYM ... FOR STORAGE MAP ... is now supported. Synonyms for storage maps can be created, altered and dropped.

The following example shows the result of the RENAME INDEX and RENAME STORAGE MAP commands.

```
SQL> show table (storage maps,index) employees
Information for table EMPLOYEES

Indexes on table EMPLOYEES:
EMPLOYEES_HASH                with column EMPLOYEE_ID
  No Duplicates allowed
  Type is Hashed Scattered
  Key suffix compression is DISABLED

EMP_EMPLOYEE_ID               with column EMPLOYEE_ID
  No Duplicates allowed
  Type is Sorted
  Key suffix compression is DISABLED
  Node size 430

EMP_LAST_NAME                 with column LAST_NAME
  Duplicates are allowed
  Type is Sorted
  Key suffix compression is DISABLED

Storage Map for table EMPLOYEES:
  EMPLOYEES_MAP

SQL> rename storage map EMPLOYEES_MAP to EMP_STORAGE_MAP;
SQL> rename index EMPLOYEES_HASH to EMP_ID_HASH;
SQL> show table (storage maps,index) employees
Information for table EMPLOYEES

Indexes on table EMPLOYEES:
EMP_EMPLOYEE_ID               with column EMPLOYEE_ID
  No Duplicates allowed
  Type is Sorted
  Key suffix compression is DISABLED
  Node size 430

EMP_ID_HASH                   with column EMPLOYEE_ID
  No Duplicates allowed
  Type is Hashed Scattered
  Key suffix compression is DISABLED

EMP_LAST_NAME                 with column LAST_NAME
  Duplicates are allowed
  Type is Sorted
  Key suffix compression is DISABLED
```

## Oracle® Rdb for OpenVMS

Storage Map for table EMPLOYEES:

EMP\_STORAGE\_MAP

SQL> show storage map

User Storage Maps in database with filename mf\_personnel\_sql

CANDIDATES\_MAP

COLLEGES\_MAP

DEGREES\_MAP

DEPARTMENTS\_MAP

EMP\_STORAGE\_MAP

JOBS\_MAP

JOB\_HISTORY\_MAP

SALARY\_HISTORY\_MAP

WORK\_STATUS\_MAP

SQL> show index

User indexes in database with filename mf\_personnel\_sql

COLL\_COLLEGE\_CODE

DEG\_COLLEGE\_CODE

DEG\_EMP\_ID

DEPARTMENTS\_INDEX

EMP\_EMPLOYEE\_ID

EMP\_ID\_HASH

EMP\_LAST\_NAME

JH\_EMPLOYEE\_ID

JOB\_HISTORY\_HASH

SH\_EMPLOYEE\_ID

EMPLOYEES\_HASH

A synonym for index EMP\_ID\_HASH

SQL> show system function

Functions in database with filename mf\_personnel\_sql

CANDIDATES\_MAP

COLLEGES\_MAP

DEGREES\_MAP

DEPARTMENTS\_MAP

EMP\_STORAGE\_MAP

JOBS\_MAP

JOB\_HISTORY\_MAP

SALARY\_HISTORY\_MAP

WORK\_STATUS\_MAP

EMPLOYEES\_MAP

A synonym for function EMP\_STORAGE\_MAP

SQL>

---

# **Chapter 10**

## **Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.5**

# 10.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.5

## 10.1.1 AIJ Extend Additional Information In Operator Notification New Feature

Bug 8286207

In order to help understand the impact of AIJ extension operations on the system, a new optional feature can be used to display additional OPCOM messages and perform validation (read checking of initialized data) during AIJ extend operations.

If the logical name RDM\$BIND\_AIJ\_EXTEND\_ADDITIONAL\_INFO is defined to a value of "1" (Oracle recommends that this logical be defined system-wide if you intend to take advantage of this feature), and if the database is configured to send operator notifications, additional OPCOM messages will be generated during an AIJ extend operation. This optional feature also includes a validation of the initialization pattern written to the AIJ file. If an invalid pattern is detected, the process performing the extension will bugcheck.

The following example shows the format of the additional messages that indicate the old and new physical EOF locations, the number of blocks of the AIJ file being initialized, the number of IO operations required for the initialization and the ID of the process performing the extension and initialization.

```
%%%%%%%%%% OPCOM 25-MAR-2009 02:31:45.95 %%%%%%%%%%%  
Message from user SUPERDOME on BRDBRY  
Oracle Rdb V7.2-340 Event Notification for Database  
$1$DGA301:[SUPERDOME.V72]FOO.RDB;3
```

```
AIJ journal 512 block extension in progress (new size is 95552 blocks)
```

```
%%%%%%%%%% OPCOM 25-MAR-2009 02:31:45.95 %%%%%%%%%%%  
Message from user SUPERDOME on BRDBRY  
Oracle Rdb V7.2-340 Event Notification for Database  
$1$DGA301:[SUPERDOME.V72]FOO.RDB;3
```

```
AIJ new PEOF = 95616, old PEOF = 95040, init count = 576, PID = 0000049E
```

```
%%%%%%%%%% OPCOM 25-MAR-2009 02:31:45.95 %%%%%%%%%%%  
Message from user SUPERDOME on BRDBRY  
Oracle Rdb V7.2-340 Event Notification for Database  
$1$DGA301:[SUPERDOME.V72]FOO.RDB;3
```

```
AIJ initialize IO count = 3, validating VBN 95041 to 95616, PID = 0000049E
```

## 10.1.2 Default Behavior Change, New Syntax for RMU/RECOVER/CONFIRM

Bugs 7656967 and 8242546

For RMU/RECOVER/NOCONFIRM, where the user cannot be prompted for a decision on how to proceed, the default behavior in the case where the next Oracle Rdb AIJ file to be rolled forward is missing or is not specified in the correct sequence, has been changed. This also affects RMU/RECOVER executed in a batch job where /NOCONFIRM is the default and /CONFIRM cannot be specified unless it is used with the new parameters /CONFIRM=ABORT or /CONFIRM=CONTINUE.

The old behavior for RMU/RECOVER/NOCONFIRM and RMU/RECOVER executed in batch jobs was to assume the user wanted to continue rolling forward if there was an AIJ file sequence gap. This could cause loss of data and invalid indexes due to not rolling forward the database transactions contained in the missing AIJ file. If the user did not do a full RMU/VERIFY of the database following the recovery, he would not be immediately aware of these problems.

To protect the integrity of Rdb databases, the new default behavior for RMU/RECOVER/NOCONFIRM and RMU/RECOVER executed in batch jobs is to terminate the RMU/RECOVER at the point where the out-of-sequence AIJ file is detected. To override this new default behavior, the user can continue to roll forward and ignore the missing AIJ file either by specifying the existing command syntax RMU/RECOVER/CONFIRM to get a prompt on whether to continue rolling forward if there is an AIJ sequence gap, or by specifying the new syntax RMU/CONFIRM=CONTINUE if he does not want the prompt or is executing the RMU/RECOVER in a batch job.

The user can specify the new command syntax RMU/CONFIRM=ABORT if he wants to terminate the RMU/RECOVER at the point where the out-of-sequence AIJ file is detected. This is the new default behavior if /NOCONFIRM is specified or if RMU/RECOVER is executed in a batch job. For interactive recoveries, /CONFIRM, which prompts the user, will continue to be the default.

The new syntax added to the RMU/RECOVER command is the following.

- Do not prompt the user if a sequence gap is detected on the next AIJ file to be rolled forward but ignore the missing AIJ file and continue rolling forward.

```
/CONFIRM=CONTINUE
```

- Do not prompt the user if a sequence gap is detected on the next AIJ roll forward but end the database recover at this point. This is the same as the new default behavior for RMU/RECOVER/NOCONFIRM and RMU/RECOVER in batch.

```
/CONFIRM=ABORT
```

The default for interactive recoveries continues to be /CONFIRM which prompts the user to see if he wants to continue.

The following example shows the new "/CONFIRM=CONTINUE" syntax used so that RMU/RECOVER will continue rolling forward if a sequence gap is detected. This used to be the default behavior if /NOCONFIRM was specified or for batch jobs but is no longer the default behavior.

```
RMU/RECOVER/CONFIRM=CONTINUE/LOG/ROOT=user$test:foo faijbck1,faijbck2,faijbck4
%RMU-I-LOGRECDB, recovering database file DEVICE:[DIRECTORY]FOO.RDB;1

%RMU-I-LOGOPNAIJ, opened journal file DEVICE:[DIRECTORY]FAIJBCK4.AIJ;1
  at 25-FEB-2009 17:26:04.00
%RMU-W-AIJSEQAFT, incorrect AIJ file sequence 8 when 7 was expected
%RMU-I-AIJONEDONE, AIJ file sequence 8 roll-forward operations completed
%RMU-I-LOGRECOVR, 1 transaction committed
%RMU-I-LOGRECOVR, 0 transactions rolled back
```



## Oracle® Rdb for OpenVMS

```
%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 9
%RMU-I-AIJALLDONE, after-image journal roll-forward operations completed
%RMU-I-LOGSUMMARY, total 3 transactions committed
%RMU-I-LOGSUMMARY, total 0 transactions rolled back
%RMU-I-LOGSUMMARY, total 0 transactions ignored
%RMU-I-AIJSUCCEs, database recovery completed successfully
%RMU-I-AIJFNLSEQ, to start another AIJ file recovery, the sequence number
needed will be 9
%RMU-I-AIJNOENABLED, after-image journaling has not yet been enabled
```

The following example shows the new `"/CONFIRM=ABORT"` syntax used so that `RMU/RECOVER` will not continue rolling forward if a sequence gap is detected. This is now the default behavior if `/NOCONFIRM` is specified or for batch jobs. Note that the exit status of `RMU` will be `"%RMU-E-AIJRECESQ"` if the recovery is aborted due to a sequence gap. It is always a good policy to check the exit status of `RMU`, especially when executing `RMU` in batch jobs.

```
RMU/RECOVER/CONFIRM=ABORT/LOG/ROOT=user$test:foo faijbck1,faijbck2,faijbck4
%RMU-I-LOGRECDB, recovering database file DEVICE:[DIRECTORY]FOO.RDB;1

%RMU-I-LOGOPNAIJ, opened journal file DEVICE:[DIRECTORY]FAIJBCK4.AIJ;1
at 25-FEB-2009 17:27:42.29
%RMU-W-AIJSEQAFT, incorrect AIJ file sequence 8 when 7 was expected
%RMU-E-AIJRECESQ, AIJ roll-forward operations terminated due to sequence error
%RMU-I-AIJALLDONE, after-image journal roll-forward operations completed
%RMU-I-LOGSUMMARY, total 2 transactions committed
%RMU-I-LOGSUMMARY, total 0 transactions rolled back
%RMU-I-LOGSUMMARY, total 0 transactions ignored
%RMU-I-AIJFNLSEQ, to start another AIJ file recovery, the sequence number
needed will be 7
%RMU-I-AIJNOENABLED, after-image journaling has not yet been enabled
```

This change has been made in Oracle Rdb Release 7.2.3.5.

---

# **Chapter 11**

## **Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.2**

# 11.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.2

## 11.1.1 New SQL Functions Added

Bug 7365167

This release of Oracle Rdb adds new functions to the SYS\$LIBRARY:SQL\_FUNCTIONS72.SQL script. To replace the existing library of functions first use the SQL\_FUNCTIONS\_DROP72.SQL script and then reapply using SQL\_FUNCTIONS72.SQL.

### *Description*

- ACOS returns the arc cosine of n. The argument n must be in the range of  $-1$  to  $1$ , and the function returns a DOUBLE PRECISION value in the range of  $0$  to  $\pi$ , expressed in radians. If the passed expression results in NULL, then the result of ACOS will be NULL. The following example returns the arc cosine of .3:

```
SQL> SELECT ACOS(.3) "Arc_Cosine" FROM Rdb$DATABASE;  
          Arc_Cosine  
          1.266103672779499E+000  
1 row selected
```

- ACOSH returns the hyperbolic arc cosine of n. The argument n must be equal to or greater than  $1$ . The function returns a DOUBLE PRECISION value. If either passed expression results in NULL, then the result of ACOSH will be NULL.

```
SQL> SELECT ACOSH(1.0) "Hyperbolic Arc Cosine" FROM Rdb$DATABASE;  
          Hyperbolic Arc Cosine  
          0.000000000000000E+000  
1 row selected
```

- ASIN returns the arc sine of n. The argument n must be in the range of  $-1$  to  $1$ , and the function returns a DOUBLE PRECISION value in the range of  $-\pi/2$  to  $\pi/2$ , expressed in radians. If the passed expression results in NULL, then the result of ASIN will be NULL. The following example returns the arc sine of .3:

```
SQL> SELECT ASIN(.3) "Arc_Sine" FROM Rdb$DATABASE;  
          Arc_Sine  
          3.046926540153975E-001  
1 row selected
```

- ASINH returns the hyperbolic arc sine of n. The function returns a DOUBLE PRECISION value. If either passed expression results in NULL, then the result of ASINH will be NULL.

```
SQL> SELECT ASINH (-90.0) "Hyperbolic Arc Sine" FROM Rdb$DATABASE;  
          Hyperbolic Arc Sine  
          -5.192987713658941E+000  
1 row selected
```

- ATAN returns the arc tangent of n. The argument n can be in an unbounded range and returns a value in the range of  $-\pi/2$  to  $\pi/2$ , expressed in radians. If the passed expression results in NULL, then the result of ATAN will be NULL. The following example returns the arc tangent of .3:

## Oracle® Rdb for OpenVMS

```
SQL> SELECT ATAN(.3) "Arc_Tangent" FROM Rdb$DATABASE;
          Arc_Tangent
2.914567944778671E-001
1 row selected
```

- **ATAN2** returns the arc tangent of n and m. The argument n can be in an unbounded range and returns a value in the range of  $-\pi$  to  $\pi$ , depending on the signs of n and m, expressed in radians. **ATAN2(n,m)** is the same as **ATAN(n/m)**.  
If either passed expression results in NULL, then the result of **ATAN2** will be NULL.

```
SQL> SELECT ATAN2(.3, .2) "Arc_Tangent2" FROM Rdb$DATABASE;
          Arc_Tangent2
9.827937232473291E-001
1 row selected
```

- **ATANH** returns the hyperbolic arc tangent of n (in radians). The argument n must be in the range of  $-1$  to  $1$ , and the function returns a **DOUBLE PRECISION** value.  
If either passed expression results in NULL, then the result of **ATANH** will be NULL.

```
SQL> SELECT ATANH(0.905148254) "Hyperbolic Arc Tangent" FROM Rdb$DATABASE;
          Hyperbolic Arc Tangent
1.500000001965249E+000
1 row selected
```

- **BITAND** computes an AND operation on the bits of expr1 and expr2, both of which must resolve to integers, and returns an integer. This function is commonly used with the **DECODE** function, as illustrated in the example that follows.  
If the passed expression results in NULL, then the result of **BITAND** will be NULL.  
The first bit of the mask stored in the column **RDB\$FLAGS** of the table **Rdb\$RELATIONS** indicates that this relation is a view definition. This query displays the names of any views in the database.

### **Example 11–1 Checking bits in RDB\$FLAGS column**

---

```
SQL> -- Which objects in Rdb$RELATIONS are views?
SQL> select rdb$relation_name from rdb$relations where bitand(rdb$flags, 1) = 1;
RDB$RELATION_NAME
RDBVMS$COLLATIONS
RDBVMS$INTERRELATIONS
RDBVMS$PRIVILEGES
RDBVMS$RELATION_CONSTRAINTS
RDBVMS$RELATION_CONSTRAINT_FLDS
RDBVMS$STORAGE_MAPS
RDBVMS$STORAGE_MAP_AREAS
RDBVMS$TRIGGERS
8 rows selected
```

---

This example uses the result of the **BITAND** in a **DECODE** list to display attributes of an Rdb database.

### **Example 11–2 Using BITAND with DECODE**

---

```
SQL> select
cont>   DECODE (BITAND (rdb$flags,1), 0, 'No Dictionary', 'Dictionary'),
cont>   DECODE (BITAND (rdb$flags,2), 0, 'ACL Style', 'ANSI Style'),
cont>   DECODE (BITAND (rdb$flags,64), 0, 'No Multischema', 'Multischema')
cont> from
cont>   Rdb$DATABASE;
```

```
No Dictionary   ACL Style   Multischema
1 row selected
SQL>
```

- COT returns the cotangent of n. The function returns a DOUBLE PRECISION value. If either passed expression results in NULL, then the result of COT will be NULL.

```
SQL> SELECT COT (3.14159265358979/4) "Cotangent" FROM Rdb$DATABASE;
          Cotangent
1.0000000000000002E+000
1 row selected
```

## 11.1.2 RMU /SHOW STATISTICS Enhanced LogMiner Information Display

The RMU /SHOW STATISTICS "LogMiner Information" display has been enhanced to include additional information about the state of Continuous LogMiner processes.

Information displayed includes the following fields (note that the terminal screen width must be greater than 80 columns to display all fields):

- Process.ID – The Process and Stream ID of the Continuous LogMiner process
- State – The current Continuous LogMiner process state:
  - ◆ Inactive
  - ◆ Hibernating
  - ◆ Polling
  - ◆ Extracting
- SeqNum – The AIJ sequence number currently being read
- CurrVBN – The disk block number currently being processed
- Actv – The number of active transactions being processed
- MQP\_VNO – The AIJ sequence number location of the last micro–quiet point detected
- MQP\_VBN – The disk block number location of the last micro–quiet point detected
- MQP\_TSN – The transaction sequence number of the last micro–quiet point detected
- LastTSN – The transaction sequence number of the last transaction

## 11.1.3 RMU /SHOW STATISTICS "Checkpoint Statistics" New Counters

Previously, on the "Checkpoint Statistics" screen of the "RMU /SHOW STATISTICS" utility, two of the possible checkpoint conditions were not being captured for display. These conditions now displayed are:

- "Clear" indicates that a process's checkpoint information is to be cleared in the root file.
- "Initial" indicates that a process has performed its initial checkpoint to establish a checkpoint starting location.

---

Aggregate Total

*Because some of the checkpoint conditions can occur in combination, the sum total of all of the possible checkpoint types may exceed the aggregate total "checkpoints" value displayed.*

---

## 11.1.4 SQL Enhancements: Allowing Optional Correlation Names

Bug 6847694

This release of Oracle SQL improves compatibility with Oracle RDBMS SQL. In particular, some SQL syntax sent via OCI Services for Rdb was rejected in prior releases.

Rdb SQL contains the following enhancements:

1. Oracle RDBMS does not require the correlation name for a derived table (which can be deeply nested). However, in prior releases, Oracle RDB would reject this syntax.

```
SQL> select *
cont> from (select last_name, first_name from candidates);
%SQL-F-CORNAMREQ, A correlation name is required for a derived table
SQL>
SQL> select *
cont> from (select last_name, first_name
cont>         from (select last_name, first_name from candidates));
%SQL-F-CORNAMREQ, A correlation name is required for a derived table
```

Oracle Rdb now also supports an optional correlation name.

```
SQL> select *
cont> from (select last_name, first_name from candidates);
  LAST_NAME      FIRST_NAME
  Wilson         Oscar
  Schwartz       Trixie
  Boswick        Fred
3 rows selected
SQL>
SQL> select *
cont> from (select last_name, first_name
cont>         from (select last_name, first_name from candidates));
  LAST_NAME      FIRST_NAME
  Wilson         Oscar
  Schwartz       Trixie
  Boswick        Fred
3 rows selected
```

2. Oracle RDBMS allows an optional correlation name for INSERT targets. Oracle Rdb did not permit a correlation name in this location.

```
SQL> create table fred (a integer);
SQL> insert into fred f (a) values (1);
insert into fred f (a) values (1);
      ^
%SQL-W-LOOK_FOR_STT, Syntax error, looking for:
%SQL-W-LOOK_FOR_CON,          AS, RETURNING, PLACEMENT, ;,
```

```
%SQL-F-LOOK_FOR_FIN,      found VALUES instead
```

Oracle Rdb now also supports an optional correlation name.

```
SQL> create table fred (a integer);
SQL> insert into fred f (a) values (1);
1 row inserted
```

The correlation name is useful for the RETURNING clause.

```
SQL> insert into fred f (a) values (2) returning f.a;
      A
      2
1 row inserted
```

These problems have been corrected in Oracle Rdb Release 7.2.3.2.

## 11.1.5 Performance Enhancements With Internal Lock Data Structures

In order to help performance for certain classes of applications that lock a large number of records within a transaction, several optimizations have been implemented:

- An internal hash table used to access a list of record locks owned by a database user has been increased in size to help speed access to entries within the table.
- An internal data structure used in conjunction with the hash table is now allocated in larger segments to reduce the number of memory allocations.

## 11.1.6 Change in Frequency of Cardinality Updates Might be Observed

With this release, Oracle Rdb changed slightly the frequency at which table and index cardinality changes, or workload column group rows (RDB\$WORKLOAD) were flushed to the system tables. Some users of RMU/UNLOAD/AFTER\_IMAGE or RMU/OPTIMIZE/AFTER\_IMAGE may observe one more or one less transaction being processed than in prior releases. The reason for this difference is that Rdb attempts to update the system tables on the tail end of a user's READ WRITE transaction and the change in frequency may or may not require Rdb to start an additional transaction during the DISCONNECT processing in which to write back the final statistics.

## 11.1.7 New Interactive SQL Statements

This release of Oracle Rdb adds new Interactive SQL statements to improve compatibility with SQL\*plus.

- SET FEEDBACK { ON | OFF | n }  
The existing SET FEEDBACK statement now supports a numeric option that defines the threshold at which reported line counters are displayed.

```
SQL> set feedback 2
SQL> select * from work_status;
  STATUS_CODE  STATUS_NAME  STATUS_TYPE
  0            INACTIVE   RECORD EXPIRED
  1            ACTIVE     FULL TIME
  2            ACTIVE     PART TIME
```

3 rows selected

```
SQL> set feedback 4
SQL> select * from work_status;
  STATUS_CODE  STATUS_NAME  STATUS_TYPE
  0            INACTIVE   RECORD EXPIRED
  1            ACTIVE     FULL TIME
  2            ACTIVE     PART TIME
```

- **SET LINESIZE n**

The SET LINESIZE command is synonymous with the SET LINE LENGTH command.

- **SET PAGESIZE n**

The SET PAGESIZE command is synonymous with the SET PAGE LENGTH command.

- **SET TIMING { ON | OFF }**

The SET TIMING enables a single line report of used CPU and Elapsed time for each successful SQL statement or command. This is shown in the following example.

```
SQL> start transaction;
SQL> set timing on;
SQL> select count(*)
cont> from employees
cont>     inner join job_history using (employee_id)
cont>     inner join salary_history using (employee_id)
cont>     inner join departments using (department_code)
cont>     inner join jobs using (job_code)
cont>     left outer join resumes using (employee_id)
cont>     left outer join degrees using (employee_id)
cont>     left outer join colleges using (college_code)
cont>
cont> ;

          3871
1 row selected
Timing: Elapsed:    0 00:00:00.82 Cpu:    0 00:00:00.16
SQL> set timing off;
SQL> commit;
```

---



# **Chapter 12**

## **Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.0**

# 12.1 Enhancements And Changes Provided in Oracle Rdb Release 7.2.3.0

## 12.1.1 Optional Run-Time Routine Native Compiler on I64 Enabled By Default

On Alpha and VAX systems, Oracle Rdb generates, at run time, callable subroutines of native machine instructions. These subroutines are executed during the processing of database insert/update/retrieval operations.

Oracle Rdb on I64 systems uses a hardware-portable instruction interpreter. This allowed rapid development and deployment of Oracle Rdb on Integrity OpenVMS systems. Release 7.2.1.1 of Oracle Rdb introduced a second pass optimization that converts these portable instructions to native I64 machine code for enhanced performance. This feature is enabled by default starting with Release 7.2.3 of Oracle Rdb.

To disable the run-time compiler, define `RDMS$BIND_CODE_OPTIMIZATION` to a value of "0" or use the flag `"CODE_OPTIMIZATION(0)"`. The run-time compiler is enabled by default. To re-enable, if needed, the run-time compiler on I64 systems, either deassign the logical name `RDMS$BIND_CODE_OPTIMIZATION` or define the logical name `RDMS$BIND_CODE_OPTIMIZATION` to a value of "2" or use the flag `"CODE_OPTIMIZATION(2)"`. This logical name and flag have no effect on Alpha systems. The following examples show use of the logical name and the flag to enable and disable and show the status of the optional run-time compiler on I64 systems.

```
! To disable:
$ DEFINE RDMS$BIND_CODE_OPTIMIZATION 0
$ DEFINE RDMS$SET_FLAGS "CODE_OPTIMIZATION(0)"
SQL> SET FLAGS 'CODE_OPTIMIZATION(0)';

! To enable:
$ DEFINE RDMS$BIND_CODE_OPTIMIZATION 2
$ DEFINE RDMS$SET_FLAGS "CODE_OPTIMIZATION(2)"
SQL> SET FLAGS 'CODE_OPTIMIZATION(2)';

! Show current setting enabled and disabled:
SQL> ATT 'FI MF_PERSONNEL';
SQL> SET FLAGS 'CODE_OPTIMIZATION(2)';
SQL> SHOW FLAGS

Alias RDB$DBHANDLE:
Flags currently set for Oracle Rdb:
    PREFIX,WARN_DDL,INDEX_COLUMN_GROUP,MAX_SOLUTION
    ,MAX_RECURSION(100),REFINE_ESTIMATES(127)
    ,CODE_OPTIMIZATION(2),NOBITMAPPED_SCAN

SQL> SET FLAGS 'CODE_OPTIMIZATION(0)';
SQL> SHOW FLAGS

Alias RDB$DBHANDLE:
Flags currently set for Oracle Rdb:
    PREFIX,WARN_DDL,INDEX_COLUMN_GROUP,MAX_SOLUTION
    ,MAX_RECURSION(100),REFINE_ESTIMATES(127),NOBITMAPPED_SCAN
SQL>
```

There is CPU overhead required to compile a subroutine and not every subroutine can be compiled, at present, into native Itanium routines. Additional virtual memory is required for the functionality of the optional run-time compiler.

If you detect a difference in the functionality or behavior of Oracle Rdb when this feature is enabled, you may define the logical name to revert to the prior (interpreted) execution model.

## 12.1.2 Temporary Table Improvements

In prior versions of Oracle Rdb, all temporary table data structures were always allocated in 32-bit process (P0) virtual address space. In several cases, the 1GB size of P0 space drastically limited the number of rows that could be stored or would result in unexpected errors when deleting or modifying rows. Additionally, performance when accessing rows within temporary tables would degrade as additional rows were inserted.

The following enhancements have been made to address these cases:


- With the exception of the actual data rows themselves, most of the data structures related to temporary tables have been moved to 64-bit process (P2) virtual address space. The data content remains in 32-bit address space.
- Use of the logical name RDMS\$TTB\_HASH\_SIZE to control the size of an internal hash table may no longer be required for consistent performance. The hash table is now initially sized larger than it was in the past and will automatically extend as needed to accommodate the rows in the temporary table with more consistent performance.

## 12.1.3 New SIGN Builtin Function

This release of Oracle Rdb enhances the Conditional expressions with support for the SIGN builtin function.

### Syntax

**conditional-expr =**



*Usage Notes*

- SIGN returns an INTEGER value.
- SIGN accepts any numeric (fixed or floating) or interval value expression.
- If the value expression evaluates to NULL, then the SIGN function returns NULL.
- If the value expression evaluates to a negative value, then SIGN returns -1; if the value is positive then SIGN returns 1; otherwise a zero will be returned.
- The SQL\_FUNCTIONS script continues to add the SIGN function to the database. However, this function is now deprecated and is retained only for backward compatibility with applications built using that function.

*Examples*

This example computes delayed departures from the LAYOVER table.

*Example 12-1 Using SIGN Builtin Function*


---

```
SQL> select arr_date,
cont>         dep_date,
cont>         DECODE (SIGN ((dep_date - arr_date) day(9)),
cont>             -1, 'date error - can not depart before arrival',
cont>             0, 'same day departure',
cont>             1, 'delayed')
cont> from LAYOVER;
```

2005-12-22	2006-01-20	delayed
2005-12-23	2005-12-25	delayed
2006-01-30	2006-02-01	delayed
2006-02-06	2006-02-09	delayed
2006-01-24	2006-01-26	delayed
2006-02-02	2006-02-19	delayed
2007-02-10	2007-02-16	delayed
2007-02-20	2007-02-26	delayed
2007-05-29	2007-06-08	delayed
2007-06-12	2007-06-26	delayed
2007-05-15	2007-05-21	delayed
2007-09-10	2007-09-14	delayed
2007-09-04	2007-09-06	delayed
2007-09-19	2007-09-20	delayed
2007-09-21	2007-09-24	delayed

15 rows selected


---

## 12.1.4 Enhanced Simple CASE Expression


This release of Oracle Rdb enhances the Simple CASE expression to support value lists and partial predicates.

*Syntax*


**simple-case-expr =**



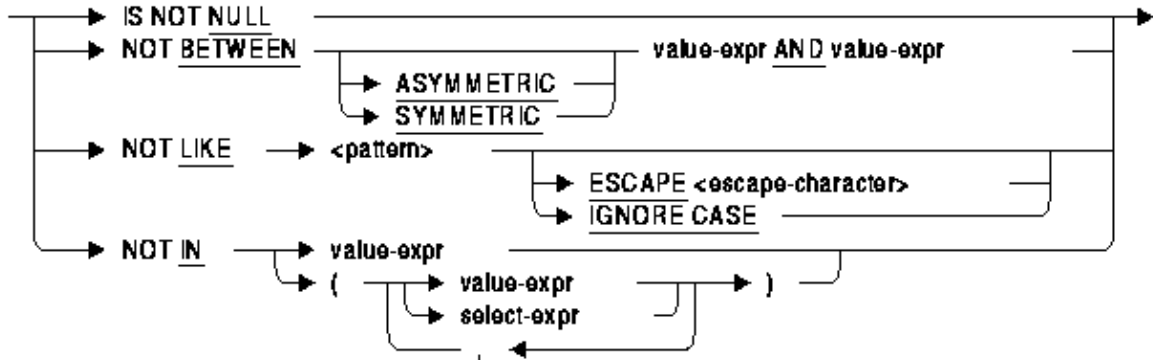
**with-values =**



**partial\_predicate1 =**



partial\_predicate2 =



### Usage Notes

- The WHEN clause can now include a list of value expressions. In prior versions, the WHEN clause was limited to just one value expression. Each value expression is separated by a comma.
- The WHEN clause can now include partial predicates. In prior versions, the default operator that was applied to the case primary expression of each WHEN value expression was equality (=). With this release, this can be overridden. Multiple partial predicates can also be listed.  
A WHEN clause partial predicate can begin with one of the following operators: IS NULL, IS NOT NULL, BETWEEN, NOT BETWEEN, IN, NOT IN, LIKE, NOT LIKE, STARTING WITH, NOT STARTING WITH, CONTAINING, NOT CONTAINING, <=, <, <>, >=, >, ^=, !=, and =.
- If the WHEN clause includes the value expression NULL, it assumes that this is equivalent to the IS NULL operator.

### Examples

This example calculates the calendar quarter using a CASE expression.

#### Example 12–2 Using the IN clause as a partial predicate

```
SQL> select
cont>     salary_end,
cont>     case extract (month from salary_end)
cont>         when in (1, 2, 3) then 'Q1'
cont>         when in (4, 5, 6) then 'Q2'
cont>         when in (7, 8, 9) then 'Q3'
cont>         when in (10, 11, 12) then 'Q4'
cont>         else 'UNKNOWN'
cont>     end quarter
cont> from salary_history
cont> where employee_id = '00164'
cont> order by salary_start;
SALARY_END    QUARTER
2-Mar-1981    Q1
21-Sep-1981   Q3
14-Jan-1983   Q1
NULL          UNKNOWN
4 rows selected
```

SQL>

A simple list of values could also be used to achieve the same result.

```
SQL> select
cont>     salary_end,
cont>     case extract (month from salary_end)
cont>         when 1, 2, 3 then 'Q1'
cont>         when 4, 5, 6 then 'Q2'
cont>         when 7, 8, 9 then 'Q3'
cont>         when 10, 11, 12 then 'Q4'
cont>         else 'UNKNOWN'
cont>     end quarter
cont> from salary_history
cont> where employee_id = '00164'
cont> order by salary_start;
SALARY_END    QUARTER
  2-Mar-1981   Q1
 21-Sep-1981   Q3
 14-Jan-1983   Q1
        NULL   UNKNOWN
4 rows selected
SQL>
```

---

This example shows the use of partial predicates to process the LAST\_NAME column and provide a customized ordering of the EMPLOYEES table.

### *Example 12–3 Using Partial predicates in the CASE expression*

---

```
SQL> select
cont>     concat (trim (first_name), ' ',
cont>             nvl2 (middle_initial,
cont>                   concat (middle_initial, '.', last_name),
cont>                   last_name)) as names
cont> from employees
cont> order by
cont>     case last_name
cont>         when starting with 'Mc'
cont>         then
cont>             substring (last_name from 3)
cont>         when starting with 'Mac'
cont>         then
cont>             substring (last_name from 4)
cont>         when containing ''''
cont>         then
cont>             substring (last_name from position (''' in last_name)+1)
cont>         else
cont>             last_name
cont>     end;
NAMES
Louie A.Ames
Aruwa D'Amico
Leslie Q.Andriola
Joseph Y.Babbin
Dean G.Bartlett
.
.
.
```

Brian Wood  
 Al F.Ziemke  
 100 rows selected

---

This example shows the use of value lists to make the simple case more readable. Using input values for the month and year, it returns the number of days in that month.

#### ***Example 12–4 Using lists of values in the Simple CASE expression***

---

```
SQL> set flags 'trace';
SQL> begin
cont> declare :year_val, :month_val int;
cont> set :year_val = extract (year from current_date);
cont> for :month_val in 1 to 12
cont> do
cont>     trace
cont>         case :month_val
cont>             when 1, 3, 5, 7, 8, 10, 12 then
cont>                 31
cont>             when 4, 6, 9, 11 then
cont>                 30
cont>             when 2 then
cont>                 DECODE (MOD (:year_val, 4), 0, 29, 28)
cont>             else
cont>                 NULL
cont>         end;
cont> end for;
cont> end;
~Xt: 31
~Xt: 29
~Xt: 31
~Xt: 30
~Xt: 31
~Xt: 30
~Xt: 31
~Xt: 31
~Xt: 30
~Xt: 31
~Xt: 31
~Xt: 30
~Xt: 31
~Xt: 30
~Xt: 31
```

---

## **12.1.5 Changes in Generated Query Outline ID**

This release of Oracle SQL has added various optimizations to the Simple CASE expression as well as the DECODE and ABS functions. These changes might, in some cases, change the generated intermediate query. This will have no effect on the query results but will change the query id generated for a query outline. In such cases, the query outline should be re-created to ensure matching by the new outline id.

Alternately, queries using these query outlines can be modified to use the OPTIMIZE USING clause to use the name of the query outline. Then such changes as described here do not affect the selection of the query outline. If a query outline is used in conjunction with the RMU Unload command then the /OPTIMIZE=USING clause can be used to specify the name of the query outline for use by RMU.

---

Note



*This query outline id is also displayed by the SET FLAGS 'WATCH\_OPEN' flag as shown in this example:*

```
SQL> set flags 'watch_open';
SQL> select decode (employee_id, '00164', 'Toliver', 'Others') from employees;
~Xo: Start Request ECB31E522CC71247B16B18660AD42F1D (unnamed)
.
.
.
```

*This is the query id generated by Oracle Rdb V7.1.*

```
SQL> select decode (employee_id, '00164', 'Toliver', 'Others') from employees;
~Xo: Start Request 77D353270F7842D60FCA3D6CC11378D3 (unnamed)
.
.
.
```

## 12.1.6 ALTER INDEX ... MAINTENANCE IS ENABLED DEFERRED Syntax is Now Active

Bug 6880752

In prior versions of Oracle Rdb, the ALTER INDEX ... MAINTENANCE IS ENABLED DEFERRED syntax was not active. Instead it produced a WISH\_LIST error. This clause has now been activated for this release.

The following example shows the WISH\_LIST error generated in prior versions.

```
SQL> alter index t1_idx maintenance is enabled deferred;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDB-F-WISH_LIST, feature has not been implemented
```

When ALTER INDEX ... MAINTENANCE IS ENABLED DEFERRED is used on an index, it has the following actions:

- If the index is MAINTENANCE IS ENABLED IMMEDIATE mode, that index is changed to a build-pending index. The side effect of this change is that the table upon which the index is defined can now only be read. The index partitions may now be rebuilt using REBUILD PARTITION or a combination of TRUNCATE PARTITION and BUILD PARTITION clauses. A warning message is issued upon successful execution of the ALTER INDEX command so that the database administrator is aware that the index is incomplete and also that the table upon which the index is defined is write-disabled so that INSERT, UPDATE and DELETE statements cannot be used.

```
SQL> alter index t1_idx maintenance is enabled deferred;
%RDB-W-META_WARN, metadata successfully updated with the reported warning
-RDMS-W-IDXBLDAPEND, index in build pending state - maintenance is disabled
SQL>
```

After each partition is rebuilt, a final ALTER INDEX ... MAINTENANCE IS ENABLED IMMEDIATE statement will be required to make the index active and allow updates to the table. This final step validates all partitions, rolls up partition cardinality information and clears the build-pending state.

- If the index is MAINTENANCE IS ENABLED DEFERRED mode then there are no changes made to the index and the ALTER INDEX statement quietly succeeds.
- If the index is MAINTENANCE IS DISABLED then an error is raised. This ALTER INDEX operation is not compatible with such indices. However, ALTER INDEX ... TRUNCATE ALL PARTITIONS or ALTER INDEX ... TRUNCATE PARTITION can be used to modify a disabled index into a build-pending index.

```
SQL> alter index t1_idx maintenance is enabled deferred;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-INDMAINTDIS, maintenance on index T1_IDX has been disabled
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.2.3.

## 12.1.7 SQL Precompiler and Module Language Compiler /ARCHITECTURE Command Line Qualifier

For improved performance of generated code, the SQL Language Precompiler and Module Language Compilers include support for the /ARCHITECTURE=keyword command line qualifier on OpenVMS Alpha systems. At present, the /ARCHITECTURE qualifier is ignored on Itanium systems.

The /ARCHITECTURE=keyword qualifier specifies the lowest version of the Alpha architecture where this code will run which can allow the compiler to generate more efficient code with the tradeoff that code may not run on older systems.

All Alpha processors implement a core set of instructions and, in some cases, the following extensions:

- Byte/word extension (BWX) – The instructions that comprise the BWX extension are LDBU, LDWU, SEXTB, SEXTW, STB, and STW.
- Square-root and floating-point convert extension (FIX) – The instructions that comprise the FIX extension are FTOIS, FTOIT, ITOFF, ITOFS, ITOFT, SQRTF, SQRTG, SQRTS, and SQRTT.
- Count extension (CIX) – The instructions that comprise the CIX extension are CTLZ, CTPOP, and CTTZ.
- Multimedia extension (MVI) – The instructions that comprise the MVI extension are MAXSB8, MAXSW4, MAXUB8, MAXUW4, MINSB8, MINSW4, MINUB8, MINUW4, PERR, PKLB, PKWB, UNPKBL, and UNPKBW.

The Alpha Architecture Reference Manual describes the extensions in detail.

The keyword specified with the /ARCHITECTURE qualifier determines which instructions the compiler can generate and which coding rules it must follow.

- GENERIC – Generate instructions that are appropriate for all Alpha processors. This option is the default and is equivalent to /ARCH=EV4.
- HOST – Generate instructions for the processor that the compiler is running on (for example, EV56

## Oracle® Rdb for OpenVMS

- instructions on an EV56 processor, EV7 instructions on an EV7 processor, and so on).
- EV4 – Generate instructions for the EV4 processor (21064, 20164A, 21066, and 21068 chips). Applications compiled with this option will not incur any emulation overhead on any Alpha processor.
  - EV5 – Generate instructions for the EV5 processor (some 21164 chips). (Note that the EV5 and EV56 processors both have the same chip number – 21164.) Applications compiled with this option will not incur any emulation overhead on any Alpha processor.
  - EV56 – Generate instructions for EV56 processors (some 21164 chips). This option permits the compiler to generate any EV4 instruction, plus any instructions contained in the BWX extension. Applications compiled with this option may incur emulation overhead on EV4 and EV5 processors.
  - PCA56 – Generate instructions for PCA56 processors (21164PC chips). This option permits the compiler to generate any EV4 instruction plus any instructions contained in the BWX and MVI extensions. Applications compiled with this option may incur emulation overhead on EV4 and EV5 processors.
  - EV6 – Generate instructions for EV6 processors (21264 chips). This option permits the compiler to generate any EV4 instruction, any instruction contained in the BWX and MVI extensions, plus any instructions added for the EV6 chip. These new instructions include a floating–point square root instruction (SQRT), integer/floating–point register transfer instructions, and additional instructions to identify extensions and processor groups. Applications compiled with this option may incur emulation overhead on EV4, EV5, EV56, and PCA56 processors.
  - EV67 or EV68 – Generate instructions for EV67 and EV68 processors (21264A chips). This option permits the compiler to generate any EV6 instruction, plus the new bit count instructions (CTLZ, CTPOP, and CTTZ). However, the precompilers do not currently generate any of the new bit count instructions and the EV67 and EV68 have identical instruction scheduling models so the EV67 and EV68 are essentially identical to the EV6. Applications compiled with this option may incur emulation overhead on EV4, EV5, EV56, and PCA56 processors.
  - EV7 – Generate instructions for the EV7 processor (21364 chip). This option permits the compiler to generate any EV67 instruction. There are no additional instructions available on the EV7 processor but the compiler does have different instruction scheduling and prefetch rules for tuning code for the EV7. Applications compiled with this option may incur emulation overhead on EV4, EV5, EV56, and PCA56 processors.

The OpenVMS Alpha operating system includes an instruction emulator. This capability allows any Alpha chip to execute and produce correct results from Alpha instructions – even if some of the instructions are not implemented on the chip. Applications using emulated instructions will run correctly but may incur significant emulation overhead at run time.

Of the available extension types, the Byte/word extension (BWX) will often be beneficial for increased performance of Rdb–based applications. In addition, for those Alpha implementations that support quad–issue of instructions (the EV6 and later processors), the compiler does have different instruction scheduling and prefetch rules for tuning code.

For highest levels of performance of generated code, Oracle recommends that the /ARCHITECTURE switch be specified with the keyword that most closely matches the lowest processor type of the machine where the program will execute.

---

### Language Compiler Support for /ARCHITECTURE

*If specified, the /ARCHITECTURE qualifier is passed on the command line to the specified language compiler by the SQL Precompiler. The language compiler being used must support the /ARCHITECTURE qualifier and the architecture keyword value when the*

*/ARCHITECTURE qualifier is specified.*

---

## 12.1.8 PERFT4\_RDB Example Program

Accessing performance information in a tabular fashion for Oracle Rdb databases can be often be beneficial. In particular, data in "CSV" (comma separated values) format can be readily loaded into other applications. HP's T4 utility program TIViz expects data in an enhanced CSV format where the first 4 lines of the data file contain additional information.

SQL\$SAMPLE:PERFT4\_RDBxx.C is a sample C program that reads an RMU /SHOW STATISTICS binary file and converts all statistic values for each sample into a current rate per second. The statistics values are written to an output text file. This example is supplied as a template for writing your own program.

To use the PERFT4\_RDB program, optionally compile and link (the example program has been supplied also as a .EXE) and then create a foreign command symbol with a value of "\$SQL\$SAMPLE:PERFT4\_RDBxx.EXE" (where xx is the version of Rdb) and pass an input binary statistics file name and an output text file name. The following example command sequence demonstrates one possible way that statistics can be gathered for one hour and then formatted.

```
$ PERFT4 := $SQL$SAMPLE:PERFT4_RDB72
$ RMU /SHOW STATISTICS MFP -
    /NOINTERACTIVE -
    /OUTPUT = 2008-11-16-00-56.STATS -
    /UNTIL = "16-NOV-2008 11:00:00" -
    /TIME = 15
$ PERFT4 2008-11-16-00-56.STATS 2008-11-16-00-56.CSV
Wrote 251 records (from 16-Nov-2008 09:56:29 to 16-Nov-2008 10:59:57)
```

This example source code SQL\$SAMPLE:PERFT4\_RDBxx.C is intended solely to be used as a template for writing your own program. No support for this example template program is expressed or implied. Oracle Corporation assumes no responsibility for the functionality, correctness or use of this example program. Oracle Corporation reserves the right to change the format and contents of the Oracle Rdb RMU SHOW STATISTICS binary output file at any time without prior notice.

## 12.1.9 RMU Load Quietly Truncated String Data During Insert

Bug 6732438

In prior versions of Oracle Rdb, the RMU Load command would quietly truncate data loaded into CHAR and VARCHAR columns. This loss of data might be significant but was never reported by Oracle Rdb.

This problem has been corrected in Oracle Rdb Release 7.2.3. RMU Load now defaults to SQL dialect SQL99 which implicitly checks for and reports truncations during INSERT operations.

This behavior is controlled by a new /DIALECT qualifier.

- /NODIALECT, /DIALECT=SQL89 or /DIALECT=NONE will revert to prior behavior and no truncation error will be reported.

- /DIALECT=SQL99 (the default) will enable reporting of truncation errors. Note that truncation occurs if non-space characters are discarded during the insert.

```
$ rmu/load abc f2 f1
%RMU-I-LOADERR, Error loading row 1.
%RDB-E-TRUN_STORE, string truncated during assignment to a column
%RMU-I-DATRECREAD, 1 data records read from input file.
%RMU-I-DATRECSTO, 0 data records stored.
%RMU-F-FTL_LOAD, Fatal error for LOAD operation at 13-FEB-2008 15:39:44.40
$
```

## 12.1.10 New FETCH FIRST and OFFSET Clauses for Select Expression

This release of Oracle Rdb adds two new clauses to the SELECT expression. These clauses are defined by the draft SQL Database Language Standard 2008.

- OFFSET skip-expression

The OFFSET clause allows the database programmer to start fetching the result rows from the specified offset within the result table. OFFSET accepts a numeric value expression which may contain arbitrary arithmetic operators, function calls, subselect clauses or sequence references. The subselect clauses may not reference columns in the outer query as it is evaluated before row processing begins.

---

### Note

*Oracle recommends that the values specified for skip-expression be kept small for performance reasons. The skipped rows are still fetched and processed by the query; they are just not returned to the application.*

---

The OFFSET clause is equivalent in functionality to the SKIP clause currently supported by the LIMIT TO clause. The distinction is that OFFSET can be specified without a row limit.

OFFSET is not compatible with the SKIP (or OFFSET) sub clause of the LIMIT TO clause. However, OFFSET and LIMIT TO can be used together.

- FETCH FIRST limit-expression
- FETCH NEXT limit-expression

The FETCH FIRST clause allows the database programmer to limit the results returned from a query expression. The FETCH FIRST clause is equivalent to functionality currently supported by the LIMIT TO clause. FETCH accepts a numeric value expression which may contain arbitrary arithmetic operators, function calls, subselect clauses or sequence references. The subselect clauses may not reference columns in the outer query as it is evaluated before row processing begins.

The FETCH NEXT is identical to FETCH FIRST but allows the syntax to be more descriptive when coupled with the OFFSET clause.

If no value expression is provided for FETCH, it will default to 1 row.

The FETCH clause is not compatible with the LIMIT TO clause.

### Examples

The following examples show the use of the FETCH FIRST and OFFSET clauses.

This example uses the DEPARTMENTS table to locate the employee id of each manager and, after sorting them by their birthday, the oldest manager's name and employee id are displayed.

**Example 12–5 Using *FETCH FIRST* to find the oldest manager in the company**

```
SQL> -- select the most senior manager
SQL> select e.last_name, e.first_name, e.employee_id
cont> from departments d, employees e
cont> where d.manager_id = e.employee_id
cont> order by e.birthday
cont> fetch first row only;
  E.LAST_NAME      E.FIRST_NAME      E.EMPLOYEE_ID
  O'Sullivan      Rick              00190
1 row selected
SQL>
```


This query uses a subselect in the OFFSET clause to locate the median (or middle) row of the sorted set.

**Example 12–6 Using *OFFSET ROWS* and *FETCH NEXT* to compute the median salaried employee**


```
SQL> select e.last_name, e.first_name, employee_id, sh.salary_amount
cont> from salary_history sh inner join employees e using (employee_id)
cont> where sh.salary_end is null
cont> order by sh.salary_amount
cont> offset (select count(*)
cont>          from salary_history
cont>          where salary_end is null)/2 rows
cont> fetch next row only;
  E.LAST_NAME      E.FIRST_NAME      EMPLOYEE_ID      SH.SALARY_AMOUNT
  Boyd            Ann              00244            $24,166.00
1 row selected
SQL>
```

**Syntax**

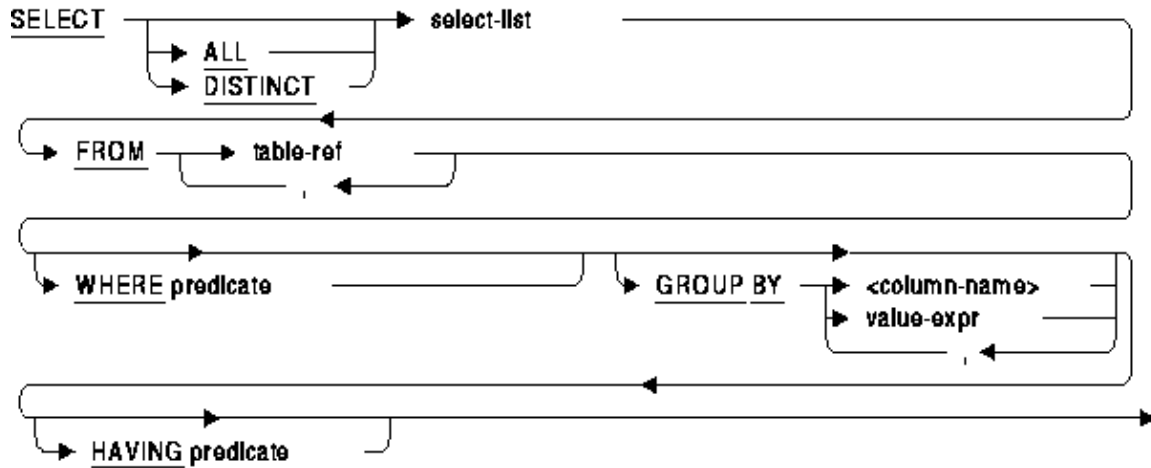
**select-expr =**



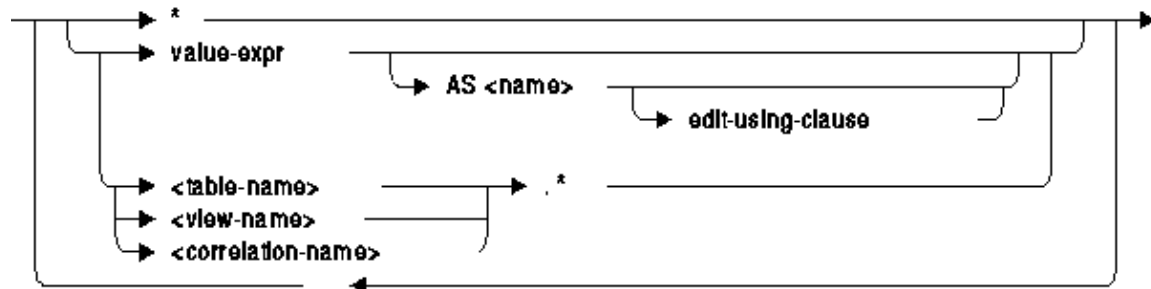
**select-merge-clause =**




**select-clause =**




**select-list =**




**edit-using-clause =**



**table-ref =**




**derived-table =**



**joined-table =**




**qualified-join =**







**cross-join =**




**join-type =**



**correlation-name-clause =**




**order-by-clause =**




**offset-clause =**



limit-to-clause =



optimize-clause =



### Usage Notes

- If ORDER BY is used in a query that includes OFFSET, FETCH FIRST (FETCH NEXT), or LIMIT TO clauses then the rows are first retrieved and sorted prior to applying the OFFSET, FETCH or LIMIT TO actions.
- A select expression may contain both OFFSET and FETCH NEXT (LIMIT TO) clauses, in which case the OFFSET is applied first and then the FETCH NEXT clause. For instance, if the query would normally return 100 rows, then an OFFSET 20 would skip over the first 20 rows and return the remaining 80. If on the other hand an OFFSET 20 and a LIMIT TO 20 were specified, then after skipping the first 20 rows the next 20 are returned.
- The OFFSET, FETCH FIRST or LIMIT TO clauses may result in no rows being retrieved.

## 12.1.11 RMU Unload Now Creates SQL\*Loader Control Files

Enhancement 2146782

This release of Oracle Rdb has enhanced RMU Unload by providing support for SQL\*Loader control files and portable data files.

The following example shows the FORMAT=CONTROL option.

```
$ RMU/UNLOAD/RECORD_DEFINITION=(FORMAT=CONTROL,FILE=EMP) -  
  SQL$DATABASE -  
  EMPLOYEES -  
  EMPLOYEES
```

This command creates a file EMP.CTL (the SQL\*Loader control file) and EMPLOYEES.DAT in a portable format to be loaded.

### *Usage Notes*

- FORMAT=CONTROL implicitly uses a portable data format as TEXT rather than binary values. The unloaded data files are similar to that generated by FORMAT=TEXT but include a NULL vector to represent NULL values ('1') and non-NULL values ('0').  
The SQL\*Loader control file uses this NULL vector to set NULL for the data upon loading.
  - When FORMAT=CONTROL is used, the output control file and associated data file are intended to be used with the Oracle RDBMS SQL\*Loader (sqlldr) command to load the data into an Oracle RDBMS database table. LIST OF BYTE VARYING (SEGMENTED STRING) columns are not unloaded.  
The file specification for the FILE option will default to a .CTL file type.  
Oracle Rdb does not support this format for RMU Load.
  - The keywords NULL, PREFIX, SEPARATOR, SUFFIX, and TERMINATOR only apply to DELIMITED\_TEXT format and may not be used in conjunction with the CONTROL keyword.
  - DATE VMS data is unloaded including the fractional seconds precision. However, when mapped to Oracle DATE type in the control file, the fractional seconds value is ignored. It is possible to modify the generated control file to use the TIMESTAMP type and add FF to the date edit mask.
-

# **Chapter 13**

## **Documentation Corrections, Additions and Changes**

This chapter provides corrections for documentation errors and omissions.

# 13.1 Documentation Corrections

## 13.1.1 RMU /SET ROW\_CACHE Command Updates

The documentation and online help for the "RMU /SET ROW\_CACHE" command inadvertently did not include the full set of allowed keywords and qualifiers.

The valid command line qualifiers for the "RMU /SET ROW\_CACHE" command are:

- **Alter** – Specifies the action to take on the named cache. You must specify the cache name and at least one other option.
- **Disable** – Disables row caching. Do not use with the **Enable** qualifier.
- **Enable** – Enables row caching. Do not use with the **Disable** qualifier.
- **Log** – Specifies whether the processing of the command is reported to SYS\$OUTPUT. Specify the **Log** qualifier to request log output and the **Nolog** qualifier to prevent it. If you specify neither, the default is the current setting of the DCL verify switch.
- **Backing\_Store\_Location=devdir** – Specify the per-database default backing store location.
- **NoBacking\_Store\_Location** – Remove the per-database default backing store location and revert back to the default backing store file location of the root file device and directory.

The valid values for the **ALTER** qualifier are:

- **NAME=cachename** – Name of the cache to be modified. The cache must already be defined in the database. This is a required parameter. This parameter accepts the wildcard characters asterisk (\*) and percent sign (%).
- **ENABLE** – Enable the cache.
- **DISABLE** – Disable the cache.
- **DROP** – Drop (delete) the cache.
- **SNAPSHOT\_SLOT\_COUNT=n** – Specify the number of snapshot slots in the cache. A value of zero disables the snapshot portion for the specified cache.
- **SLOT\_COUNT=n** – Specify the number of slots in the cache.
- **SLOT\_SIZE=n** – Specify the size (in bytes) of each slot in the cache.
- **WORKING\_SET\_COUNT=n** – Specify the number of working set entries for the cache. Valid values are from 1 to 100.
- **BACKING\_STORE\_LOCATION=devdir** – Specify the per-cache default backing store location.
- **NOBACKING\_STORE\_LOCATION** – Remove the per-cache default backing store location and revert back to the database default backing store file location.
- **SHARED\_MEMORY** – Specify the shared memory type and parameters for the cache. Valid keywords are:
  - ◆ **TYPE=PROCESS** to specify traditional shared memory global section, which means that the database global section is located in process (P0) address space and may be paged from the processes working set as needed.
  - ◆ **TYPE=RESIDENT** to specify that the database global section is memory resident in process (P0) address space using OpenVMS Alpha shared page tables, which means that a system space global section is fully resident, or pinned, in memory.
  - ◆ **RAD\_HINT= "number"** to indicate a request that memory for this shared memory should be allocated from the specified OpenVMS Alpha Resource Affinity Domain (RAD). This parameter specifies a hint to Oracle Rdb and OpenVMS about where memory should be

physically allocated. It is possible that if the memory is not available, it will be allocated from other RADs in the system. For systems that do not support RADs, no RAD\_HINT specification is valid.

The RAD\_HINT qualifier is only valid when the shared memory type is set to RESIDENT. Setting the shared memory type to SYSTEM or PROCESS explicitly disables any previously defined RAD hint.

---

Note

*OpenVMS support for RADs is available only on the AlphaServer GS series systems. For more information about using RADs, refer to the OpenVMS Alpha Partitioning and Galaxy Guide.*

---

- ◆ NORAD\_HINT disables the RAD hint.

The "/ALTER=(...)" qualifier may be specified multiple times on the command line. Each /ALTER qualifier specified operates on one unique cache if no wildcard character (% or \*) is specified. Otherwise, Each /ALTER operates on all matching cache names.

For example, the following command alters two caches:

```
$ RMU /SET ROW_CACHE MF_PERSONNEL -
      /ALTER= ( NAME = RDB$SYS_CACHE,
                SLOT_COUNT = 800) -
      /ALTER= ( NAME = RESUMES, -
                SLOT_SIZE=500, -
                WORKING_SET_COUNT = 15)
```

The following command alters caches named FOOD and FOOT (and any other cache with a 4 character name with the first three characters of "FOO" defined in the database):

```
$ RMU /SET ROW_CACHE MF_PERSONNEL -
      /ALTER= ( NAME = FOO%,
                BACKING_STORE_LOCATION=DISK$RDC:[RDC])
```

## 13.1.2 Documentation for the DEBUG\_OPTIONS Qualifier of RMU Unload

Bug 8447357

The RMU Help file and RMU Reference Manual is missing the description of the following qualifier for RMU Unload.

The DEBUG\_OPTIONS qualifier accepts a list of keyword options.

- [NO]TRACE
 

Traces the qualifier and parameter processing performed by RMU Unload. In addition, the query executed to read the table data is annotated with the TRACE statement at each Commit (controlled by Commit\_Every qualifier). When the logical name RDMS\$SET\_FLAGS is defined as "TRACE", then a line similar to the following is output after each commit is performed.

~Xt: 2009-04-23 15:16:16.95: Commit executed.

The default is NOTRACE.

```
$ RMU/UNLOAD/REC=(FILE=WS,FORMAT=CONTROL) SQL$DATABASE WORK_STATUS WS/DEBUG=
TRACE
Debug = TRACE
* Synonyms are not enabled
Row_Count = 500
Message buffer: Len: 13524
Message buffer: Size: 27, Cnt: 500, Use: 4 Flg: 00000000
%RMU-I-DATRECUNL, 3 data records unloaded.
```

- [NO]FILENAME\_ONLY

When the qualifier Record\_Definition=Format:CONTROL is used, the name of the created unload file is written to the control file (.CTL). When the keyword FILENAME\_ONLY is specified, RMU Unload will prune the output file specification to show only the file name and type. The default is NOFILENAME\_ONLY.

```
$ RMU/UNLOAD/REC=(FILE=TT:,FORMAT=CONTROL) SQL$DATABASE WORK_STATUS WS/DEBUG=
FILENAME
--
-- SQL*Loader Control File
--   Generated by: RMU/UNLOAD
--   Version:      Oracle Rdb X7.2-00
--   On:           23-APR-2009 11:12:46.29
--
LOAD DATA
INFILE 'WS.UNL'
APPEND
INTO TABLE "WORK_STATUS"
(
  STATUS_CODE              POSITION(1:1) CHAR NULLIF (RDB$UL_NB1 = '1')
,STATUS_NAME              POSITION(2:9) CHAR NULLIF (RDB$UL_NB2 = '1')
,STATUS_TYPE              POSITION(10:23) CHAR NULLIF (RDB$UL_NB3 = '1')
-- NULL indicators
,RDB$UL_NB1              FILLER POSITION(24:24) CHAR -- indicator for STATUS_CODE
,RDB$UL_NB2              FILLER POSITION(25:25) CHAR -- indicator for STATUS_NAME
,RDB$UL_NB3              FILLER POSITION(26:26) CHAR -- indicator for STATUS_TYPE
)
%RMU-I-DATRECUNL, 3 data records unloaded.
```

- [NO]HEADER

This keyword controls the output of the header in the control file. To suppress the header, use NOHEADER. The default is HEADER.

- APPEND, INSERT, REPLACE, TRUNCATE

These keywords control the text that is output prior to the INTO TABLE clause in the control file. The default is APPEND and only one of these options can be specified.

### 13.1.3 SQL\$MSGxx.DOC Is Not Alphabetical

Bug 4387383

The last paragraph of page A-3 of volume 5 of the SQL Reference Manual says the message codes in files such as SQL\$MSG71.DOC are alphabetized. However, it was found that the message codes were not alphabetized in SQL\$MSG71.DOC or SQL\$MSG72.DOC.

The cause of this problem was that the COSI message codes were appended to the end of the SQL message codes in this file.

This has been corrected in Oracle Rdb Release 7.2.4. We no longer append the COSI message codes to the SQL\$MSGnn.DOC file since the COSI message codes are available separately.

## 13.1.4 LOCK\_TIMEOUT Documentation Error in RMU Reference Manual Release 7.2

The "Oracle Rdb for OpenVMS: Oracle RMU Reference Manual Release 7.2" incorrectly implies that there is a default value for the lock timeout in seconds specified by the /LOCK\_TIMEOUT qualifier in the following sections:

- 1.10 RMU/BACKUP COMMAND
- 1.11 RMU/BACKUP/AFTER\_JOURNAL COMMAND
- 1.17 RMU COPY\_DATABASE COMMAND

In all these sections, in the description of the "/Lock\_Timeout=n" qualifier, any reference to a default value such as "The default value for the /Lock\_Timeout=n qualifier is ..." needs to be removed since there is no default value allowed for this qualifier. If you specify the /LOCK\_TIMEOUT qualifier, you have to specify the lock timeout value in seconds. If you do not specify the /LOCK\_TIMEOUT qualifier, the default is to wait indefinitely to acquire the QUIET POINT lock and any other locks needed for ONLINE execution of the command. It should also be mentioned that the LOCK\_TIMEOUT value does not only affect the QUIET POINT lock but can affect other locks RMU may need to acquire for ONLINE execution.

## 13.1.5 Back Out SQL\_PROTOCOL\_RETRY Configuration Parameter

In older versions of Oracle Rdb, remote connections would retry the connection using the old V5.1 protocol following a failed remote connection attempt. This feature was by default disabled in Oracle Rdb V7.2.2 but has now been enabled again. The configuration file parameter keyword SQL\_PROTOCOL\_RETRY has been removed. This is because no positive effect was seen from this change.

This change was originally documented in the Oracle Rdb 7.2.2 New Features Chapter but has since been removed. More work will be done in this area in the future.

## 13.1.6 Revised Example for SET OPTIMIZATION LEVEL Statement

Bug 6350960

Example 1: Setting the optimization level

The dynamic optimizer can use either FAST FIRST or TOTAL TIME tactics to return rows to the application. The default setting, FAST FIRST, assumes that applications, especially those using interactive SQL, will want to see rows as quickly as possible and possibly abort the query before completion. Therefore, if the FAST FIRST tactic is possible, the optimizer will sacrifice overall retrieval time to initially return rows quickly. This



choice can be affected by setting the OPTIMIZATION LEVEL.

The following example contrasts the query strategies selected when FAST FIRST versus TOTAL TIME is in effect. Databases and queries will vary in their requirements. Queries should be tuned to see which setting best suits the needs of the application environment. For the MF\_PERSONNEL database, there is little or no difference between these tactics but for larger tables the differences could be noticeable.

```
SQL> set flags 'STRATEGY,DETAIL';
SQL> --
SQL> -- No optimization level has been selected. The optimizer
SQL> -- selects the FAST FIRST (FFirst) retrieval tactic to
SQL> -- retrieve the rows from the EMPLOYEES table in the
SQL> -- following query:
SQL> --
SQL> select EMPLOYEE_ID, LAST_NAME
cont> from EMPLOYEES
cont> where EMPLOYEE_ID IN ('00167', '00168');
Tables:
  0 = EMPLOYEES
Leaf#01 FFirst 0:EMPLOYEES Card=100
  Bool: (0.EMPLOYEE_ID = '00167') OR (0.EMPLOYEE_ID = '00168')
  BgrNdx1 EMPLOYEES_HASH [(1:1)2] Fan=1
  Keys: r0: 0.EMPLOYEE_ID = '00168'
        r1: 0.EMPLOYEE_ID = '00167'
EMPLOYEE_ID  LAST_NAME
00167        Kilpatrick
00168        Nash
2 rows selected
SQL> --
SQL> -- Use the SET OPTIMIZATION LEVEL statement to specify that
SQL> -- you want the TOTAL TIME (BgrOnly) retrieval strategy to
SQL> -- be used.
SQL> --
SQL> SET OPTIMIZATION LEVEL 'TOTAL TIME';
SQL> select EMPLOYEE_ID, LAST_NAME
cont> from EMPLOYEES
cont> where EMPLOYEE_ID IN ('00167', '00168');
Tables:
  0 = EMPLOYEES
Leaf#01 BgrOnly 0:EMPLOYEES Card=100
  Bool: (0.EMPLOYEE_ID = '00167') OR (0.EMPLOYEE_ID = '00168')
  BgrNdx1 EMPLOYEES_HASH [(1:1)2] Fan=1
  Keys: r0: 0.EMPLOYEE_ID = '00168'
        r1: 0.EMPLOYEE_ID = '00167'
EMPLOYEE_ID  LAST_NAME
00167        Kilpatrick
00168        Nash
2 rows selected
SQL> --
SQL> -- When the SET OPTIMIZATION LEVEL 'DEFAULT' statement
SQL> -- is specified the session will revert to the default FAST FIRST
SQL> -- optimizer tactic.
SQL> --
SQL> SET OPTIMIZATION LEVEL 'DEFAULT';
SQL> select EMPLOYEE_ID, LAST_NAME
cont> from EMPLOYEES
cont> where EMPLOYEE_ID IN ('00167', '00168');
Tables:
  0 = EMPLOYEES
Leaf#01 FFirst 0:EMPLOYEES Card=100
  Bool: (0.EMPLOYEE_ID = '00167') OR (0.EMPLOYEE_ID = '00168')
```

```

BgrNdX1 EMPLOYEES_HASH [(1:1)2] Fan=1
  Keys: r0: 0.EMPLOYEE_ID = '00168'
        r1: 0.EMPLOYEE_ID = '00167'
EMPLOYEE_ID   LAST_NAME
00167         Kilpatrick
00168         Nash
2 rows selected
SQL>

```

## 13.1.7 RMU /VERIFY Process Quotas and Limits Clarification

When using the RMU/VERIFY command, a process requires a minimum of the following quotas:

- FILLM and CHANNELCNT at least 25 more than the total number of database storage areas, snapshot storage areas, and after image journals.
- Large enough BYTLM, page file quota and working set to open all of the database storage areas, snapshot storage areas, and after image journals.

## 13.1.8 Online Backup Can Be Performed With Transfer Via Memory

The following incorrect Oracle RMU BACKUP command restriction will be removed from the Oracle RMU Reference Manual.

In prior releases of the Oracle RMU Reference Manual, it states under the RMU Backup Online option that "However, an online backup operation cannot be performed if TRANSFER VIA MEMORY, also referred to as optimized page transfer, is enabled. (See the description of the SQL ALTER DATABASE statement in the Oracle Rdb SQL Reference Manual for information on optimized page transfer.)". This restriction is no longer true and will be removed from the Oracle RMU Reference Manual.

The same restriction is also listed for the Online Copy Database and for the Online Move Area commands. This restriction is no longer in place for these commands so it will be removed from the Oracle RMU Reference Manual.

## 13.1.9 Missing Example for CREATE STORAGE MAP

Bug 5655348

The SQL Reference Manual did not include an example showing the storage area attributes for a LIST storage map. The following example will appear in a future version of the Oracle Rdb V7.2 SQL Reference Manual in the CREATE STORAGE MAP section.

### *Example*

The following example shows the use of storage area attributes in a LIST storage map. The storage area attributes must be immediately following the storage area name (as in table storage maps).

## Oracle® Rdb for OpenVMS

```
SQL> create database
cont>     filename 'DB$:MULTIMEDIA'
cont>
cont>     create storage area PHOTO_AREA1
cont>         filename 'DB$:PHOTO_AREA1'
cont>         page format UNIFORM
cont>
cont>     create storage area PHOTO_AREA2
cont>         filename 'DB$:PHOTO_AREA2'
cont>         page format UNIFORM
cont>
cont>     create storage area TEXT_AREA
cont>         filename 'DB$:TEXT_AREA'
cont>         page format UNIFORM
cont>
cont>     create storage area AUDIO_AREA
cont>         filename 'DB$:AUDIO_AREA'
cont>         page format UNIFORM
cont>
cont>     create storage area DATA_AREA
cont>         filename 'DB$:DATA_AREA'
cont>         page format UNIFORM
cont> ;
SQL>
SQL> create table EMPLOYEES
cont>     (name      char(30),
cont>      dob       date,
cont>      ident     integer,
cont>      photograph list of byte varying (4096) as binary,
cont>      resume    list of byte varying (132) as text,
cont>      review    list of byte varying (80) as text,
cont>      voiceprint list of byte varying (4096) as binary
cont>     );
SQL>
SQL> create storage map EMPLOYEES_MAP
cont>     for EMPLOYEES
cont>     enable compression
cont>     store in DATA_AREA:f
SQL>
SQL> create storage map LISTS_MAP
cont>     store lists
cont>         in AUDIO_AREA
cont>             (thresholds are (89, 99, 100)
cont>             ,comment is 'The voice clips'
cont>             ,partition AUDIO_STUFF)
cont>         for (employees.voiceprint)
cont>     in TEXT_AREA
cont>         (thresholds is (99)
cont>         ,partition TEXT_DOCUMENTS)
cont>         for (employees.resume, employees.review)
cont>     in (PHOTO_AREA1
cont>         (comment is 'Happy Smiling Faces?'
cont>         ,threshold is (99)
cont>         ,partition PHOTOGRAPHIC_IMAGES_1)
cont>         ,PHOTO_AREA2
cont>         (comment is 'Happy Smiling Faces?'
cont>         ,threshold is (99)
cont>         ,partition PHOTOGRAPHIC_IMAGES_2)
cont>         )
cont>         for (employees.photograph)
cont>         fill randomly
cont>     in RDB$SYSTEM
```

```

cont>                (partition SYSTEM_LARGE_OBJECTS);
SQL>
SQL> show storage map LISTS_MAP;
LISTS_MAP
For Lists
Store clause:        STORE lists
  in AUDIO_AREA
    (thresholds are (89, 99, 100)
    ,comment is 'The voice clips'
    ,partition AUDIO_STUFF)
  for (employees.voiceprint)
  in TEXT_AREA
    (thresholds is (99)
    ,partition TEXT_DOCUMENTS)
  for (employees.resume, employees.review)
  in (PHOTO_AREAL
    (comment is 'Happy Smiling Faces?'
    ,threshold is (99)
    ,partition PHOTOGRAPHIC_IMAGES_1)
    ,PHOTO_AREA2
    (comment is 'Happy Smiling Faces?'
    ,threshold is (99)
    ,partition PHOTOGRAPHIC_IMAGES_2)
    )
  for (employees.photograph)
  fill randomly
  in RDB$SYSTEM
    (partition SYSTEM_LARGE_OBJECTS)

```

Partition information for lists map:

```

Vertical Partition: VRP_P000
Partition: (1) AUDIO_STUFF
  Fill Randomly
  Storage Area: AUDIO_AREA
  Thresholds are (89, 99, 100)
Comment:      The voice clips
Partition: (2) TEXT_DOCUMENTS
  Fill Randomly
  Storage Area: TEXT_AREA
  Thresholds are (99, 100, 100)
Partition: (3) PHOTOGRAPHIC_IMAGES_1
  Fill Randomly
  Storage Area: PHOTO_AREAL
  Thresholds are (99, 100, 100)
Comment:      Happy Smiling Faces?
Partition: (3) PHOTOGRAPHIC_IMAGES_2
  Storage Area: PHOTO_AREA2
  Thresholds are (99, 100, 100)
Comment:      Happy Smiling Faces?
Partition: (4) SYSTEM_LARGE_OBJECTS
  Fill Randomly
  Storage Area: RDB$SYSTEM

```

```

SQL>
SQL> commit;

```

## 13.1.10 RDM\$BIND\_MAX\_DBR\_COUNT Documentation Clarification

Bugs 1495227 and 3916606

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) for each database.

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 for each database 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.

---

#### Per–Database Value

*The RDM\$BIND\_MAX\_DBR\_COUNT logical name specifies the maximum number of database recovery processes to run at once for each database. For example, if there are 10 databases being recovered and the value for the RDM\$BIND\_MAX\_DBR\_COUNT logical name is 8, up to 80 database recovery processes would be started by the monitor after a node failure.*

---

The RDM\$BIND\_MAX\_DBR\_COUNT logical name is translated when the monitor process opens a database. Databases need to be closed and reopened for a new value of the logical to become effective.

### 13.1.11 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 13–1](#).

*Table 13–1 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.

## 13.1.12 Explanation of SQL\$INT in a SQL Multiversion Environment and How to Redefine SQL\$INT

Bug 2500594

In an environment running multiple versions of Oracle Rdb, for instance Rdb V7.0 and Rdb V7.1, there are now several variant SQL images, such as SQL\$70.EXE and SQL\$71.EXE. However, SQL\$INT.EXE is not variant but acts as a dispatcher using the translation of the logical name RDMS\$VERSION\_VARIANT to activate the correct SQL runtime environment. This image is replaced when a higher version of Oracle Rdb is installed. Thus, using the example above, when Rdb V7.1 is installed, SQL\$INT.EXE will be replaced with the V7.1 SQL\$INT.EXE.

If an application is linked in this environment (using V7.1 SQL\$INT) and the corresponding executable deployed to a system running Oracle Rdb V7.0 multiversion only, the execution of the application may result in the following error:

```
%IMGACT-F-SYMVECMIS, shareable image's symbol vector table mismatch
```

In order to avoid such a problem, the following alternative is suggested:

In the multiversion environment running both Oracle Rdb V7.0 and Oracle Rdb V7.1, run Oracle Rdb V7.0 multiversion by running the command procedures RDB\$SETVER.COM 70 and RDB\$SETVER RESET. This will set up the necessary logical names and symbols that establish the Oracle Rdb V7.0 environment.

For example:

```
$ @SYS$LIBRARY:RDB$SETVER 70
```

```
Current PROCESS Oracle Rdb environment is version V7.0-63 (MULTIVERSION)
Current PROCESS SQL environment is version V7.0-63 (MULTIVERSION)
Current PROCESS Rdb/Dispatch environment is version V7.0-63 (MULTIVERSION)
```

## 13.1.12 Explanation of SQL\$INT in a SQL Multiversion Environment and How to Redefine SQL\$INT

```
$ @SYS$LIBRARY:RDB$SERVER RESET
```

Now run SQL and verify that the version is correct:

```
$ sql$
SQL> show version
Current version of SQL is: Oracle Rdb SQL V7.0-63
```

Define SQL\$INT to point to the variant SQL\$SHR.EXE. Then, create an options file directing the linker to link with this newly defined SQL\$INT. An example follows:

```
$ DEFINE SQL$INT SYS$SHARE:SQL$SHR'RDMS$VERSION_VARIANT'.EXE
$ LINK TEST_APPL,SQL$USER/LIB,SYS$INPUT/option
SQL$INT/SHARE
^Z
```

The executable is now ready to be deployed to the Oracle Rdb V7.0 multiversion environment and should run successfully.

Please note that with each release of Oracle Rdb, new entry points are added to the SQL\$INT shareable image. This allows the implementation of new functionality. Therefore, applications linked with SQL\$INT from Oracle Rdb V7.1 cannot be run on systems with only Oracle Rdb V7.0 installed. This is because the shareable image does not contain sufficient entry points.

The workaround presented here allows an application to explicitly link with the Oracle Rdb V7.0 version of the image. Such applications are upward compatible and will run on Oracle Rdb V7.0 and Oracle Rdb V7.1. The applications should be compiled and linked under the lowest version.

In environments where Oracle Rdb V7.1 is installed, this workaround is not required because the SQL\$INT image will dynamically activate the appropriate SQL\$SHRxx image as expected.

### 13.1.13 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.

## 13.1.14 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.



## 13.2 Address and Phone Number Correction for Documentation

In release 7.0 or earlier documentation, the address and fax phone number listed on the Send Us Your Comments page are incorrect. The correct information is:

FAX -- 603.897.3825  
Oracle Corporation  
One Oracle Drive  
Nashua, NH 03062-2804  
USA

## 13.3 Online Document Format and Ordering Information

You can view the documentation in Adobe Acrobat format using the Acrobat Reader, which allows anyone to view, navigate, and print documents in the Adobe Portable Document Format (PDF). See <http://www.adobe.com> for information about obtaining a free copy of Acrobat Reader and for information on supported platforms.

The Oracle Rdb documentation in Adobe Acrobat format is available on MetaLink:

Top Tech Docs\Oracle Rdb\Documentation\`<bookname>`

Customers should contact their Oracle representative to purchase printed documentation.

---

# Chapter 14

## Known Problems and Restrictions

This chapter describes problems and restrictions relating to Oracle Rdb and includes workarounds where appropriate.

# 14.1 Known Problems and Restrictions in All Interfaces

This section describes known problems and restrictions that affect all interfaces. This is not an exhaustive list. Check the Oracle Bug database to see a list of all open Rdb bugs and their current status.

## 14.1.1 Possible Incorrect Results When Using Partitioned Descending Indexes

Bug 6129797

In the current release of Oracle Rdb, 7.2.4, it is possible for some queries using partitioned indexes with segments of mixed ascending and descending order to return incorrect results either on Alpha or I64 systems.

The following examples show two problems when using partitioned index with segments of mixed ascending and descending order:

```
create database file foo
  create storage area fooa
  create storage area foob;

create table mesa (id integer, m4 char (1), m5 integer);
create table rasa (id integer, r4 char (1), r5 integer);

insert into mesa (id, m4, m5) values (1, 'm', 1 );

insert into rasa (id, r4, r5) values (1, 'm', 1 );
insert into rasa (id, r4, r5) values (1, 'k', 1 );
insert into rasa (id, r4, r5) values (1, 'e', 1 );

create index x4 on mesa (id asc , m4 asc ) ;

! The following index contains ascending segments followed by descending
! segments and thus causes the query to return the wrong result.
!
! Note that the query works if both segments are either ascending or descending.
!
create index y4 on rasa (id asc , r4 desc)
  store using (id, r4)
  in fooa with limit of (1, 'g' )
  otherwise in foob ;
commit;

! Problem #1:
!
! the following query returns correctly 3 rows on Alpha but 1 row on IA64:

SQL> select m.id, m.m4, r.r4 from
  mesa m inner join rasa r on (m.id = r.id);
      1  m      m
      1  m      k
      1  m      e
3 rows selected

SQL> select m.id, m.m4, r.r4 from mesa m inner join rasa r on (m.id = r.id);
```

```

          1   m       e
1 row selected

! Problem #2:
!
! The following query using reverse scan returns 2 rows incorrectly on Alpha
! but 3 rows correctly on IA64:
!

SQL> select id, r4 from rasa where id = 1 and r4 <= 'm' order by id, r4;
Tables:
  0 = RASA
Index only retrieval of relation 0:RASA
  Index name  Y4 [2:1]   Reverse Scan
  Keys: (0.ID = 1) AND (0.R4 <= 'm')
        ID   R4
        1   k
        1   m
2 rows selected

SQL> select id, r4 from rasa where id = 1 and r4 <= 'm' order by id, r4;
Tables:
  0 = RASA
Index only retrieval of relation 0:RASA
  Index name  Y4 [2:1]   Reverse Scan
  Keys: (0.ID = 1) AND (0.R4 <= 'm')
        ID   R4
        1   e
        1   k
        1   m
3 rows selected

```

This problem is related to the construction and comparison of the descending key values during the index partitions.

The problem will be corrected in a future version of Oracle Rdb.

## 14.1.2 Remote Attach Stalls Before Detecting a Node is Unreachable

Bug 7681548

A remote attach can stall for a noticeable period, typically 75 seconds, before detecting a node is unreachable.

The following example shows the expected error message when attempting to access a database on a node that is not reachable. The problem is that when the value of the parameter `SQL_NETWORK_TRANSPORT_TYPE` in the file `RDB$CLIENT_DEFAULTS.DAT` is not specifically set to `DECNET` (in UPPER CASE), a stall of typically 75 seconds will happen before you get the expected error message.

```

SQL> attach 'file 1::disk1:[dbdir]db';
%SQL-F-ERRATTDEC, Error attaching to database 1::disk1:[dbdir]db
-RDB-F-IO_ERROR, input or output error
-SYSTEM-F-UNREACHABLE, remote node is not currently reachable

```

There are two possible ways to avoid the stall and get the error message after a user configurable period of time or instantly: decrease the value of the TCPIP parameter TCP\_KEEPINIT, or explicitly specify SQL\_NETWORK\_TRANSPORT\_TYPE as DECNET (in UPPER CASE).

- The default behavior when attempting to connect to an unreachable node via TCPIP is to stall 75 seconds before returning an error. The stall time is configurable in TCPIP via the parameter TCP\_KEEPINIT which is expressed in units of 500 ms. The default value of TCP\_KEEPINIT is 150 which corresponds to a 75 second stall.
- When connecting via DECnet, the error message is typically returned instantly so a significant stall will not be seen in this case. However, the value of the parameter SQL\_NETWORK\_TRANSPORT\_TYPE is case sensitive so for DECnet to be selected as the transport, "DECNET" must be specified in UPPER CASE. Failing to do so will result in connecting via the DEFAULT method which is to first try connecting via DECnet and if that fails attempt to connect via TCPIP and hence a 75 second stall will take place unless TPC\_KEEPINIT is set to a value lower than 150.

### 14.1.3 Case Sensitive Values in RDB\$CLIENT\_DEFAULTS.DAT

Bug 7681548

Various characteristics for network access to remote databases can be specified by entering parameters and values in a file named RDB\$CLIENT\_DEFAULTS.DAT. The following keywords that have character strings as their values only take effect if the values are specified in UPPER CASE:

SQL\_NETWORK\_TRANSPORT\_TYPE, SQL\_MESSAGE\_VECTOR\_RETURN\_TYPE, and SQL\_DEFAULTS\_RESTRICTION. The result of including one or more lower case characters in the value of one of these parameters is the same as if the parameter was not specified at all (for example, the default behavior would be applied and no error message would be issued).

In the following example, DECnet is specified with the last three characters in lower case. The result will be that the value of the parameter SQL\_NETWORK\_TRANSPORT\_TYPE will be DEFAULT and not the intended value DECNET.

```
SQL_NETWORK_TRANSPORT_TYPE DECnet
```

This problem can be avoided by specifying the values for SQL\_NETWORK\_TRANSPORT\_TYPE, SQL\_MESSAGE\_VECTOR\_RETURN\_TYPE, and SQL\_DEFAULTS\_RESTRICTION in RDB\$CLIENT\_DEFAULTS.DAT using UPPER CASE.

In the next major release of Oracle Rdb, the values in RDB\$CLIENT\_DEFAULTS.DAT will be case insensitive.

### 14.1.4 Standalone WITH Clause in Compound Statements Now Deprecated

In prior versions of Oracle Rdb, it was permitted to follow the BEGIN keyword in a top level compound statement or stored routine with a WITH HOLD clause to specify that the procedure treated all FOR loops as HOLD cursors.

Unfortunately, this syntax conflicts with recent syntax added to the ANSI and ISO SQL Database Language standard. Support for this new syntax (known as subquery factoring) is used by Oracle tools accessing Oracle Rdb through OCI Services for Rdb. Therefore, to accommodate this change Oracle will remove the WITH HOLD syntax as a standalone clause after the BEGIN keyword. The alternate syntax, available since Oracle Rdb V7.1, is to use the PRAGMA clause which allows the WITH HOLD clause to be specified.

Any applications or SQL command files must be modified prior to the next major release of Oracle Rdb. At that time, the old syntax will be removed and the new WITH clause (aka subquery factoring) will be introduced.

The following example shows the old syntax, which now produces a deprecated message.

```
SQL> begin
cont> with hold preserve none
%SQL-I-DEPR_FEATURE, Deprecated Feature: WITH HOLD no longer supported in
this context - use PRAGMA (WITH HOLD) instead
cont> trace 'a';
cont> end;
```

It should be replaced with the following syntax which provides the same behavior.

```
SQL> begin
cont> pragma (with hold preserve none)
cont> trace 'a';
cont> end;
```

## 14.1.5 Calling DECC\$CRTL\_INIT

In cases where user-supplied code is being called by Oracle Rdb (such as an external function, a module called implementing the Oracle Backup API, or a user-supplied output routine for the Oracle Rdb LogMiner), if calls are made to certain DECC\$SHR RTL routines, it may be required to first call DECC\$CRTL\_INIT.

DECC\$CRTL\_INIT is a C run time library routine that allows developers to call the HP C RTL from other languages or to use the HP C RTL when the main function is not in C. It initializes the run-time environment and establishes both an exit and condition handler. The Oracle Rdb main images are not written in C and should not be expected to have called DECC\$CRTL\_INIT prior to the user's code being invoked. The requirement for DECC\$CRTL\_INIT in certain cases exists in all versions of Oracle Rdb.

A shareable image need only call this function if it contains an HP C function call for signal handling, environment variables, I/O, exit handling, a default file protection mask, or if it is a child process that should inherit context. Although many of the initialization activities are performed only once, DECC\$CRTL\_INIT can safely be called multiple times.

See the HP C Run-Time Library Reference Manual for OpenVMS Systems manual for additional information.

## 14.1.6 Application and Oracle Rdb Both Using SYS\$HIBER

In application processes that use Oracle Rdb and the SYS\$HIBER system service (possibly via RTL routines such as LIB\$WAIT), it is very important that the application ensures that the event being waited for has

actually occurred. Oracle Rdb utilizes \$HIBER/\$WAKE sequences for interprocess communication and synchronization.

Because there is just a single process-wide "hibernate" state along with a single process-wide "wake pending" flag, Oracle Rdb must assume that it "shares" use of the hibernate/wake state with the user's application code. To this end, Oracle Rdb generally will re-wake the process via a pending wake request after using a hibernate sequence.

Oracle Rdb's use of the \$WAKE system service will 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, applications that use HIBER/WAKE facilities must 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 includes a 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.

In order to prevent application hangs, inner-mode users of SYS\$HIBER must take explicit steps to ensure that a pending wake is not errantly "consumed". The general way of accomplishing this is to issue a SYS\$WAKE to the process after the event is complete if a call to SYS\$HIBER was done. Rdb takes this step and therefore application programs must be prepared for cases where a wakeup might appear unexpectedly.



## 14.1.7 Unexpected RCS Termination

It has been observed in internal testing of Rdb Release 7.2.2 that if the Record Cache Server (the RCS) terminates in an uncontrolled fashion this may, under some conditions, cause corruption of the database and/or the After Image Journal file.

When the RCS terminates, the database is shut down and a message like the following is written to the monitor log:

```
6-DEC-2007 15:04:17.02 - Received Record Cache Server image termination from
22ED5144:1
- database name "device:[directory]database.RDB:1" [device] (1200,487,0)
- abnormal Record Cache Server termination detected
- starting delete-process shutdown of database:
  - %RDMS-F-RCSABORTED, record cache server process terminated abnormally
- sending process deletion to process 22ED10F9
- sending process deletion to process 22ECED59
- sending process deletion to process 22EC0158
- sending process deletion to process 22EB9543 (AIJ Log server)
- database shutdown waiting for active users to terminate
```

A future attempt to roll forward the AIJ following a restore of a database backup might fail with a bugcheck dump if this problem has happened.

The only currently known situation where this problem has been observed is if the logical name RDM\$BIND\_RCS\_VALIDATE\_SECS is defined to some value and the logical name RDM\$BIND\_RCS\_LOG\_FILE at the same time is undefined or defined incorrectly.

To prevent this problem, Oracle recommends any customer using the Row Cache feature to either avoid defining the logical name RDM\$BIND\_RCS\_VALIDATE\_SECS or if this logical name for any reason needs to be defined, to make sure RDM\$BIND\_RCS\_LOG\_FILE is correctly defined (i.e. defined with the /SYSTEM and /EXECUTIVE qualifiers and is pointing to a valid file name in an existing directory on a cluster accessible device with sufficient free space).

This recommendation applies to all versions of Oracle Rdb.

## 14.1.8 Possible Incorrect Results When Using Partitioned Descending Indexes on I64

When running on I64 systems using Rdb Release 7.2, it is possible when using partitioned descending indexes for some queries to return incorrect results. Alpha systems are not effected by this problem.

The following example shows this difference in behavior between Alpha and I64 when using partitioned descending indexes:

```
SQL> CREATE DATABASE FILE FOO
cont>         CREATE STORAGE AREA FOOA
cont>         CREATE STORAGE AREA FOOB;
SQL>
SQL> CREATE TABLE MESA (ID INTEGER, M4 CHAR (1), M5 INTEGER);
SQL> CREATE TABLE RASA (ID INTEGER, R4 CHAR (1), R5 INTEGER);
SQL>
SQL> INSERT INTO MESA (ID, M4, M5) VALUES (1, 'M', 1 );
```

```

1 row inserted
SQL> INSERT INTO RASA (ID, R4, R5) VALUES (1, 'M', 1 );
1 row inserted
SQL>
SQL> CREATE INDEX X4 ON MESA (ID ASC , M4 DESC)
cont>         STORE USING (ID, M4)
cont>         IN FOOA WITH LIMIT OF (1, 'G')
cont>         OTHERWISE IN FOOB ;
SQL>
SQL> CREATE INDEX Y4 ON RASA (ID ASC , R4 DESC)
cont>         STORE USING (ID, R4)
cont>         IN FOOA WITH LIMIT OF (1, 'G' )
cont>         OTHERWISE IN FOOB ;
SQL>
SQL> COMMIT;

! This query correctly returns 1 row
! on Alpha but returns 0 rows on I64:

SQL> SELECT M.ID, M.M4, R.R4 FROM
cont> MESA M INNER JOIN RASA R ON (M.ID = R.ID);
0 rows selected
SQL>

```

This problem is related to the construction and comparison of the descending key values with Oracle Rdb running on I64. This problem will be corrected in a future Rdb 72 release.

## 14.1.9 Response 'QUIT' to RMU Restart Prompt Loops

Bug 6001187

On Itanium, responding to a restart prompt with 'QUIT' causes the RMU command to loop creating bugchecks.

This problem does not appear on Alpha and shall be fixed in a future release on Itanium.

## 14.1.10 Changes for Processing Existence Logical Names

This release of Oracle Rdb will change the handling of so called "existence" logical names used to tune the Rdb environment. These existence logical names could in past versions be defined to any value to enable their effect. The Rdb documentation in most cases described using the value 1 or YES as that value and this change is upward compatible with the documentation.

Rdb now treats these logical names (see the list below) as Boolean logicals and accepts a string starting with "Y", "y", "T", "t" or "1" to mean TRUE. All other values will be considered to be FALSE. This change allows process level definitions to override definitions in higher logical name tables which was not possible previously.

Oracle recommends that customers examine all procedures that define the following logical names to ensure that their values conform to these rules prior to upgrading to Oracle Rdb V7.2.1.1 or later to avoid unexpected changes in behavior.

- RDMS\$AUTO\_READY
- RDMS\$DISABLE\_HIDDEN\_KEY

- RDMS\$DISABLE\_MAX\_SOLUTION
- RDMS\$DISABLE\_REVERSE\_SCAN
- RDMS\$DISABLE\_TRANSITIVITY
- RDMS\$DISABLE\_ZIGZAG\_BOOLEAN
- RDMS\$ENABLE\_BITMAPPED\_SCAN
- RDMS\$ENABLE\_INDEX\_COLUMN\_GROUP
- RDMS\$MAX\_STABILITY
- RDMS\$USE\_OLD\_COST\_MODEL
- RDMS\$USE\_OLD\_COUNT\_RELATION
- RDMS\$USE\_OLD\_SEGMENTED\_STRING
- RDMS\$USE\_OLD\_UPDATE\_RULES

## 14.1.11 Patch Required When Using VMS V8.3 and Dedicated CPU Lock Manager

During qualification testing of Oracle Rdb Release 7.2.1 on OpenVMS V8.3 systems, a problem with the use of Extended Lock Value Blocks and the OpenVMS Dedicated CPU Lock Manager feature was discovered.

To avoid this problem, Oracle strongly recommends that customers wishing to use Oracle Rdb and the OpenVMS Dedicated CPU Lock Manager feature with OpenVMS V8.3 install one of the following architecture-specific patch kit (or subsequent replacement if superseded) prior to using Oracle Rdb Release 7.2.1 on OpenVMS V8.3 systems:

- VMS83I\_SYS-V0200 (I64)
- VMS83A\_SYS-V0100 (Alpha)

## 14.1.12 SQL Module or Program Fails with %SQL-F-IGNCASE\_BAD

Bug 2351258

A SQL Module or Pre-compiled SQL program built with Rdb 6.1 or earlier may fail when running under Rdb 7.2 if the program submits queries that involve certain kinds of character operations on parameters in the queries. For example, a LIKE operator in the WHERE clause of a SQL statement requires SQL to look for character- or string-matching wildcard characters. Another example is the use of IGNORE CASE which causes SQL to equivalence upper and lower case characters for the character set in use.

The following example shows a portion of a SQL module language program that queries a PERSONNEL database.

```
DECLARE MANL_NAME_LIST CURSOR FOR
  SELECT DISTINCT E.LAST_NAME, E.FIRST_NAME, J.JOB_CODE, J.DEPARTMENT_CODE, E.CITY
FROM    DB1_HANDLE.EMPLOYEES E, DB1_HANDLE.JOB_HISTORY J
WHERE J.EMPLOYEE_ID = E.EMPLOYEE_ID
      AND E.STATUS_CODE = STATUS_CODE
      AND E.CITY LIKE CITYKEY IGNORE CASE
ORDER BY E.EMPLOYEE_ID DESC, E.LAST_NAME DESC
```

```
PROCEDURE SQL_OPN_NAME_LIST
SQLCODE
CITYKEY          CHAR(20)
STATUS_CODE     CHAR(1);
OPEN MANL_NAME_LIST;
```

If the SQL Module containing the code above is compiled and linked into an executable using a pre-7.0 version of Rdb, it will run properly against that version. However if the same program is run in an Rdb 7.2 environment, a call to the SQL\_OPN\_NAME\_LIST procedure will return a SQLCODE of -1. The RDB\$MESSAGE\_VECTOR will contain a code associated with the following message:

```
%SQL-F-IGNCASE_BAD, IGNORE CASE not supported for character set
```

To workaround this problem, re-link the program using a 7.2 version of SQL\$INT.EXE and/or SQL\$USER.OLB.

### 14.1.13 External Routine Images Linked with PTHREAD\$RTL

The OpenVMS Guide to the POSIX Threads Library describes that it is not supported to dynamically activate the core run-time library shareable image PTHREAD\$RTL. Oracle has found in testing that a shareable image supplied for use as an External Routine that is linked with PTHREAD\$RTL can be expected to cause a hang during dynamic image activation on OpenVMS I64 systems. This problem has not been observed on OpenVMS Alpha systems.

To avoid this problem in any case where the shareable image used for an Rdb External Routine is linked with PTHREAD\$RTL, the main program image must likewise be linked with PTHREAD\$RTL. This requirement applies to customer built application main programs as well as the main interactive SQL image.

The shareable image RDB\$NATCONN\_FUNC72.EXE supplied with OCI Services for Oracle Rdb (part of SQL/Services) is one such shareable image that is linked with PTHREAD\$RTL. Customer built applications that utilize External Routines from the RDB\$NATCONN\_FUNC72.EXE image must ensure that the main image is linked with PTHREAD\$RTL. The external routines that a user may call that use functions from RDB\$NATCONN\_FUNC72.EXE include:

- TO\_CHAR
- TO\_NUMBER
- TO\_DATE

You can use the OpenVMS command ANALYZE/IMAGE to determine whether an image depends upon PTHREAD\$RTL. For more information, see the OpenVMS documentation.

### 14.1.14 SQL Procedure External Location Should Be Upper Case

Bug 4722422

When using External Routines, it is important that all declarations for the same shareable image use the exact same strings for the image file specification. Failure to use the same string content may result in multiple copies of the image being activated or failure to correctly call the external routine.

The "ALTER FUNCTION ... LOCATION" command can be used to alter the existing function location string without having to drop and recreate the function.

The following example shows the same string for the EXTERNAL LOCATION specifications:

```
create procedure sys$asctim(
    out :timlen smallint by reference,
    out :timbuf char(23) by descriptor,
    in  :timadr date vms by reference,
    in  :cvtflag integer by value);
    external location 'SYS$SHARE:SYS$PUBLIC_VECTORS.EXE'
    language general general parameter style;

create procedure sys$gettim(
    in  :timadr date vms by reference);
    external location 'SYS$SHARE:SYS$PUBLIC_VECTORS.EXE'
    language general general parameter style;
```

## 14.1.15 Using Databases from Releases Earlier than V7.0

You cannot convert or restore databases earlier than the Oracle Rdb V7.0 format directly to Oracle Rdb V7.2 format. The RMU Convert command for Oracle Rdb V7.2 supports conversions from Oracle Rdb V7.0 and V7.1 format databases only. If you have an Oracle Rdb V3.0 through V6.1 format database, you must convert it to at least Oracle Rdb V7.0 format and then convert it to Oracle Rdb V7.2 format. For example, if you have a V4.2 format database, you must convert it first to at least Oracle Rdb V7.0 format, then convert it to Oracle Rdb V7.2 format.

If you attempt to convert or restore a database that is prior to Oracle Rdb V7.0 format directly to Oracle Rdb V7.2 format, Oracle RMU generates an error.

## 14.1.16 Partitioned Index with Descending Column and Collating Sequence

Bug 2797443

A known problem exists in which a query can return wrong results (number of rows returned is incorrect). This can happen on a table that has a multi-column, partitioned index in which one of the columns is sorted in descending order and the column has an associated collating sequence.

The following example can be used to demonstrate the problem.

```
$ sql$
create database file mf_collating.rdb alloc 10
    collating sequence french french
    create storage area area1 alloc 10
    create storage area area2 alloc 10
    create storage area area3 alloc 10;
create table tabl (id tinyint, r3 char (3));
insert into tabl (id, r3) values (1, 'a');
insert into tabl (id, r3) values (1, 'b');
insert into tabl (id, r3) values (1, 'f');
create index y3 on tabl (id asc, r3 desc)
    store using (id, r3)
```

```

in area1 with limit of (1, 'k')
in area2 with limit of (1, 'e')
otherwise in area3 ;
commit;

set flags 'strategy';

! Here is a query that returns the correct rows using sequential rather
! than indexed access.

select id, r3 from tabl where id = 1 and r3 <= 'e'
  optimize for sequential access;
Conjunct          Get          Retrieval sequentially of relation TAB1
  ID    R3
    1    a
    1    b
2 rows selected

! Here is the same query without the sequential access restriction.
! Note in the query strategy that index Y3 is used for data retrieval.
! This query ought to (but does not) return the same set of rows as
! for the sequential access query.

select id, r3 from tabl where id = 1 and r3 <= 'e';
Leaf#01 FFirst TAB1 Card=3
  BgrNdx1 Y3 [2:1] Fan=16
0 rows selected

```

## 14.1.17 Domain–Qualified TCP/IP Node Names in Distributed Transactions

Bug 3735144

When using TCP/IP for Oracle Rdb remote connections, distributed transactions involving databases on nodes which are not on the same subnet may not work.

Remote Rdb has the capability to make remote connections via TCP/IP in lieu of DECnet. (See the Oracle Rdb OpenVMS Installation and Configuration Guide for how to set this up.) However, distributed transactions involving remote databases connected to via TCP/IP have been difficult. This is because Rdb relies on OpenVMS DECdtm for distributed transaction support and DECdtm requires DECnet for off–node communication. (This is an OpenVMS and not an Rdb restriction. Contact Hewlett–Packard OpenVMS Support for more details.)

OpenVMS provides a capability to run DECnet over TCP/IP so that OpenVMS services which require DECnet (like DECdtm) can operate in an environment where a TCP/IP network is used as the communications backbone. This capability allows DECdtm (and hence Rdb) to manage distributed transactions via TCP/IP. (See HP's OpenVMS DECnet–Plus documentation set for how to configure and use this capability.)

However, for a transaction involving a remote database, Rdb only provides the SCSNODE name of the remote node to DECdtm. For example, consider the following SQL attaches to two remote databases using TCP/IP:

```
SQL> attach 'alias db1 filename node1.a.b.c::db_root:db1 user 'me' using
```

```

''pw'';
SQL> attach 'alias db2 filename node1.a.b.c::db_root:db2 user ''me'' using
''pw'';

```

In the above example, Rdb can successfully connect to both remote databases using the TCP/IP address "node1.a.b.c." but when multiple databases are attached, Rdb implicitly uses distributed transactions via DECdtm. Since Rdb only passes DECdtm the SCSNODE name retrieved from the RDBSERVERnn at the other end of the connection, DECdtm does not, in general, have the information it needs to resolve the remote reference. It will only be able to do so if the SCSNODE name and the TCP/IP node name are the same and the local node is on the same subnet (i.e. ".a.b.c" in the example). Otherwise, after the second attach is made, the following error message will be received as soon as a transaction is started:

```

SQL> set trans read write;
%RDB-F-SYS_REQUEST_CAL, error from system services request - called from 100001
-RDB-E-DECDTMERR, DECdtm system service call error
-IPC-E-BCKTRNSFAIL, failure on the back translate address request

```

There are three potential workarounds:

- If distributed transactions are unimportant to the application, they can be disabled by defining the logical name `SQL$DISABLE_CONTEXT` to `TRUE`. Rdb will then not call DECdtm and the node name resolution problem will not be seen. However, it will be the problem of the application to maintain database integrity in the event that a commit succeeds on one database and not on another. See the Rdb Guide to Distributed Transactions for more information.
- If all the nodes involved in the distributed transaction are in the same domain, then TCP/IP can resolve the node with only the first part of the node provided that the SCSNODE name is identical to it. In the example above, this would mean that the remote node had an SCSNODE name of "NODE1" and that the local node was on TCP/IP subnet ".a.b.c".
- It may also be possible to define a DNS/BIND alias name for the remote node's SCSNODE name to the local node's TCP/IP database. This should allow the SCSNODE name passed by Rdb Dispatch to be translated successfully. For example, assuming HP TCP/IP Services for OpenVMS is the TCP/IP protocol stack then a command like the following could be used on the local node:

```
$ TCP SET HOST NODE1.A.B.C/address=nnn.nnn.nnn.nnn/alias=NODE1_SCS
```

Where "nnn.nnn.nnn.nnn" is the IP address and "NODE1\_SC" the OpenVMS SCSNODE name of the remote node. See the HP DECnet-Plus documentation set for more information on how to maintain TCP/IP domain databases.

## 14.1.18 ILINK-E-INVORINI Error on I64

When linking an application with multiple modules, the following error message may be returned:

```

%ILINK-E-INVORINI, incompatible multiple initializations for overlaid section
  section: VMSRDB
  module: M1
  file: DKA0:[BLD]M1.OBJ;1
  module: M2
  file: DKA0:[BLD]SYS.OLB;1

```

On I64 systems, it is not allowed to have a program section that attempts to be initialized a subsequent time where the non-zero portions of the initializations do not match. This is a difference from OpenVMS Alpha and VAX systems where the linker permitted such initializations.

If the modules specified are SQL module language or precompiler produced, the application build procedures usually need to be modified. Typically, the solution is to initialize the database handles in only one of the modules. The SQLMOD command line qualifiers `/NOINITIALIZE_HANDLES` and `/INITIALIZE_HANDLES` are used to specify whether or not alias definitions are coerced into alias references.

## 14.1.19 New Attributes Saved by RMU/LOAD Incompatible With Prior Versions

Bug 2676851

To improve the behavior of unloading views, Oracle Rdb Release 7.1.2 changed the way view columns were unloaded so that attributes for view computed columns, `COMPUTED BY` and `AUTOMATIC` columns were saved. These new attributes are not accepted by prior releases of Oracle Rdb.

The following example shows the reported error trying to load a file from V7.1.2 under V7.1.0.4.

```
%RMU-F-NOTUNLFILE, Input file was not created by RMU UNLOAD
%RMU-I-DATRECSTO,    0 data records stored.
%RMU-F-FTL_LOAD, Fatal error for LOAD operation at 21-OCT-2003 16:34:54.20
```

You can work around this problem by using the `/RECORD_DEFINITION` qualifier and specifying the `FORMAT=DELIMITED` option. However, this technique does not support `LIST OF BYTE VARYING` column unloading.

## 14.1.20 SYSTEM-F-INSFMEM Fatal Error With SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED in Galaxy Environment

When using the `GALAXY SUPPORT IS ENABLED` feature in an OpenVMS Galaxy environment, a `%SYSTEM-F-INSFMEM`, *insufficient dynamic memory error* may be returned when mapping record caches or opening the database. One source of this problem specific to a Galaxy configuration is running out of Galaxy Shared Memory regions. For Galaxy systems, `GLX_SHM_REG` is the number of shared memory region structures configured into the Galaxy Management Database (GMDB).

While the default value (for OpenVMS versions through at least V7.3-1) of 64 regions might be adequate for some installations, sites using a larger number of databases or row caches when the `SHARED MEMORY IS SYSTEM` or `LARGE MEMORY IS ENABLED` features are enabled may find the default insufficient.

If a `%SYSTEM-F-INSFMEM`, *insufficient dynamic memory error* is returned when mapping record caches or opening databases, Oracle Corporation recommends that you increase the `GLX_SHM_REG` parameter by 2 times the sum of the number of row caches and number of databases that might be accessed in the Galaxy at one time. As the Galaxy shared memory region structures are not very large, setting this parameter to a higher than required value does not consume a significant amount of physical memory. It also may avoid a later reboot of the Galaxy environment. This parameter must be set on all nodes in the Galaxy.



## Galaxy Reboot Required

---

*Changing the `GLX_SHM_REG` system parameter requires that the OpenVMS Galaxy environment be booted from scratch. That is, all nodes in the Galaxy must be shut down and then the Galaxy reformed by starting each instance.*

---

## 14.1.21 Oracle Rdb and OpenVMS ODS-5 Volumes

OpenVMS Version 7.2 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.

## 14.1.22 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.

For 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 alter column f2 primary key not deferrable;
SQL> alter table t1 alter column f1 references t2 not deferrable;
```

When deleting from the PRIMARY KEY table, 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.

```
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 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 alter column f1
cont>  check (f1 in (select * from t2 where f2=f1)) not deferrable;
```

or:

```
SQL> alter table t1 alter column f1
cont>  check (f1=(select * from t2 where f2=f1)) not deferrable;
```

In both cases the retrieval strategy will look like this:

```
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 [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 after insert on t1
cont> when (not exists (select * from t2 where f2=f1))
cont> (error) for each row;
SQL> create trigger t1_update 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 before delete on t2
cont> when (exists (select * from t1 where f1=f2))
cont> (error) for each row;
SQL> create trigger t2_modify 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 NULL to indicate the absence of a value, the above triggers could easily be modified to allow for these semantics.

## 14.1.23 Carryover Locks and NOWAIT Transaction Clarification

In NOWAIT transactions, the BLAST (Blocking AST) 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 carryover locks. There can be a delay before the transactions with carryover locks detect the presence of the NOWAIT transaction and release their carryover locks. You can detect this condition by examining the stall messages. If the "Waiting for NOWAIT signal (CW)" stall message appears frequently, the application is probably experiencing a decrease in performance, and you should consider disabling the carryover lock behavior.

## 14.1.24 Unexpected Results Occur During Read-Only Transactions on a Hot Standby Database

When using Hot Standby, it is typical to use the standby database for reporting, simple queries, and other read-only transactions. If you are performing these types of read-only transactions on a standby database, be sure you can tolerate a READ COMMIT level of isolation. This is because the Hot Standby database might be updated by another transaction before the read-only transaction finishes, and the data retrieved might not be what you expected.

Because Hot Standby does not write to the snapshot files, the isolation level achieved on the standby database for any read-only transaction is a READ COMMITTED transaction. This means that nonrepeatable reads and phantom reads are allowed during the read-only transaction:

- **Nonrepeatable read operations:** Allows the return of different results within a single transaction when an SQL operation reads the same row in a table twice. Nonrepeatable reads can occur when another transaction modifies and commits a change to the row between transactions. Because the standby database will update the data when it confirms a transaction has been committed, it is very possible to see an SQL operation on a standby database return different results.
- **Phantom read operations:** Allows the return of different results within a single transaction when an SQL operation retrieves a range of data values (or similar data existence check) twice. Phantoms can occur if another transaction inserted a new record and committed the insertion between executions of the range retrieval. Again, because the standby database may do this, phantom reads are possible.

Thus, you cannot rely on any data read from the standby database to remain unchanged. Be sure your read-only transactions can tolerate a READ COMMIT level of isolation before you implement procedures that read and use data from a standby database.

## 14.1.25 Row Cache Not Allowed While Hot Standby Replication is Active

The row cache feature may not be enabled on a hot standby database while replication is active. The hot standby feature will not start if row cache is enabled.

This restriction exists because rows in the row cache are accessed via logical dbkeys. However, information transferred to the standby database via the after image journal facility only contains physical dbkeys. Because there is no way to maintain rows in the cache via the hot standby processing, the row cache must be disabled when the standby database is open and replication is active.

A new command qualifier, ROW\_CACHE=DISABLED, has been added to the RMU Open command. To open the hot standby database prior to starting replication, use the ROW\_CACHE=DISABLED qualifier on the RMU Open command.

## 14.1.26 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. 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 a 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 shareable image. SQL IMPORT and RMU Load operations do, however, call the OpenVMS SORT run-time library.

At the beginning of a sort operation, the SORT code allocates memory for working space. The SORT code uses this space for buffers, in-memory copies of the data, and sorting trees.

SORT does not directly consider the processes 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 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 the new pages are faulted in. Once the sort operation completes and SORT returns back to Oracle Rdb, 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. Since that 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 for Oracle Rdb sort operations. When the operation can not be done in the available memory, SORT uses temporary disk files to hold the data as it is being sorted. The Oracle Rdb7 Guide to Database Performance and Tuning contains more detailed information about sort work files.

The logical name RDMS\$BIND\_SORT\_WORKFILES specifies how many work files sort is to use if work files are required. The default is 2 and the maximum number is 36. The work files can be individually controlled by the SORTWORKn logical names (where n ranges from "0" through "Z"). 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 36 logical names, "SORTWORK0" through "SORTWORKZ".

Normally, SORT places work files in the your SYSSCRATCH directory. By default, SYSSCRATCH is the same device and directory as the SYS\$LOGIN location. Spreading the I/O load over multiple disks and/or controllers improves efficiency as well as performance by taking advantage of more system resources and helps prevent disk I/O bottlenecks. Specifying that a your work files 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 (for example, while Oracle Rdb builds indexes for a very large table). Using the "SORTWORK0" through "SORTWORKZ" logical names can help you avoid this problem.

Note that SORT 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 36 sort work files, it is possible to exceed the capacity of the first SORT file device and the sort operation fails never having accessed the remaining 35 sort work files.

At this time, more than 10 sort work files will only be used by the Oracle Rdb sort interface as used by the CREATE INDEX, ALTER INDEX and the clauses UNION DISTINCT, ORDER BY, GROUP BY and SELECT DISTINCT. The RMU and SQL IMPORT interfaces use the OpenVMS SORT interface which does not currently support more than 10 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 allocation within Oracle Rdb.

## 14.1.27 Control of Sort Work Memory Allocation

Oracle Rdb uses a built-in SORT32 package to perform many sort operations. Sometimes, these sorts exhibit a significant performance problem when initializing work memory to be used for the sort. This behavior can be experienced, for example, when a very large sort cardinality is estimated, but the actual sort cardinality is small.

In rare cases, it may be desirable to artificially limit the sort package's use of work memory. Two logicals have been created to allow this control. In general, there should be no need to use either of these logicals and misuse of them can significantly impact sort performance. Oracle recommends that these logicals be used carefully and sparingly.

The logical names are:

*Table 14–1 Sort Memory Logicals*

Logical	Definition
RDMS\$BIND_SORT_MEMORY_WS_FACTOR	Specifies a percentage of the process's working set limit to be used when allocating sort memory for the built-in SORT32 package. If not defined, the default value is 75 (representing 75%), the maximum value is 75 (representing 75%), and the minimum value is 2 (representing 2%). Processes with vary large working set limits can sometimes experience significant page faulting and CPU consumption while initializing sort memory. This logical name can restrict the sort work memory to a percentage of the processes maximum working set.

RDMS\$BIND_SORT_MEMORY_MAX_BYTES	Specifies an absolute limit to be used when allocating sort memory for the built-in SORT32 package. If not defined, the default value is unlimited (up to 1GB), the maximum value is 2147483647 and the minimum value is 32768.
----------------------------------	---

## 14.1.28 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 columns 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 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 Example 2-2 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 Example 2-1 and Example 2-2.

The following example shows that the EMP\_LAST\_NAME index is used for retrieval. Any update performed will possibly be subject to the Halloween problem.

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

The following example 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.

```
SQL> set flags 'strategy';
SQL> declare emp2 cursor for
cont> select * from employees where last_name >= 'M'
cont> order by last_name for update of last_name;
SQL> open emp2;
Temporary relation      Conjunct      Get
Retrieval by index of relation EMPLOYEEES
  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 cannot perform the normal protection against the Halloween problem.



## 14.2 SQL Known Problems and Restrictions

This section describes known problems and restrictions for the SQL interface.

### 14.2.1 SET FLAGS CRONO\_FLAG Removed

The SET FLAGS statement and RDMS\$SET\_FLAGS logical name no longer accept the obsolete keyword CRONO\_FLAG. This keyword has been removed. Please update all scripts and applications to use the keyword CHRONO\_FLAG.

### 14.2.2 Interchange File (RBR) Created by Oracle Rdb Release 7.2 Not Compatible With Previous Releases

To support the large number of new database attributes and objects, the protocol used by SQL EXPORT and SQL IMPORT has been enhanced to support more protocol types. Therefore, this format of the Oracle Rdb release 7.2 interchange files can no longer be read by older versions of Oracle Rdb.

Oracle Rdb continues to provide upward compatibility for interchange files generated by older versions.

Oracle Rdb has never supported backward compatibility, however, it was sometimes possible to use an interchange file with an older version of IMPORT. However, this protocol change will no longer permit this usage.

### 14.2.3 Single Statement LOCK TABLE is Not Supported for SQL Module Language and SQL Precompiler

The new LOCK TABLE statement is not currently supported as a single statement within the module language or embedded SQL language compiler.

Instead you must enclose the statement in a compound statement. That is, use BEGIN... END around the statement as shown in the following example. This format provides all the syntax and flexibility of LOCK TABLE.

This restriction does not apply to interactive or dynamic SQL.

The following extract from the module language listing file shows the reported error if you use LOCK TABLE as a single statement procedure. The other procedure in the same module is acceptable because it uses a compound statement that contains the LOCK TABLE statement.

```
1 MODULE sample_test
2 LANGUAGE C
3 PARAMETER COLONS
4
5 DECLARE ALIAS FILENAME 'mf_personnel'
6
7 PROCEDURE a (SQLCODE);
8 LOCK TABLE employees FOR EXCLUSIVE WRITE MODE;
%SQL-F-WISH_LIST, (1) Feature not yet implemented - LOCK TABLE requires compound
statement
```

```

9
10 PROCEDURE b (SQLCODE);
11 BEGIN
12 LOCK TABLE employees FOR EXCLUSIVE WRITE MODE;
13 END;

```

To workaroud this problem of using LOCK TABLE for SQL module language or embedded SQL application, use a compound statement in an EXEC SQL statement.

## 14.2.4 Multistatement or Stored Procedures May Cause Hangs

Long-running multistatement or stored procedures can cause other users in the database to hang if the procedures obtain resources needed by those other users. Some resources obtained by the execution of a multistatement or stored procedure are not released until the multistatement or stored procedure finishes. Thus, any-long running multistatement or stored procedure can cause other processes to hang. This problem can be encountered even if the statement contains SQL 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 hangs forever.

Session 1:

```

SQL> attach 'filename MF_PERSONNEL';
SQL> create function LIB$WAIT (in real by reference)
cont> returns integer;
cont> external name LIB$WAIT location 'SYS$SHARE:LIBRTL.EXE'
cont> language general general parameter style variant;
SQL> commit;

```

```

.
.
.

```

\$ SQL

```

SQL> attach 'filename 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 = LIBWAIT (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, you can see that the backup process is waiting for a lock held in the first session:

```
$ RMU/SHOW LOCKS /MODE=BLOCKING MF_PERSONNEL
```

```

.
.
.

```

## Oracle® Rdb for OpenVMS

Resource: nowait signal

ProcessID	Process Name	Lock ID	System ID	Requested	Granted
20204383	RMU BACKUP.....	5600A476	00010001	CW	NL
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 are released.

### 14.2.5 Use of Oracle Rdb from Shareable Images

If code in the image initialization routine of a shareable image makes any calls into Oracle Rdb, through SQL or any other means, access violations or other unexpected behavior may occur if Oracle Rdb images have not had a chance to do their own initialization.

To avoid this problem, applications must take one of the following steps:

- Do not make Oracle Rdb calls from the initialization routines of shareable images.
- Link in such a way that the RDBSHR.EXE image initializes first. You can do this by placing the reference to RDBSHR.EXE and any other Oracle Rdb shareable images last in the linker options file.

This is not a bug; it is a restriction resulting from the way OpenVMS image activation works.

## 14.3 Oracle RMU Known Problems and Restrictions

This section describes known problems and restrictions for the RMU interface.

### 14.3.1 RMU Convert Fails When Maximum Relation ID is Exceeded

If, when relation IDs are assigned to new system tables during an RMU Convert to a V7.2 database, the maximum relation ID of 8192 allowed by Oracle Rdb is exceeded, the fatal error %RMU-F-RELMAXIDBAD is displayed and the database is rolled back to the prior database version. Contact your Oracle support representative if you get this error. Note that when the database is rolled back, the fatal error %RMU-F-CVTROLSUC is displayed to indicate that the rollback was successful but caused by the detection of a fatal error and not requested by the user.

This condition only occurs if there are an extremely large number of tables defined in the database or if a large number of tables were defined but have subsequently been deleted.

The following example shows both the %RMU-F-RELMAXIDBAD error message if the allowed database relation ID maximum of 8192 is exceeded and the %RMU-F-CVTROLSUC error message when the database has been rolled back to V7.0 since it cannot be converted to V7.2:

```
$rmu/convert mf_personnel
%RMU-I-RMUTXT_000, Executing RMU for Oracle Rdb V7.2
Are you satisfied with your backup of
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 and your backup of
  any associated .aij files [N]? Y
%RMU-I-LOGCONVRT, database root converted to current structure level
%RMU-F-RELMAXIDBAD, ROLLING BACK CONVERSION - Relation ID exceeds maximum
  8192 for system table RDB$RELATIONS
%RMU-F-CVTROLSUC, CONVERT rolled-back for
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 to version V7.0
```

The following example shows the normal case when the maximum allowed relation ID is not exceeded:

```
$rmu/convert mf_personnel
%RMU-I-RMUTXT_000, Executing RMU for Oracle Rdb V7.2
Are you satisfied with your backup of
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 and your backup of
  any associated .aij files [N]? Y
%RMU-I-LOGCONVRT, database root converted to current structure level
%RMU-S-CVTDBSUC, database DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1
  successfully converted from version V7.0 to V7.2
%RMU-I-CVTCOMSUC, CONVERT committed for
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 to version V7.2
```

### 14.3.2 RMU Unload /After\_Journal Requires Accurate AIP Logical Area Information

The RMU Unload /After\_Journal command uses the on-disk area inventory pages (AIPs) to determine the appropriate type of each logical area when reconstructing logical dbkeys for records stored in mixed-format storage areas. However, the logical area type information in the AIP is generally unknown for logical areas created prior to Oracle Rdb release 7.0.1. If the RMU Unload /After\_Journal command cannot determine the logical area type for one or more AIP entries, a warning message is displayed for each such area and may ultimately return logical dbkeys with a 0 (zero) area number for records stored in mixed-format storage areas.

In order to update the on-disk logical area type in the AIP, the RMU Repair utility must be used. The INITIALIZE=LAREA\_PARAMETERS=optionfile qualifier option file can be used with the TYPE qualifier. For example, to repair the EMPLOYEES table of the MF\_PERSONNEL database, you would create an options file that contains the following line:

```
EMPLOYEES /TYPE=TABLE
```

For partitioned logical areas, the AREA=name qualifier can be used to identify the specific storage areas that are to be updated. For example, to repair the EMPLOYEES table of the MF\_PERSONNEL database for the EMPID\_OVER storage area only, you would create an options file that contains the following line:

```
EMPLOYEES /AREA=EMPID_OVER /TYPE=TABLE
```

The TYPE qualifier specifies the type of a logical area. The following keywords are allowed:

- TABLE  
Specifies that the logical area is a data table. This would be a table created using the SQL CREATE TABLE syntax.
- B-TREE  
Specifies that the logical area is a B-tree index. This would be an index created using the SQL CREATE INDEX TYPE IS SORTED syntax.
- HASH  
Specifies that the logical area is a hash index. This would be an index created using the SQL CREATE INDEX TYPE IS HASHED syntax.
- SYSTEM  
Specifies that the logical area is a system record that is used to identify hash buckets. Users cannot explicitly create these types of logical areas.

---

Note

*This type should NOT be used for the RDB\$SYSTEM logical areas. This type does NOT identify system relations.*

---

- BLOB  
Specifies that the logical area is a BLOB repository.

There is no explicit error checking of the type specified for a logical area. However, an incorrect type may cause the RMU Unload /After\_Journal command to be unable to correctly return valid, logical dbkeys.

### 14.3.3 Do Not Use HYPERSORT with RMU Optimize After\_Journal Command

The OpenVMS Alpha V7.1 operating system introduced the high-performance Sort/Merge utility (also known as HYPERSORT). This utility takes advantage of the OpenVMS 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.

Because of this, 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.

### **14.3.4 Changes in EXCLUDE and INCLUDE Qualifiers for RMU Backup**

The RMU Backup command no longer accepts both the Include and Exclude qualifiers in the same command. This change removes the confusion over exactly what gets backed up when Include and Exclude are specified on the same line, but does not diminish the capabilities of the RMU Backup command.

To explicitly exclude some storage areas from a backup, use the Exclude qualifier to name the storage areas to be excluded. This causes all storage areas to be backed up except for those named by the Exclude qualifier.

Similarly, the Include qualifier causes only those storage areas named by the qualifier to be backed up. Any storage area not named by the Include qualifier is not backed up. The Noread\_only and Noworm qualifiers continue to cause read-only storage areas and WORM storage areas to be omitted from the backup even if these areas are explicitly listed by the Include qualifier.

Another related change is in the behavior of EXCLUDE=\*. In previous versions, EXCLUDE=\* caused all storage areas to be backed up. Beginning with V7.1, EXCLUDE=\* causes only a root backup to be done. A backup created by using EXCLUDE=\* can be used only by the RMU Restore Only\_Root command.

### **14.3.5 RMU Backup Operations Should Use Only One Type of Tape Drive**

When using more than one tape drive for an RMU Backup command, all of the tape drives must be of the same type (for example, all the tape drives must be TA90s or TZ87s or TK50s). Using different tape drive types (for example, 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 of 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 valid. This may be the case, for example, when using a TA90 and a TA90E.

Oracle Corporation recommends that, on a regular basis, you test your backup and recovery procedures and environment using a test system. You should restore the database and then recover using AIJs to simulate failure recovery of the production system.

Consult the Oracle Rdb7 Guide to Database Maintenance, the Oracle Rdb7 Guide to Database Design and Definition, and the Oracle RMU Reference Manual for additional information about Oracle Rdb backup and restore operations.

## 14.3.6 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 Rdb7 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 Oracle Rdb. 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 Oracle Rdb. 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 roll back 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, then the introduction of these errors is considered to be part of the normal operation of Oracle Rdb. 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.



## 14.4 Known Problems and Restrictions in All Interfaces for Release 7.0 and Earlier

The following problems and restrictions from release 7.0 and earlier still exist.

### 14.4.1 Converting Single-File Databases

Because of a substantial increase in the database root file information for V7.0, you should ensure that you have adequate disk space before you use the RMU Convert command with single-file databases and V7.0 or higher.

The size of the database root file of any given database increases a maximum of about 600 disk blocks. The actual increase depends mostly on the maximum number of users specified for the database.

### 14.4.2 Row Caches and Exclusive Access

If a table has a row-level cache defined for it, the Row Cache Server (RCS) may acquire a shared lock on the table and prevent any other user from acquiring a Protective or Exclusive lock on that table.

### 14.4.3 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:

- ◆ Reserve the table for SHARED WRITE
- ◆ Close the database and disable row cache for the duration of the exclusive transaction
- ◆ 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.

### 14.4.4 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)
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.

## 14.4.5 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 less than the smallest existing page size and the database has been manually opened or users are active, the add operation fails with the following errors:

```
%RDMS-F-NOEUACCESS, unable to acquire exclusive access to database
```

or

```

%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, any recovery scenario fails. 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.

## 14.4.6 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 is not 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.

## 14.4.7 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 catch up to the application and are not able to process database pages that are logically ahead of the application in the RDB\$CHANGES system relation. 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 relation 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 V4.0 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.

## 14.5 SQL Known Problems and Restrictions for Oracle Rdb Release 7.0 and Earlier

The following problems and restrictions from Oracle Rdb Release 7.0 and earlier still exist.

### 14.5.1 ARITH\_EXCEPT or Incorrect Results Using LIKE IGNORE CASE

When you use LIKE...IGNORE CASE, programs linked under Oracle Rdb V4.2 and V5.1, 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 (V6.0 or higher.)

### 14.5.2 Different Methods of Limiting Returned Rows from Queries

You can establish the query governor for rows returned from a query by using either the SQL SET QUERY LIMIT statement or a logical name. This note describes the differences between the two mechanisms.

If you define the RDMS\$BIND\_QG\_REC\_LIMIT logical name to a small value, the query often fails with no rows returned regardless of the value assigned to the logical. 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 relations 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 relation) is sufficient to read each column definition.

To see an indication of the queries executed against the system relations, define the RDMS\$DEBUG\_FLAGS logical name 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 SQLOPTIONS=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. 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 relations as part of query processing.

### 14.5.3 Suggestions for Optimal Use of 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 relation 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 the NUMBER OF BUFFERS IS clause in SQL to change the number of buffers).
- ◆ To distribute the disk I/O load, store the storage areas for the indexes on separate disk drives. Note that using the same storage area for multiple indexes results 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 Rdb7 Guide to Database 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 ten parallel process index creations (Index1, Index2, ... Index10) and one process with ten 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 ten 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 ten of the indexes serially.

*Table 14–2 Elapsed Time for Index Creations*

<b>Index Create Job</b>	<b>Elapsed Time</b>
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
All10	00:03:26.66

## 14.5.4 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
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.

## 14.5.5 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 stale 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 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 helps to stabilize the cursor to some degree.

## 14.5.6 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.



## Oracle® Rdb for OpenVMS

- ◆ 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.

[|\\_Contents](#)