

Table of Contents

Oracle® Rdb for OpenVMS	1
Release Notes	2
November 2001	3
Contents	4
Preface	5
Purpose of This Manual	6
Intended Audience	7
Document Structure	8
Chapter 1 Installing Oracle Rdb Release 7.1.0.1	9
1.1 Requirements	10
1.2 Invoking VMSINSTAL	11
1.3 Stopping the Installation	12
1.4 After Installing Oracle Rdb	13
1.5 Alpha EV68 Processor Support Added	14
1.6 Maximum OpenVMS Version Check Added	15
Chapter 2 Software Errors Fixed in Oracle Rdb Release 7.1.0.1	16
2.1 Software Errors Fixed That Apply to All Interfaces	17
2.1.1 Excessive Disk I/O for DROP TABLE and TRUNCATE TABLE	17
2.1.2 LIST Storage Map Not Updated Upon ALTER or DROP TABLE	17
2.1.3 ARBs Exhausted	17
2.1.4 CLEAN BUFFER COUNT Parameter Not Obeyed	18
2.1.5 DETECTED ASYNCHRONOUS PREFETCH THRESHOLD Not Obeyed	18
2.1.6 Page Locks Not Demoted at End of Transaction When FAST COMMIT Enabled	18
2.1.7 Bitmapped Scan Causes Bugcheck on Transaction Termination	18
2.1.8 Problems With Column Outlines	19
2.1.9 Count Scan Optimization Incorrectly Returning Count of 0	20
2.1.10 Disabling AIJ When Row Cache Recovery Required	20
2.1.11 Bitmapped Scan Problem With Large Indexes	21
2.1.12 Query With Range List OR Predicates Returns Wrong Results	22
2.1.13 Database Corruption Using Cluster With Galaxy and Non-Galaxy Nodes	23
2.1.14 Performance Problems when RDM\$BIND SNAP QUIET POINT Defined to 0	24
2.1.15 Workload Ignored When Loaded with RMU/INSERT OPTIMIZER STATISTICS	24
2.1.16 Descending Sort Not Producing Correct Ordering for BIGINT and DATE Columns	25
2.1.17 Bitmapped Scan Incorrectly Chosen by Optimizer	25
2.1.18 Cannot Connect With Remote Access When Using a Logical	26
2.1.19 Query Joining Derived Tables of Union Legs With Empty Tables Returns Wrong Results	27

Table of Contents

2.1.20 Left Outer Join Query With OR Predicate Returns Wrong Results	29
2.1.21 Query Using Match Strategy With DISTINCT Function Returns Wrong Results	30
2.1.22 GROUP BY Query With SUM Aggregate Returns Wrong Results	32
2.1.23 ROLLBACK Hangs Under DECdtm When Called From an ACMS CANCEL Procedure	34
2.1.24 COMPUTED BY Columns Now Automatically Reserve Referenced Tables	34
2.2 SQL Errors Fixed	36
2.2.1 Command Line Recall Expanded to 255 Lines	36
2.2.2 New Minimum Value for the INTERVAL Leading Precision	36
2.2.3 Incorrect Processing of CASE Expression	36
2.2.4 ALTER TABLE Not Dropping NOT NULL Constraints When NULL Clause Used	37
2.2.5 Some Constraint Definitions Not Supported for AUTOMATIC Columns	38
2.2.6 %RDB-E-NO_DIST_BATCH_U Error When Executing SET TRANSACTION	39
2.2.7 Select With Identical "not in" Clauses	39
2.2.8 Keyword Matching Now Reported by Interactive SQL	39
2.2.9 CREATE MODULE Bugchecks When a Subselect is Used as a Parameter DEFAULT	40
2.2.10 Obsolete Metadata Errors When Using Rdb SQL V7.1 to Access Oracle Rdb V7.0 Databases	40
2.2.11 SQL\$PRE and SQL\$MOD Performance Improvements	41
2.2.12 Incompatible Character Sets Not Detected by SQL Interface	41
2.2.13 SQLMOD Fails to Set Default Character Set Correctly	41
2.3 Oracle RMU Errors Fixed	44
2.3.1 RMU Extract Not Formatting View Column Expressions Correctly	44
2.3.2 RMU/UNLOAD/AFTER JOURNAL Fragmented Records Clarification	44
2.3.3 RMU/DUMP/BACKUP Did Not Check the VMS BYPASS Privilege	45
2.3.4 RMU/BACKUP Invalid Volume 1 Tape Label When Used With COMPAQ SLS	45
2.3.5 RMU/ANALYZE/CARDINALITY Fails on Databases With Local Temporary Tables	46
2.3.6 File Name Not Displayed by RMU /RESTORE for Extend Failure	47
2.3.7 RMU/SHOW STATISTICS Allowed Suspend of Disabled ABS	47
2.3.8 RMU/COPY/BLOCKS PER PAGE Can Corrupt Copied Database Uniform Areas	47
2.3.9 DROPPed Storage Area and RMU /VERIFY in Cluster	48
2.3.10 RMU /VERIFY Checks All Storage Area Files First	48
2.3.11 RMU/SHOW STATISTICS Multi-Page Report File	49
2.3.12 Area Locks Demoted Statistic Not Always Correctly Incremented	49
2.3.13 RMU /BACKUP /ONLINE /NOQUIET POINT Fails	49
2.4 LogMiner Errors Fixed	50
2.4.1 LogMiner Compresses Pre-Delete Record Content	50
2.5 Optimizer Problems Fixed in Oracle Rdb Release 7.1.0	51
2.5.1 Query Having OR Compound Predicates With Subquery Returns Wrong Results	51
2.5.2 Query Using OR/AND Predicates With EXISTS Clause Returns Wrong Results	51
2.5.3 Query Using German Collating Sequence Returns Wrong Results	53
2.5.4 Left Outer Join Query Returns Wrong Results When ON Clause Evaluates to False	53
2.5.5 Query With Two IN Clauses on Two Subqueries Returns Wrong Results	54
2.5.6 Query Having Same SUBSTRINGs Within CASE Expression Returns Wrong Results	55
2.5.7 Aggregate Query With Nested MIN Function Returns Wrong Results	57
2.5.8 Query with UNION Subselect Returns Wrong Results	58
2.5.9 Query with CONCATENATE in BETWEEN Clause Returns Wrong Results	59
2.5.10 ORDER BY Query With GROUP BY on Two Joined Derived Tables Returns Wrong Results	60

Table of Contents

2.5.11 Left Outer Join Query With CONCATENATE Returns Wrong Results	62
2.5.12 Query With UNION in German Collating Sequence Returns Wrong Results	63
2.5.13 Query With OR Predicate on Aggregate Column Returns Wrong Results	64
2.5.14 Query With Equality Predicate Included in IN Clause Returns Wrong Results	66
2.5.15 Match Strategy on Columns of Different Size, Using Collating Sequence, Returns Wrong Results	67
2.5.16 Left Outer Join Query With CAST Function on USING Column Bugchecks	68
2.5.17 Query Using Constant Values in OR Predicates Returns Wrong Results	69
Chapter 3 Enhancements	71
3.1 Enhancements Provided in Oracle Rdb Release 7.1.0.1	72
3.1.1 SQL Now Supports a Native ABS Function	72
3.1.2 New DUMP Output Format for LogMiner	73
3.1.3 Data and SPAM Prefetch Screens Added to RMU/SHOW STATISTICS	74
3.1.4 RMU/SHOW STATISTICS Stall Log Lock Information Optional	75
3.1.5 New Option for the GET DIAGNOSTICS Statement	75
3.1.6 Alternate Outline Ids	76
3.1.7 Field Widths Wider on Row Cache Overview Display	79
3.1.8 FOR Counted Loop Enhancements	79
3.1.9 Enhancements to SET DISPLAY Statement for Interactive SQL	81
3.1.10 New BITSTRING Built In Function	83
3.1.11 New SET PAGE LENGTH Command for Interactive SQL	84
3.1.12 New ALTER CONSTRAINT Statement	84
3.1.13 DECLARE Variable Now Supports CHECK Constraint	87
3.1.14 RMU/SHOW STATISTICS Active User Stall Messages Sorted by Process ID	88
3.1.15 RMU/REPAIR /INITIALIZE ONLY LAREA TYPE Keyword	88
3.1.16 RMU/SHOW STATISTICS Cluster Data Collection Performance Enhancement	89
3.1.17 RMU Extract has Enhanced Extract of Conditional Expressions	89
Chapter 4 Documentation Corrections, Additions and Changes	90
4.1 Documentation Corrections	91
4.1.1 DROP INDEX Now an Online Table Operation	91
4.2 Address and Phone Number Correction for Documentation	92
4.3 Online Document Format and Ordering Information	93
4.3.1 Documentation in Adobe Acrobat Format	93
4.3.2 Documentation in HTML format	93
4.4 Documentation for This Release	94
4.5 Updated Documentation for Oracle Rdb-related Products	97
4.6 New and Changed Features in Oracle Rdb Release 7.1	98
4.6.1 PERSONA is Supported in Oracle SQL/Services	98
4.6.2 NEXTVAL and CURRVAL Pseudocolumns Can Be Delimited Identifiers	98
4.6.3 Only=select list Qualifier for the RMU Dump After Journal Command	98

Table of Contents

<u>4.7 Oracle Rdb7 and Oracle CODASYL DBMS Guide to Hot Standby Databases</u>	100
<u>4.7.1 Restrictions Lifted on After-Image Journal Files</u>	100
<u>4.7.2 Changes to RMU Replicate After Journal ... Buffer Command</u>	100
<u>4.7.3 Unnecessary Command in the Hot Standby Documentation</u>	101
<u>4.7.4 Change in the Way RDMAIJ Server is Set Up in UCX</u>	101
<u>4.7.5 CREATE INDEX Operation Supported for Hot Standby</u>	102
<u>4.8 Oracle Rdb7 for OpenVMS Installation and Configuration Guide</u>	103
<u>4.8.1 Suggestion to Increase GH_RSRVPGCNT Removed</u>	103
<u>4.8.2 Prerequisite Software</u>	103
<u>4.8.3 Defining the RDBSERVER Logical Name</u>	103
<u>4.9 Guide to Database Design and Definition</u>	105
<u>4.9.1 Lock Timeout Interval Logical Incorrect</u>	105
<u>4.9.2 Example 4-13 and Example 4-14 Are Incorrect</u>	105
<u>4.10 Oracle Rdb7 SQL Reference Manual</u>	106
<u>4.10.1 Clarification of the DDLDONOTMIX Error Message</u>	106
<u>4.10.2 Node Specification Allowed on Root FILENAME Clauses</u>	106
<u>4.10.3 Incorrect Syntax Shown for Routine-Clause of the CREATE MODULE Statement</u>	107
<u>4.10.4 Omitted SET Statements</u>	107
<u>4.10.4.1 QUIET COMMIT</u>	107
<u>4.10.4.2 COMPOUND TRANSACTIONS</u>	108
<u>4.10.5 Size Limit for Indexes with Keys Using Collating Sequences</u>	109
<u>4.10.6 Clarification of SET FLAGS Option DATABASE_PARAMETERS</u>	109
<u>4.10.7 Incorrect Syntax for CREATE STORAGE MAP Statement</u>	110
<u>4.10.8 Use of SQL_SOLCA Include File Intended for Host Language File</u>	111
<u>4.10.9 Missing Information on Temporary Tables</u>	112
<u>4.11 Oracle RMU Reference Manual, Release 7.0</u>	113
<u>4.11.1 RMU Unload After Journal Null Bit Vector Clarification</u>	113
<u>4.11.2 New Transaction_Mode Qualifier for Oracle RMU Commands</u>	115
<u>4.11.3 RMU Server After Journal Stop Command</u>	116
<u>4.11.4 Incomplete Description of Protection Qualifier for RMU Backup After Journal Command</u>	117
<u>4.11.5 RMU Extract Command Options Qualifier</u>	117
<u>4.11.6 RDM\$SNAP QUIET_POINT Logical is Incorrect</u>	117
<u>4.11.7 Using Delta Time with RMU Show Statistics Command</u>	117
<u>4.12 Oracle Rdb7 Guide to Database Performance and Tuning</u>	118
<u>4.12.1 Dynamic OR Optimization Formats</u>	118
<u>4.12.2 Oracle Rdb Logical Names</u>	118
<u>4.12.3 Waiting for Client Lock Message</u>	118
<u>4.12.4 RDM\$STTB_HASH_SIZE Logical Name</u>	119
<u>4.12.5 Error in Updating and Retrieving a Row by Dbkey Example 3-22</u>	120
<u>4.12.6 Error in Calculation of Sorted Index in Example 3-46</u>	121
<u>4.12.7 Documentation Error in Section C.7</u>	121
<u>4.12.8 Missing Tables Descriptions for the RDBEXPERT Collection Class</u>	122
<u>4.12.9 Missing Columns Descriptions for Tables in the Formatted Database</u>	122
<u>4.12.10 A Way to Find the Transaction Type of a Particular Transaction Within the Trace Database</u>	130
<u>4.12.11 Using Oracle TRACE Collected Data</u>	130

Table of Contents

4.12.12 AIP Length Problems in Indexes that Allow Duplicates	131
4.12.13 RDM\$BIND MAX DBR COUNT Documentation Clarification	133
4.13 Oracle Rdb7 Guide to SQL Programming	134
4.13.1 Location of Host Source File Generated by the SQL Precompiler	134
4.13.2 Remote User Authentication	135
4.13.3 Additional Information About Detached Processes	135
4.14 Guide to Using Oracle SQL/Services Client APIs	137
4.15 Updates to System Relations	138
4.15.1 Clarification on Updates to the RDB\$LAST ALTERED Column for the RDB\$DATABASE System Relation	138
4.15.2 Missing Descriptions of RDB\$FLAGS	138
4.16 Error Messages	141
4.16.1 Clarification of the DDLDONOTMIX Error Message	141
Chapter 5 Known Problems and Restrictions	142
5.1 Known Problems and Restrictions in All Interfaces	143
5.1.1 RDB-E-ARITH EXCEPT Error From the Rdb Optimizer	143
5.1.2 RMU Fails to Perform OPTIMIZER STATISTICS Actions on Some Databases	143
5.1.3 Possible RMU Bugcheck or Failure to Notify Triggering of User Defined Events	144
5.1.4 Optimization of Check Constraints	144
5.1.5 Using Databases from Releases Earlier Than V5.1	146
5.1.6 PAGE TRANSFER VIA MEMORY Disabled	146
5.1.7 Carryover Locks and NOWAIT Transaction Clarification	147
5.1.8 Unexpected Results Occur During Read-Only Transactions on a Hot Standby Database	147
5.1.9 IMPORT Unable to Import Some View Definitions	147
5.1.10 Both Application and Oracle Rdb Using SYSSHIBER	148
5.1.11 Bugcheck Dump Files with Exceptions at COSI CHF SIGNAL	149
5.1.12 Read-only Transactions Fetch AIP Pages Too Often	150
5.1.13 Row Cache Not Allowed While Hot Standby Replication is Active	150
5.1.14 Excessive Process Page Faults and other Performance Considerations During Oracle Rdb Sorts	150
5.1.15 Control of Sort Work Memory Allocation	151
5.1.16 The Halloween Problem	152
5.2 SQL Known Problems and Restrictions	154
5.2.1 Interchange File (RBR) Created by Oracle Rdb Release 7.1 Not Compatible With Previous Releases	154
5.2.2 Unexpected NO META UPDATE Error Generated by DROP MODULE ... CASCADE When Attached by PATHNAME	154
5.2.3 Problem Exporting and Importing Sequences with ANSI-Style Databases	154
5.2.4 System Relation Change for International Database Users	154
5.2.5 Single Statement CALL Does Not Support Truncated Parameter List or DEFAULT Keyword	155
5.2.6 Single Statement LOCK TABLE is Not Supported for SQL Module Language and SQL Precompiler	155
5.2.7 Restriction for CREATE STORAGE MAP Statement on Table with Data	156
5.2.8 Multistatement or Stored Procedures May Cause Hangs	156

Table of Contents

5.2.9 Use of Oracle Rdb from Shareable Images.....	157
5.3 Oracle RMU Known Problems and Restrictions.....	159
5.3.1 RMU/CONVERT Fails to Correctly Define the RDB\$WORKLOAD Table.....	159
5.3.2 RMU Convert Fails When Maximum Relation ID is Exceeded.....	159
5.3.3 RMU Unload /After Journal Requires Accurate AIP Logical Area Information.....	160
5.3.4 Do Not Use HYPERSORT with RMU Optimize After Journal Command.....	161
5.3.5 Changes in EXCLUDE and INCLUDE Qualifiers for RMU Backup.....	161
5.3.6 Default for RMU CRC Qualifier Changing in Future Release.....	162
5.3.7 RMU Backup Operations Should Use Only One Type of Tape Drive.....	162
5.3.8 RMU/VERIFY Reports PGSPAMENT or PGSPMCLST Errors.....	163
5.4 Known Problems and Restrictions in All Interfaces for Release 7.0 and Earlier.....	164
5.4.1 Converting Single-File Databases.....	164
5.4.2 Row Caches and Exclusive Access.....	164
5.4.3 Exclusive Access Transactions May Deadlock with RCS Process.....	164
5.4.4 Strict Partitioning May Scan Extra Partitions.....	164
5.4.5 Restriction When Adding Storage Areas with Users Attached to Database.....	165
5.4.6 Support for Single-File Databases to Be Dropped in a Future Release.....	165
5.4.7 Multiblock Page Writes May Require Restore Operation.....	165
5.4.8 Network Link Failure Does Not Allow DISCONNECT to Clean Up Transactions.....	166
5.4.9 Replication Option Copy Processes Do Not Process Database Pages Ahead of an Application.....	166
5.5 SQL Known Problems and Restrictions for Oracle Rdb Release 7.0 and Earlier.....	168
5.5.1 SQL Does Not Display Storage Map Definition After Cascading Delete of Storage Area.....	168
5.5.2 ARITH EXCEPT or Incorrect Results Using LIKE IGNORE CASE.....	168
5.5.3 Different Methods of Limiting Returned Rows from Queries.....	168
5.5.4 Suggestions for Optimal Use of SHARED DATA DEFINITION Clause for Parallel Index Creation.....	170
5.5.5 Side Effect When Calling Stored Routines.....	171
5.5.6 Considerations When Using Holdable Cursors.....	172

Oracle® Rdb for OpenVMS

Release Notes

Release 7.1.0.1

November 2001

Oracle Rdb Release Notes, Release 7.1.0.1 for OpenVMS

Copyright © 1984, 2001 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 is prohibited.

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

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

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

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

Oracle is a registered trademark, and Oracle Rdb, Rdb7, Oracle SQL/Services, Oracle7, Oracle Expert, and Oracle Rally are trademarks 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 7.1.0.1. 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 7.1.0.1.

Document Structure

This manual consists of five chapters:

Chapter 1	Describes how to install Oracle Rdb Release 7.1.0.1.
Chapter 2	Describes software errors corrected in Oracle Rdb Release 7.1.0.1.
Chapter 3	Describes enhancements introduced in Oracle Rdb Release 7.1.0.1.
Chapter 4	Provides information not currently available in the Oracle Rdb documentation set.
Chapter 5	Describes problems, restrictions, and workarounds known to exist in Oracle Rdb Release 7.1.0.1.

Chapter 1

Installing Oracle Rdb Release 7.1.0.1

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

NOTE

All Oracle Rdb Release 7.1 kits are full kits. There is no need to install any prior release of Oracle Rdb when installing new Rdb Release 7.1 kits.

1.1 Requirements

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

- Oracle Rdb must be shutdown before you install this update kit. That is, the command file `SYS$STARTUP:RMONSTOP71.COM` should be executed before proceeding with this installation. If you have an OpenVMS cluster, you must shutdown the Rdb 7.1 monitor on all nodes in the cluster before proceeding.
- The installation requires approximately 200,000 blocks for OpenVMS Alpha systems.

1.2 Invoking VMSINSTAL

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

```
@SYS$UPDATE:VMSINSTAL RDBAMVE1071 device-name OPTIONS N
```

device-name

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

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

```
DKA400:[RDBAMVE1071.KIT]
```

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

```
MTA0:
```

OPTIONS N

This parameter prints the release notes.

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

```
$ @SYS$UPDATE:VMSINSTAL RDBAMVE1071 MTA0: OPTIONS N
```

1.3 Stopping the Installation

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

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

1.4 After Installing Oracle Rdb

This update provides a new Oracle Rdb Oracle TRACE facility definition. Any Oracle TRACE selections that reference Oracle Rdb will need to be redefined to reflect the new facility version number for the updated Oracle Rdb facility definition, "RDBVMSV7.1-01".

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=RDBVMSV7.1-01 -  
_ $ /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.

1.5 Alpha EV68 Processor Support Added

For this release of Rdb, Oracle Rdb Release 7.1.0.1, the Alpha EV68 processor is the newest processor supported.

1.6 Maximum OpenVMS Version Check Added

As of Oracle Rdb7 Release 7.0.1.5, a maximum OpenVMS version check has been added to the product. Oracle Rdb has always had a minimum OpenVMS version requirement. With 7.0.1.5 and for all future Oracle Rdb releases, we have expanded this concept to include a maximum VMS version check and a maximum supported processor hardware check. The reason for this check is to improve product quality.

OpenVMS Version 7.3 is the maximum supported version of OpenVMS.

As of Oracle Rdb Release 7.1, the Alpha EV68 processor is supported.

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

Chapter 2

Software Errors Fixed in Oracle Rdb Release 7.1.0.1

This chapter describes software errors that are fixed by Oracle Rdb Release 7.1.0.1.

2.1 Software Errors Fixed That Apply to All Interfaces

2.1.1 Excessive Disk I/O for DROP TABLE and TRUNCATE TABLE

Bug 989292

In prior releases of Oracle Rdb, the DROP TABLE and TRUNCATE TABLE statements performed excessive disk I/O when the table contained LIST OF BYTE VARYING columns. When this data type is present, these operations must read the table to locate the LIST data. In prior releases, a DELETE operation was also performed on the table. While this achieved the delete of the LIST data, it also caused constraints, and possibly triggers, to be executed along with updating indices as each row was deleted.

This problem was corrected in Oracle Rdb7 Release 7.0.4 and was inadvertently left out of the Release Notes. The DROP TABLE and TRUNCATE TABLE statements no longer cause constraints and triggers to be executed for the table and indices are no longer updated when processing the LIST OF BYTE VARYING columns. The result is that I/O required for DROP TABLE and TRUNCATE TABLE is significantly reduced, especially for tables stored in UNIFORM format storage areas.

2.1.2 LIST Storage Map Not Updated Upon ALTER or DROP TABLE

Bug 908343

Database administrators can use CREATE STORAGE MAP to establish special storage area mapping for LIST OF BYTE VARYING columns. The LIST storage map can be used to place all or some of the columns of the table in specified storage areas. However, it has been reported that this storage map is not updated when a DROP TABLE or an ALTER TABLE ... DROP COLUMN is executed.

The LIST data is deleted from the database, however, the name of the table or column is left in the storage map. This leads to confusion later when RMU/EXTRACT is used to process the storage map. Further, if columns from the table were the only data stored in that partition, Rdb would not delete the logical area when the table was dropped.

These problems have been corrected in Oracle Rdb Release 7.1. Oracle Rdb now implicitly updates the LIST storage map when you drop a referenced table or column.

2.1.3 ARBs Exhausted

It was possible for a database to run out of AIJ Request Blocks (ARBs) if many processes were abnormally terminated. If a process had an ARB allocated at the time it was terminated, the Database Recovery Process (DBR) would fail to free the ARB allocated to the process. This problem was introduced in Oracle Rdb Release 7.0.1.2.

Symptoms of this problem include:

- Processes looping. RMU/SHOW STATISTICS would show processes stalling waiting for the AIJ lock or writing the same AIJ block over and over.
- More AIJ activity due to processes flushing the ARBs more often in attempts to make ARBs available.

- The "AIJ Journal Information" screen displayed by RMU/SHOW STATISTICS would show the available ARB count ("ARB.Avail:") to be few or none.

To avoid the problem, avoid terminating processes via the DCL STOP /IDENTIFICATION command. When the problem occurs, the database must be closed and re-opened on each node where the problem is being seen to reset the free ARB lists.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.4 CLEAN BUFFER COUNT Parameter Not Obeyed

When the Asynchronous Batch Write feature is being used, Oracle Rdb is supposed to inspect the tail of the least recently used (LRU) buffer queue to determine if there are any modified buffers at the end of the queue. The CLEAN BUFFER COUNT parameter specifies how many buffers are to be inspected. If any are found then those buffers are supposed to be written to disk. However, when unmarking buffers, Oracle Rdb would unmark buffers at the end of the modified queue instead of the LRU queue. That could cause buffers that were just modified to be immediately written, even if they were the most recently accessed buffers. This could cause the buffer to have to be modified again and thus written again.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. Instead of writing the buffers at the tail of the modified queue, Oracle Rdb now writes the modified buffers at the end of the LRU queue.

2.1.5 DETECTED ASYNCHRONOUS PREFETCH THRESHOLD Not Obeyed

The detected asynchronous prefetch (DAPF) feature is supposed to initiate asynchronous prefetch (APF) requests if it detects consecutive pages being fetched from a storage area. The THRESHOLD parameter declares how many consecutive buffers read in a sequence will trigger an APF request. However, Oracle Rdb would not actually initiate APF requests until the THRESHOLD count plus half the DEPTH number of buffers were sequentially read.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. DAPF will now be triggered when THRESHOLD number of consecutive buffers are read in a sequence.

2.1.6 Page Locks Not Demoted at End of Transaction When FAST COMMIT Enabled

When using the FAST COMMIT feature, at the end of a transaction, page locks were not being demoted. Page locks are always demoted at the end of a transaction when the FAST COMMIT feature is not enabled. In some applications, demoting page locks at the end of a transaction can significantly reduce the incidence of deadlocks involving page locks.

This situation has been improved in Oracle Rdb Release 7.1.0.1. When the FAST COMMIT feature is enabled, at the end of a transaction, any buffer that does not contain a modified page will have its page locks demoted.

2.1.7 Bitmapped Scan Causes Bugcheck on Transaction Termination

Bug 1978724

A problem with the way bitmapped scan uses indexes in the dynamic optimizer to carry out the scan caused bugchecks on transaction or session termination.

The call stacks of these bugcheck dumps may include the following:

```
KOD$ROLLBACK + 00000154
%COSI-F-BUGCHECK, internal consistency failure
```

or

```
KOD$PREPARE + 00000288
```

This problem may occur when the dynamic optimizer determines that a query may be satisfied by three or more indexes, the first priority index chosen being a non-ranked index (that is, either a normal sorted or a hashed index). At least two of the remaining indexes have to be sorted ranked indexes for the optimizer to choose to implement the 'bitmapped scan' optimization.

An example of the portion of the strategy dump from a query that will exhibit this behavior follows:

```
Leaf#01 FFirst CLIENT_DATA Card=5001      Bitmapped scan
  BgrNdx1 HASHED_1 [(1:1)2] Fan=1
  BgrNdx2 RANKED_3 [1:1] Fan=82
  BgrNdx3 RANKED_2 [1:1] Fan=82
  BgrNdx4 NON_RANKED_1 [1:1] Fan=82
```

A possible workaround for this problem is to disable bitmapped scans by either:

```
set flags 'nobitmapped_scan;
```

or

```
$ define RDMS$DISABLE_BITMAPPED_SCAN "1"
```

Disabling bitmapped scan optimization does not stop bitmapped indexes from being used for data retrieval.

Another possible workaround is to either change the first index chosen by the dynamic optimizer to a ranked index or to disable that index entirely.

This problem does not cause any data corruption in your database.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.8 Problems With Column Outlines

Two problems have been found with the creation of outlines on COMPUTED BY columns.

1. Bugcheck dumps may be seen when trying to create outlines on COMPUTED BY columns that use aggregate functions such as MAX or MIN.

For example, attempting to create an outline on the following COMPUTED BY column would generate a bugcheck dump.

```
F1 computed by (select MAX(job_end) from JOB_HISTORY)
```

There is no known workaround for this problem.

2. If two or more COMPUTED BY columns exist on the same table, and at least one of these columns has an outline created on it, it is possible that when the optimizer tries to optimize a query using these outlines, the query optimization will fail and the query will be aborted with the following error message:

```
%RDMS-F-LEVEL_MISMATCH, the table/subquery nesting levels in the query outline
do not match the query
```

This problem may occur when a query references at least two COMPUTED BY columns from the same table and one of these has an outline stored for it.

Possible workarounds for this problem are to drop the offending outline or to disable outlines by using the SET FLAGS 'IGNORE_OUTLINES' statement.

These problems have been corrected in Oracle Rdb Release 7.1.0.1.

2.1.9 Count Scan Optimization Incorrectly Returning Count of 0

Bug 2020109

A problem in the new COUNT SCAN optimization used with ranked indexes may cause incorrect results to be returned by COUNT. Depending on the distribution of keys within the ranked index nodes and the search criteria provided to the COUNT statement, the COUNT statement may incorrectly return a value of 0.

This problem will only occur when the optimizer uses count scan optimization on a sorted ranked index where the search criteria provided in the selection expression for the COUNT statement generates a search key that does not match an existing key within the index. Depending on key distribution, the scan may, infrequently, terminate prematurely resulting in an incorrect value of 0 being returned.

A possible workaround for this problem is to disable count scan optimization by using the SET FLAGS statement or logical name, as in the following example.

```
SQL> SET FLAGS 'NOCOUNT_SCAN' ;

or

$ DEFINE RDMS$SET_FLAGS 'NOCOUNT_SCAN'
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.10 Disabling AIJ When Row Cache Recovery Required

Bug 1831040

When after-image journaling is manually disabled on a closed database that had Row Caching active and requires recovery, it is possible to render the database unusable. For example, consider the following sequence of events:

1. Database is running with Row Caching enabled
2. AIJ files not backed up and eventually fill
3. User processes deleted or system fails
4. User enters RMU /SET AFTER_JOURNAL /DISABLE command

At this point, a warning message is displayed, but the database can not be opened because the DBR process will fail when attempting to access the after image journal files.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. Attempts to disable journaling will now result in a fatal error and journaling will not be disabled when Row Cache recovery is required. The following example demonstrates this condition.

```
$ RMU/SET AFTER/DISABLE MF_PERSONNEL.RDB
%RMU-W-DBRABORTED, database recovery process terminated abnormally
%RMU-F-MUSTRECDB, database must be closed or recovered
%RMU-F-FTL_SET, Fatal error for SET operation at 11-SEP-2001 22:52:22.37
```

2.1.11 Bitmapped Scan Problem With Large Indexes

Bug 2030599

A problem in the new bitmapped scan optimization used with ranked indexes may infrequently cause Rdb to return zero records even when matching records exist.

This problem may be found only when either the data records associated with the keys stored in the ranked indexes span more than 131070 pages or if the data records span over 3 or more areas. In addition, the existence of this problem depends strongly on the distribution of those records and the selection criteria used to match records across the indexes.

Bitmap scan optimization may be chosen by the optimizer when two or more ranked indexes are found that may satisfy all or part of the selection criteria of a query.

Dumping the query strategy using the 'STRATEGY' debug flag will show those queries that have been optimized this way. At the end of the LEAF information of the strategy dump will be the phrase 'Bitmapped scan', as in the following example.

```
Leaf#01 FFirst CUSTOMER_DATA Card=5065237 Bitmapped scan
  BgrNdx1 ADDR_INDEX [1:1] Fan=82      (index scan#2)
  BgrNdx2 NAME_INDEX [1:1] Fan=82      (index scan#3)
  BgrNdx3 POSTCODE_INDEX [1:1] Fan=82  (index scan#4)
```

A possible workaround for this problem is to disable bitmapped scan optimization by using the SET FLAGS statement or logical name.

For example:

```
SQL> SET FLAGS 'NOBITMAPPED_SCAN' ;

or

$ DEFINE RDMS$SET_FLAGS 'NOBITMAPPED_SCAN'
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.12 Query With Range List OR Predicates Returns Wrong Results

Bug 1329838

The following query with range list OR predicates returns wrong results.

```

set flags 'strategy,detail';

select t,m,p,b from a
  where (t='S' and (m='N' or p='Q')) or (t='Z' and (m='N' or b='A'))
 order by t,m,p,b;
Tables:
  0 = A
Sort: 0.T(a), 0.M(a), 0.P(a), 0.B(a)
Conjunct: ((0.T = 'S') AND ((0.M = 'N') OR (0.P = 'Q'))) OR ((0.T = 'Z') AND ((
  0.M = 'N') OR (0.B = 'A')))
OR index retrieval          ! <== Let's call this "Outer"
Conjunct: (0.B = 'A') OR (0.M = 'N') OR (0.M = 'N')
OR index retrieval          ! <== let's call this "Inner"
Get      Retrieval by index of relation 0:A
  Index name  BTY_X [(2:2)]
  Keys: (0.B = 'A') AND (0.T = 'Z')
Conjunct: NOT (0.B = 'A') AND ((0.M = 'N') OR (0.M = 'N')) ! <== incorrect
Get      Retrieval by index of relation 0:A
  Index name  MTZ_X [(2:2)2]
  Keys: r0: (0.M = 'N') AND (0.T = 'S')
      r1: (0.M = 'N') AND (0.T = 'Z')
Conjunct: NOT ((0.B = 'A') OR (0.M = 'N') OR (0.M = 'N')) ! <== incorrect
Get      Retrieval by index of relation 0:A
  Index name  PZY_X [1:1]
  Keys: 0.P = 'Q'
T      M      P      B
S      M      Q      B
S      M      Q      NULL
S      N      P      B
S      N      P      NULL
S      N      Q      B
S      N      Q      NULL
S      N      NULL   B
S      N      NULL   NULL
S      NULL   Q      B
S      NULL   Q      NULL
10 rows selected

```

where the sequential access gives the correct result:

```

select t,m,p,b from a
  where (t='S' and (m='N' or p='Q')) or (t='Z' and (m='N' or b='A'))
 order by t,m,p,b optimize for sequential access;
T      M      P      B
S      M      Q      A      <= missing row
S      M      Q      B
S      M      Q      NULL
S      N      P      A      <= missing row
S      N      P      B
S      N      P      NULL
S      N      Q      A      <= missing row
S      N      Q      B
S      N      Q      NULL
S      N      NULL   A      <= missing row
S      N      NULL   B

```

```

S      N      NULL      NULL
S      NULL    Q          A          <= missing row
S      NULL    Q          B
S      NULL    Q          NULL
15 rows selected

```

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

1. The main select query contains a where clause with range list OR predicates that involves four columns, each testing equality with a constant literal value. In this example, we use the column names B, M, P, and T.
2. The column T is a common segment between index BTY_X and MTZ_X, where BTY_X is an index on columns B, T and Y; MTZ_X is an index on columns M, T, and Y. The column P is defined as a leading segment in PZY_X.
3. The main OR predicate has the left branch which contains an AND between "T='S'" and another secondary OR predicate "(m='N' or p='Q')". The right branch contains an AND between "T='Z'" and another secondary OR predicate "(m='N' or b='A')".
4. The OR predicates are arranged in such a way that the strategy of the optimizer uses the range list retrieval "MTZ_X [(2:2)2]" on keys "r0: (0.M = 'N') AND (0.T = 'S')" and "r1: (0.M = 'N') AND (0.T = 'Z')" in the second leg of the "inner" OR index retrieval under the first leg of the "outer" OR index retrieval.
5. The NOT filter, created at the top of the second leg of the "inner" OR index retrieval, does not contain the equality predicate "0.T = 'Z'" from the first leg.
6. The NOT filter, created at the top of the second leg of the "outer" OR index retrieval, does not contain the predicates "(0.T = 'S')" and "(0.T = 'Z')" from the range list predicates of the first leg.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.13 Database Corruption Using Cluster With Galaxy and Non-Galaxy Nodes

It was possible for page updates to be lost when the following conditions were true:

- The database had GALAXY SUPPORT IS ENABLED.
- The database had GLOBAL BUFFERS ENABLED.
- The database was being accessed concurrently by both OpenVMS Galaxy and non-Galaxy nodes.
- The database was often being closed and reopened on one or more of the Galaxy nodes, but never closed on all of the Galaxy nodes at the same time.

In the above situation, it was possible for updates made by a non-Galaxy node to be lost if the non-Galaxy node closed the database and pages modified by the non-Galaxy node were also present in the global buffer pool being shared by the Galaxy nodes, and those pages in the Galaxy global buffer pool were not being used by any of the Galaxy nodes at the time the database was closed by the non-Galaxy node.

Any of the following actions can be taken to workaround the problem:

- Disable GALAXY SUPPORT.
- Disable GLOBAL BUFFERS.
- Manually open the database on all Galaxy nodes and keep the database open on all Galaxy nodes until all users accessing the database from the Galaxy nodes detach from the database.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.14 Performance Problems when RDM\$BIND_SNAP_QUIET_POINT Defined to 0

Bug 884004

When the logical name RDM\$BIND_SNAP_QUIET_POINT was defined to 0, it would cause Oracle Rdb to write out modified buffers and demote all page buffer locks when a READ ONLY transaction was started. This would defeat the optimizations utilized by the FAST COMMIT feature and would also cause additional locking and page buffer I/O.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. When the RDM\$BIND_SNAP_QUIET_POINT logical is defined to 0 and a process is holding the quiet point lock when starting a READ ONLY transaction, the quiet point lock will be retained. Thus buffers will not be flushed and page locks will not be released when starting a READ ONLY transaction. If a backup process requests the quiet point lock, and the logical RDM\$BIND_SNAP_QUIET_POINT is defined to 0, then any READ ONLY transactions will immediately write out modified buffers and release the quiet point lock.

2.1.15 Workload Ignored When Loaded with RMU/INSERT OPTIMIZER_STATISTICS

In previous versions of Oracle Rdb, if workload statistics were loaded into a database using the *RMU/INSERT OPTIMIZER_STATISTICS* command, the workload would be ignored by the optimizer.

The use of workload statistics can be observed by setting the *ESTIMATES* debug flag as shown in the following example.

```
SQL> set flags 'estimates';
SQL> select * from t1 where f1=1;
Solutions tried 1
Solutions blocks created 1
Created solutions pruned 0
Cost of the chosen solution    3.0000000E+00
Cardinality of chosen solution  1.0000000E+00
~O: Workload statistics used
      F1          F2
      1          1
1 row selected
```

After loading workload statistics with the *RMU/INSERT* command, a query that should use statistics will fail to show the *~O: Workload statistics used* message. This indicates that the statistics are being ignored.

The problem can be identified by examining the data loaded into the *RDB\$WORKLOAD* system table. If the *RDB\$CREATED* and *RDB\$LAST_ALTERED* columns have the same value, as shown in the following example, then workload statistics will be ignored.

```
SQL> select rdb$created,rdb$last_altered from rdb$workload;
RDB$CREATED          RDB$LAST_ALTERED
19-OCT-2001 00:33:53.27  19-OCT-2001 00:33:53.27
1 row selected
```

The problem can be corrected by manually updating the *RDB\$LAST_ALTERED* column, as shown in the following example. New attaches will commence using the workload values.

```
SQL> update rdb$workload set rdb$last_altered=current_timestamp
cont>where rdb$relation_name='...';
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.16 Descending Sort Not Producing Correct Ordering for BIGINT and DATE Columns

Bugs 2064232 and 2058531

Oracle Rdb Release 7.1 introduced a new fast sort facility (known as QSORT) which is used when the number of rows to be sorted are few and the sort keys are simple.

Unfortunately, QSORT did not correctly handle descending sorts for 64 bit values, such as BIGINT, DATE (both VMS and ANSI formats), TIME, TIMESTAMP and INTERVAL.

A workaround for this problem is to disable QSORT and revert to the normal sort interface by defining the following logical name to the value zero (0).

```
$ DEFINE RDMS$BIND_MAX_QSORT_COUNT 0
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.17 Bitmapped Scan Incorrectly Chosen by Optimizer

A problem in the way the Rdb optimizer determines when to use the new bitmapped scan optimization used with ranked indexes may infrequently cause Rdb to return wrong results.

The optimizer may sometimes incorrectly choose to carry out bitmapped scans that are not appropriate given the selection criteria of the query in relation to the columns available to be used within the ranked index columns. See the following example:

```
SQL> att 'file personnel';
SQL> CREATE TABLE bmtest(A INTEGER,B INTEGER,C INTEGER,D INTEGER,E INTEGER);
SQL> INSERT INTO bmtest VALUES(1,1,10,100,1000);
1 row inserted
SQL> INSERT INTO bmtest VALUES(2,1,10,100,1000);
1 row inserted

SQL> SET FLAGS 'STRATEGY';
SQL> SEL * FROM bmtest WHERE B=1;
Conjunct      Get      Retrieval sequentially of relation BMTEST
      A          B          C          D          E
      1          1          10         100         1000
      2          1          10         100         1000
2 rows selected
SQL> SEL * FROM bmtest WHERE B=1 AND D=100;
Conjunct      Get      Retrieval sequentially of relation BMTEST
      A          B          C          D          E
      1          1          10         100         1000
      2          1          10         100         1000
2 rows selected

SQL> SET FLAGS 'NOSTRATEGY';

SQL> CREATE INDEX bmtest_BCA ON bmtest(B,C,A) TYPE IS SORTED RANKED;
SQL> CREATE INDEX bmtest_DEA ON bmtest(D,E,A) TYPE IS SORTED RANKED;

SQL> SET FLAGS 'STRATEGY';
SQL> SEL * FROM bmtest WHERE B=1;
```

```

Leaf#01 FFirst BMTEST Card=0
  BgrNdx1 BMTEST_BCA [1:1] Fan=12
      A          B          C          D          E
      1          1          10         100         1000
      2          1          10         100         1000
2 rows selected
SQL> SEL * FROM bmtest WHERE B=1 AND D=100;
Leaf#01 FFirst BMTEST Card=0      Bitmapped scan
  BgrNdx1 BMTEST_BCA [1:1] Fan=12
  BgrNdx2 BMTEST_DEA [1:1] Fan=12
      A          B          C          D          E
      1          1          10         100         1000
1 row selected

```

The last query shows that bitmapped scan has been used but returns incorrect results. Bitmapped scan should not be invoked unless the query provides equality checks for all the columns in the ranked index.

A possible workaround for this problem is to disable bitmapped scan optimization by using the SET FLAGS statement or logical name. See the following example:

```

SQL> SET FLAGS 'NOBITMAPPED_SCAN';

or

$ DEFINE RDMS$SET_FLAGS 'NOBITMAPPED_SCAN'

```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.18 Cannot Connect With Remote Access When Using a Logical

Bug 451582

If a logical is used to specify the path in a remote attach, an Rdb 7.1 client fails to connect to the remote database. Depending on the way the database name is specified, either a `-RDB-E-BAD_DB_FORMAT` or `-RDB-F-NONODE` is returned. This problem is similar to Bug 451582. The following example shows the problem behavior and the workarounds.

```

ALPHA4> define ll malibu::disk$users:[remote_account]
ALPHA4> sql
SQL> attach 'filename ll:v70db';
%SQL-F-ERRATTDEC, Error attaching to database ll:my_db
-RDB-E-BAD_DB_FORMAT, ll:v70db does not reference a database known to Rdb
-RMS-E-FNF, file not found
SQL> attach 'filename ll:v70db.rdb';
%SQL-F-ERRATTDEC, Error attaching to database ll:my_db.rdb
-RDB-F-SYS_REQUEST, error from system services request
-RDMS-F-BADDBNAME, can't find database root ALPHA4::DISK$USERS:[REMOTE_ACCOUNT]`
-RDMS-F-NONODE, no node name is allowed in the file specification
SQL> attach 'filename malibu::disk$users:[remote_account]my_db.rdb';
SQL>
SQL> exit;
ALPHA4> define ll malibu::disk$users:[remote_account]my_db.rdb
%DCL-I-SUPERSEDE, previous value of LL has been superseded
ALPHA4> sql
SQL> attach 'filename ll';
SQL>

```


As a workaround, either don't use the logical to specify the path or include the database name in the logical.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.19 Query Joining Derived Tables of Union Legs With Empty Tables Returns Wrong Results

Bug 1818374

The following query, joining two derived tables containing union legs with empty tables, returns wrong results of 0 rows, instead of 1 row.

```
set flags 'strategy,detail';
select c1
  from (select v1.c1 from
        t_02,
        (select * from t_01
         union all
         select * from t_02
        ) v1
       inner join
        (select * from tt_01
         union all
         select * from tt_02
        ) as v2
       on (v1.c1 = v2.c1 and v1.c2 = v2.c2)) as tmp
 where tmp.c1 = 110759;
```

Tables:

```
0 = T_02
1 = T_01
2 = T_02
3 = TT_01
4 = TT_02
```

Merge of 1 entries

Merge block entry 1

Cross block of 3 entries

Cross block entry 1

Index only retrieval of relation 0:T_02

Index name T_02_NDX [0:0]

Cross block entry 2

Merge of 1 entries

Merge block entry 1

Merge of 2 entries

Merge block entry 1

Conjunct: 1.C1 = 110759

Index only retrieval of relation 1:T_01

Index name T_01_NDX [1:1]

Keys: <mapped field> = 110759

Merge block entry 2

Leaf#01 FFirst 2:T_02 Card=1

Bool: 2.C1 = 110759

BgrNdx1 T_02_NDX [1:1] Fan=17

Keys: <mapped field> = 110759

Cross block entry 3

Conjunct: 1.C1 = 110759

Merge of 1 entries

Merge block entry 1

Merge of 2 entries

Merge block entry 1

Conjunct: (<mapped field> = 3.C1) AND (<mapped field> = 3.C2)

Index only retrieval of relation 3:TT_01

```

Index name TT_01_NDX [2:2]
Keys: (<mapped field> = <mapped field>) AND (<mapped field> =
      <mapped field>)
Merge block entry 2
Conjunct: (<mapped field> = 4.C1) AND (<mapped field> = 4.C2)
Index only retrieval of relation 4:TT_02
Index name TT_02_NDX [2:2]
Keys: (<mapped field> = <mapped field>) AND (<mapped field> =
      <mapped field>)
0 rows selected

```

where the tables are defined as :

```

! table t_01 is empty
create table t_01 (C1 INTEGER);
create index t_01_ndx on t_01 (C1) ;

! table t_02 has 1 row
create table t_02 (C1 INTEGER, C2 TINYINT);
create index t_02_ndx on t_02 (C1) ;

insert into t_02 values (110759,9);

! table tt_01 is empty
create table tt_01 (C1 INTEGER, C2 TINYINT);
create index tt_01_ndx on tt_01 (C1, C2);

! table tt_02 has 2 rows
create table tt_02 (C1 INTEGER, C2 TINYINT);
create index tt_02_ndx on tt_02 (C1, C2);

insert into tt_02 values (110759,4);
insert into tt_02 values (110759,9);

```

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

1. The main query selects the column of a derived table with an equality predicate.
2. The main derived table joins a non-empty table (t_02) and an inner join.
3. The inner join involves a derived table of union between an empty table (t_01) and a non-empty table (t_02), and another derived table of union between an empty table (tt_01) and a non-empty table (tt_02).

As a workaround, the query works if the empty tables are loaded with some data as in the following example.

```

insert into t_01 values (110759);

select c1
  from (select v1.c1 from
        t_02,
        (select * from t_01
         union all
         select * from t_02
        ) v1
       inner join
        (select * from tt_01
         union all
         select * from tt_02
        ) as v2
       on (v1.c1 = v2.c1 and v1.c2 = v2.c2)) as tmp
 where tmp.c1 = 110759;
      C1
      110759

```

1 row selected

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.20 Left Outer Join Query With OR Predicate Returns Wrong Results

Bug 1837522

The following left outer join query with an OR predicate, having an equality predicate of a column and a constant value on the left side, and an equality predicate of a column and a subquery on the right side, returns wrong results. It should find 3 rows, but it only finds 2 rows.

```
set flags 'strategy,detail';
sel job_code, job_start, c1.employee_id, c2.employee_id
  from
  job_history as c1
 left outer join
 employees as c2 on (c1.employee_id = c2.employee_id)
 where
   c1.job_code = 'JNTR' or
   c1.job_start =
   (select max(job_start) from job_history as c3)
 ;
```

Tables:

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

Cross block of 2 entries

Cross block entry 1

Aggregate: 0:MAX (2.JOB_START)

Get Retrieval by index of relation 2:JOB_HISTORY

Index name JH_EMPLOYEE_ID [0:0]

Cross block entry 2

Conjunct: 0.JOB_START = <agg0>

Conjunct: 0.JOB_START = <agg0>

Match (Left Outer Join)

Outer loop

Conjunct: (0.JOB_CODE = 'JNTR') OR (0.JOB_START = <agg0>)

Get Retrieval by index of relation 0:JOB_HISTORY

Index name JH_EMPLOYEE_ID [0:0]

Inner loop (zig-zag)

Index only retrieval of relation 1:EMPLOYEES

Index name EMP_EMPLOYEE_ID [0:0]

C1.JOB_CODE	C1.JOB_START	C1.EMPLOYEE_ID	C2.EMPLOYEE_ID
PRSD	3-Jan-1983	00225	00225
DMGR	3-Jan-1983	00241	00241

2 rows selected

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

1. The main query is a left outer join between 2 tables with an ON clause.
2. The WHERE clause contains an OR predicate, with the left side branch being a simple equality predicate on a column, and the right branch using a sub-query in the equality predicate.

As a workaround, the query works if the left and right side of the OR predicate is swapped. For example:

```
sel job_code, job_start, c1.employee_id, c2.employee_id
  from
```

```

job_history as c1
left outer join
employees as c2
on (c1.employee_id = c2.employee_id)
   where
      c1.job_start =
        (select max(job_start) from job_history as c3)
      or
      c1.job_code = 'JNTR'
;
C1.JOB_CODE      C1.JOB_START      C1.EMPLOYEE_ID      C2.EMPLOYEE_ID
JNTR              2-Jan-1977         00223                00223
PRSD              3-Jan-1983         00225                00225
DMGR              3-Jan-1983         00241                00241
3 rows selected

```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.21 Query Using Match Strategy With DISTINCT Function Returns Wrong Results

Bugs 1891938 and 1894192

A query using the match strategy with the Distinct Function returns the wrong results, as in the following example.

```

set flags 'strategy,detail';
select count(*) from
( select distinct
      t1.ACCOUNT_ID,
      t1.SECURITY_ID
  from    T1 t1,
          T2 t2
  where   t1.SECURITY_ID = t2.SECURITY_ID
) as t ;
Tables:
  0 = T1
  1 = T2
Merge of 1 entries
Merge block entry 1
Reduce: 0.SECURITY_ID, 0.ACCOUNT_ID
Sort: 0.SECURITY_ID(a), 0.ACCOUNT_ID(a)
Conjunct: 0.SECURITY_ID = 1.SECURITY_ID
Match
Outer loop
  Sort: 1.SECURITY_ID(a)
  Get Retrieval sequentially of relation 1:T2
Inner loop (zig-zag)
  Index only retrieval of relation 0:T1
  Index name  T1_NDX1 [0:0]
ACCOUNT_ID  SECURITY_ID
A1          DE0005557508
1 row selected

```

where the tables are defined as :

```

create table T1 (
  ACCOUNT_ID      CHAR (2),
  SECURITY_ID      CHAR (12) );
create index T1_NDX      on T1 (ACCOUNT_ID, SECURITY_ID);

create table T2 (SECURITY_ID      CHAR (12) );

```

with the following contents:

```
select SECURITY_ID from T2;
```

```
SECURITY_ID
DE0005128003
DE0005557508
2 rows selected
```

```
select ACCOUNT_ID,SECURITY_ID from T1;
ACCOUNT_ID  SECURITY_ID
A1          DE0005557508
PP          DE0005128003
2 rows selected
```

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

1. The main query selects from a derived table.
2. The derived table is the output of a distinct query from T1 and T2 with a join column predicate.
3. The join column of table T1 is the second segment in index T1_NDX which is ordered by the first segment ACCOUNT_ID.
4. The order of the join column of table T2 is ascending and different from that of T2.

As a workaround, the query works if the query outline is used to apply cross strategy instead of match, as in the following example.

```
select * from
( select
  distinct
    t1.ACCOUNT_ID,
    t1.SECURITY_ID
  from
    T1 t1,
    T2 t2
  where
    t1.SECURITY_ID = t2.SECURITY_ID
) as t ;
~S: Outline "QO_325EFDCEDEBFFFA8_00000000" used
Tables:
  0 = T1
  1 = T2
Merge of 1 entries
Merge block entry 1
Reduce: 0.ACCOUNT_ID, 0.SECURITY_ID
Sort: 0.ACCOUNT_ID(a), 0.SECURITY_ID(a)
Cross block of 2 entries
Cross block entry 1
  Get Retrieval sequentially of relation 1:T2
Cross block entry 2
  Conjunct: 0.SECURITY_ID = 1.SECURITY_ID
  Index only retrieval of relation 0:T1
    Index name  T1_NDX [0:0]
-- Rdb Generated Outline : 31-JUL-2001 11:23
create outline QO_325EFDCEDEBFFFA8_00000000
id '325EFDCEDEBFFFA85200828890C4E5BA'
mode 0
as (
  query (
-- For loop
    subquery (
      subquery (
        T2 1  access path sequential
          join by cross to
        T1 0  access path index
          T1_NDX
          -- <=== change from match to cross
```


with the following content in the tables:

```
select * From t1;
ID_PRODUCTO
      8
1 row selected
```

```
select * From t2;
ID_PRODUCTO  FEC_EXPIRACION
           8           20000801
           8           20010628
2 rows selected
```

As a workaround, the query works if the predicate "OR a.id_producto = 20" is commented out from the view, as in the following example.

```
create view bug_view_good ( id_producto, total_dep, estado ) as
select
    a.id_producto,
    case
        when a.id_producto = 20 then 20
        else 15
    end as total_dep,
    case
        when b.fec_expiracion > 20001231 then 'A'
        when (a.id_producto = 15
            OR a.id_producto = 20
            ) and
            b.fec_expiracion <= 20001231
        then 'V'
    end as estado
from t1 a, t2 b
where
    a.id_producto = b.id_producto ;
```

```
select estado, sum(total_dep) from bug_view_good group by estado;
```

Tables:

```
0 = T1
1 = T2
```

Aggregate: 0:SUM (CASE (WHEN (0.ID_PRODUCTO = 20) THEN 20 ELSE 15))

Sort: CASE (WHEN (1.FEC_EXPIRACION > 20001231) THEN 'A' WHEN ((0.ID_PRODUCTO = 15) AND (1.FEC_EXPIRACION <= 20001231)) THEN 'V' ELSE NULL)(a)

Conjunct: 0.ID_PRODUCTO = 1.ID_PRODUCTO

Match

```
Outer loop      (zig-zag)
  Index only retrieval of relation 0:T1
    Index name   T1_NDX [0:0]
  Inner loop      (zig-zag)
    Get          Retrieval by index of relation 1:T2
      Index name  T2_NDX [0:0]
```

ESTADO

```
A           15
V           15
```

2 rows selected

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

1. The main query contains a GROUP BY clause and SUM aggregate function.
2. The SUM aggregate function is defined in the view as a CASE expression.
3. The column in the GROUP BY clause is defined in the view as a CASE expression which contains the same predicate from the CASE expression of the SUM aggregate.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.23 ROLLBACK Hangs Under DECdtm When Called From an ACMS CANCEL Procedure

Bug 1905068

Under certain situations, the CANCEL procedure in an ACMS application would cause the ACMS server process to hang in the RDB dispatch layer. This problem can only occur under the following circumstances:

1. The ACMS application is using 2 phase commit under DECdtm either explicitly (i.e. with a SYSS\$START_TRAN call) or implicitly (by attaching to multiple Rdb databases).
2. The CANCEL procedure contains a SYSS\$ABORT_TRAN call or ROLLBACK statement.
3. The ACMS server process has an outstanding pending request which is blocked (e.g. waiting for rows locked by another user).

If all three of these conditions occurred, the ACMS server process would hang in the CANCEL procedure even after the condition that caused the original blocking cleared.

The only workaround is to stop the ACMS server process.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.1.24 COMPUTED BY Columns Now Automatically Reserve Referenced Tables

Bug 1253235

In previous versions of Rdb, it was possible that an application could fail if a reference to a COMPUTED BY or view column required a table not specified in the RESERVING clause of the SET or DECLARE TRANSACTION statement.

The application developer may not know that a column requires these extra tables as part of the transaction, or the definition of the view or COMPUTED BY column may be changed to reference different tables after the application is in production.

The following example shows an example where a COMPUTED BY column (PRICE) requires access to a table (CASE_TABLE) that was not referenced by the RESERVING clause.

```
SQL> set transaction read only
cont>     reserving REPORT_VIEW for shared read;
SQL> select * from REPORT_VIEW order by LINE_NUM;
%RDB-E-UNRES_REL, relation CASE_TABLE in specified request is not a
relation reserved in specified transaction
SQL> rollback;
SQL> set transaction read only
cont>     reserving REPORT_VIEW, CASE_TABLE for shared read;
SQL> select * from REPORT_VIEW order by LINE_NUM;
CASE_NUM          LINE_NUM          PRICE
-----          -
1                1                7270.00
1                2                14540.00
2 rows selected
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1. Rdb now automatically reserves tables

referenced by COMPUTED BY columns for SHARED READ.

2.2 SQL Errors Fixed

2.2.1 Command Line Recall Expanded to 255 Lines

In prior releases of Oracle Rdb, the command line recall was limited to the last 20 lines. This limit has been lifted to 255 (the maximum supported by OpenVMS) for this release of Rdb.

If more recall is required then SQL provides the EDIT command to edit whole statements. This interface currently saves the last 20 commands for edit but the SET EDIT KEEP statement can be used to expand this number.

2.2.2 New Minimum Value for the INTERVAL Leading Precision

In prior releases of Oracle Rdb, the minimum value for the interval leading precision was restricted to two digits. This restriction has been removed: an interval leading precision of 1 is now supported.

The following example shows the support for the lower precision value.

```
SQL> create table TIME_CLOCK
cont>     (employee_id char(5),
cont>     clock_on          timestamp (2),
cont>     clock_off         timestamp (2),
cont>     shift_duration
cont>         computed by (clock_off - clock_on) hour (1) to minute);
SQL>
SQL> show table (column) TIME_CLOCK
Information for table TIME_CLOCK

Columns for table TIME_CLOCK:
Column Name          Data Type          Domain
-----
EMPLOYEE_ID          CHAR(5)
CLOCK_ON             TIMESTAMP(2)
CLOCK_OFF            TIMESTAMP(2)
SHIFT_DURATION       INTERVAL
                     HOUR (1) TO MINUTE
Computed:            by (clock_off - clock_on) hour (1) to minute
```

As in previous releases, if no precision is provided then a default of 2 digits will be used.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.2.3 Incorrect Processing of CASE Expression

Bug 850442

In prior releases of Oracle Rdb, the SQL interface incorrectly processed CASE expressions which included statistical functions (i.e. COUNT, MAX, MIN, AVG, STDDEV, VARIANCE and SUM).

The following example, which imbeds statistical functions in a CASE expression, caused Rdb to bugcheck:

```
select
  case
    when count(employee_id) >= 1
```

```

        then '1'
    when count(employee_id) = 0
        then '2'
    else '3'
end
from employees;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TEST]RDSBUGCHK.DMP;
%SQL-I-BUGCHKDMP, generating bugcheck dump file USER2:[TEST]SQLBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=0000000000000098, PC=000000000038B948, PS=0000001B

```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

This improved handling of statistical functions also corrects some query strategies. The following example implements a simple ABS functionality. Due to the erroneous handling of the statistical function, an extra subselect was present as shown in the optimizer STRATEGY display.

```

SQL> set flags 'strategy';
SQL> select
cont>     case
cont>         when AVG (salary_amount) < 0 then - AVG (salary_amount)
cont>         else AVG (salary_amount)
cont>     end
cont> from SALARY_HISTORY;
Cross block of 2 entries
  Cross block entry 1
    Aggregate      Get      Retrieval sequentially of relation SALARY_HISTORY
  Cross block entry 2
    Aggregate      Get      Retrieval sequentially of relation SALARY_HISTORY

  2.652896707818930E+004
1 row selected

```

The corrected SQL query now only requires a single table access.

```

Aggregate      Get      Retrieval sequentially of relation SALARY_HISTORY

  2.652896707818930E+004
1 row selected

```

Applications that encounter this type of unexpected optimizer strategy will need to be recompiled, and stored procedures and functions will need to be recreated.

2.2.4 ALTER TABLE Not Dropping NOT NULL Constraints When NULL Clause Used

In Oracle Rdb Release 7.1, new syntax was introduced to indicate that a column should allow NULL values. For instance,

```
create table MY_TABLE (my_column integer NULL);
```

This syntax is accepted for compatibility with Oracle RDBMS and on CREATE and ALTER TABLE prevents the use of the NOT NULL constraint syntax.

When used on ALTER TABLE ... ALTER COLUMN, this clause should also drop any (and all) NOT NULL constraints defined for the column. This was not performed by Rdb Release 7.1.

The following example shows that the NOT NULL constraint is now dropped by ALTER TABLE.

```

SQL> create table MY_TABLE (a integer not null);
SQL>
SQL> show table (constraint) MY_TABLE
Information for table MY_TABLE

Table constraints for MY_TABLE:
MY_TABLE_A_NOT_NULL
  Not Null constraint
  Column constraint for MY_TABLE.A
  Evaluated on UPDATE, NOT DEFERRABLE
Source:
  MY_TABLE.A NOT null

```

```

Constraints referencing table MY_TABLE:
No constraints found

```

```

SQL>
SQL> alter table MY_TABLE
cont>      alter column A NULL;
SQL>
SQL> show table (constraint) MY_TABLE
Information for table MY_TABLE

```

```

Table constraints for MY_TABLE:
No constraints found

```

```

Constraints referencing table MY_TABLE:
No constraints found

```

```

SQL>

```

This problem has been corrected in Oracle Rdb Release 7.1.0.1. This clause now implicitly drops NOT NULL constraints for the column.

NOTE: Other constraints that prevent NULL values, such as CHECK and PRIMARY KEY, are not affected by this statement. The NULL clause is not a constraint and so is not active beyond the CREATE and ALTER TABLE statements.

2.2.5 Some Constraint Definitions Not Supported for AUTOMATIC Columns

In Oracle Rdb Release 7.1, attempts to define UNIQUE, PRIMARY KEY or FOREIGN KEY constraints for AUTOMATIC columns would result in an error.

In the following example, the programmer desired an automatically generated unique number as a PRIMARY KEY:

```

SQL> create sequence s1;
SQL> create table t(a automatic as s1.nextval primary key);
%SQL-F-PKCONSNOTCB, Computed column may not be a primary key

```

Only NOT NULL and CHECK constraints were allowed for AUTOMATIC columns.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. All types of constraints are now permitted for AUTOMATIC columns.

2.2.6 %RDB-E-NO_DIST_BATCH_U Error When Executing SET TRANSACTION

Bug 1921672

If a SET TRANSACTION statement was executed to start a distributed transaction (2 phase commit) and which specified certain table partitions, an error was inappropriately returned. Specifically, if partition 14 was named, Rdb would return a %RDB-E-NO_DIST_BATCH_U error and not start the transaction.

For example, suppose an interactive SQL session has two databases attached (this implicitly starts a DECdtm distributed transaction), the following SQL would fail as shown.

```
SQL>SET TRANSACTION READ WRITE WAIT ISOLATION LEVEL READ COMMITTED -  
RESERVING DB2.MY_TABLE PARTITION(14) FOR EXCLUSIVE WRITE;  
%RDB-E-NO_DIST_BATCH_U, no distributed transaction is allowed with the  
recovery mechanism disabled
```

This query will now execute normally and start a distributed transaction.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.2.7 Select With Identical "not in" Clauses

A SQL query which contained two identical "not in" clauses would cause an application to crash, terminate or bugcheck.

This problem started in Oracle Rdb V7.0.

An example of this type of query follows.

```
select count(*) from JOBS  
  where JOB_CODE not in ('A', 'B')  
     and JOB_CODE not in ('A', 'B');
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.2.8 Keyword Matching Now Reported by Interactive SQL

In prior versions of Oracle Rdb, the keyword abbreviation and matching support in interactive SQL would discard extraneous characters from a token if an expected keyword matched the leading prefix. This was confusing in some cases. Interactive SQL now generates an informational message to clearly state the substitution.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

This example shows the informational message generated when extra characters are trimmed from the keyword.

```
SQL> create trigger mytrigger  
cont>      after updatete on mytable2  
%SQL-I-SPELLCORR, identifier UPDATETE replaced with UPDATE  
cont>      (insert into mytable values (mytable2.a, 'Any', 'Value'))
```

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

2.2.9 CREATE MODULE Bugchecks When a Subselect is Used as a Parameter DEFAULT

In a CREATE MODULE definition, if a subselect was used as a parameter DEFAULT, the create module bugchecked with the following error message:

```
%SQL-F-BUGCHK, There has been a fatal error. Please contact your Oracle support representative. SQL$BLRXPR - 15
```

An example follows:

```
SQL> create module DEF_MOD
cont>
cont> procedure DEF1
cont>      (in :a integer
cont>      default (select count(*) from rdb$database));
cont> trace :a;
cont>
cont> end module;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file
device:[directory]SQLBUGCHK.DMP;
%SQL-F-BUGCHK, There has been a fatal error. Please contact your Oracle support representative. SQL$BLRXPR - 15
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1. The CREATE MODULE definition no longer bugchecks.

2.2.10 Obsolete Metadata Errors When Using Rdb SQL V7.1 to Access Oracle Rdb V7.0 Databases

Bug 1994383

When using Oracle Rdb SQL V7.1 to access an Oracle Rdb V7.0 database, obsolete metadata errors were generated when trying to CREATE a TABLE, a VIEW, and/or a DOMAIN.

Specifically, when CREATing a TABLE or a VIEW, the following error message would be generated:

```
CREATE TABLE T (id int);
%RDB-F-OBSOLETE_METADA, request references metadata objects that no longer exist
-RDMS-F-TABNOTDEF, relation RDB$SEQUENCES is not defined in database
CREATE VIEW V as select * from employees;
%RDB-F-OBSOLETE_METADA, request references metadata objects that no longer exist
-RDMS-F-TABNOTDEF, relation RDB$SEQUENCES is not defined in database
```

When trying to CREATE a domain, the following error message would be generated:

```
create domain dom_test int;
%RDB-F-OBSOLETE_METADA, request references metadata objects that no longer exist
-RDMS-F-TABNOTDEF, relation RDB$TYPES is not defined in database
```

These problems have been corrected. SEQUENCES and TYPES are Release 7.1 features and the Rdb SQL code base has been corrected to insure that queries utilizing these features are only performed against V7.1 databases. Thus, error messages are no longer generated.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.2.11 SQL\$PRE and SQL\$MOD Performance Improvements

Bug 2032243

The performance of the SQL precompiler and the SQL module language compiler has been improved in Oracle Rdb Release 7.1.0.1. This improvement is typically seen as a dramatic reduction in CPU consumption and elapsed time when using the compilers.

Note as well that the size of the SQL\$PRE71.EXE and SQL\$MOD71.EXE images has been reduced by nearly 50%.

2.2.12 Incompatible Character Sets Not Detected by SQL Interface

In prior versions of Oracle Rdb, the SQL UNION operator would accept incompatible character sets for merging. This incompatibility was only detected at runtime by the Rdb server.

```
SQL> select _dec_mcs'aa' from rdb$database
cont> union
cont> select _kanji'bb' from rdb$database;
%RDB-E-CONVERT_ERROR, invalid or unsupported data conversion
-RDMS-E-CSETBADASSIGN, incompatible character sets prohibit the requested
assignment
```

With this release of Rdb, the SQL interface now detects this error and reports an error indicating the incompatibilities.

```
SQL> select _dec_mcs'aa' from rdb$database
cont> union
cont> select _kanji'bb' from rdb$database;
%SQL-F-INCCSCON, Incompatible character set concatenation between DEC_MCS and
KANJI
```

In addition, SQL now derives a new target character set for the UNION select values by using a character set that is compatible with both. For instance, DEC_KANJI includes the full ASCII character set so it will be chosen as the result character set when ASCII and DEC_KANJI are merged in a UNION operator.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.2.13 SQLMOD Fails to Set Default Character Set Correctly

A problem within SQLMOD prevented the correct default character set from being set for the module compilation if a character set other than DEC_MCS was specified as the DEFAULT CHARACTER SET in the module header.

A check of the listing file will show that the default character set has not been set correctly and, due to this problem, SQL-F-INCCSASS errors may be raised during the module compilation.

For example, the following module tries to set the default character set to SHIFT_JIS, however, the compilation of the module results in compilation errors.

```
$type A.SQLMOD
```

```

.
.
.
DECLARE MODULE
    DIALECT SQL92
    DEFAULT CHARACTER SET SHIFT_JIS
    NATIONAL CHARACTER SET SHIFT_JIS
    IDENTIFIER CHARACTER SET SHIFT_JIS
    LITERAL CHARACTER SET SHIFT_JIS
    DISPLAY CHARACTER SET SHIFT_JIS
    AUTHORIZATION RDB$DBHANDLE
    CHARACTER LENGTH CHARACTERS
.
.
.
INSERT INTO SHIFTJIS_TABLE ( SHIFTJIS_COL1)
VALUES
    (:LAST_NAME);
.
.
.
$ SQLMOD/LIST=A.LIS A.SQLMOD
.
.
.
    ( SHIFTJIS_COL1)
    1
%SQL-F-INCCSASS, (1) Incompatible character set assignment between
SHIFTJIS_COL1 and :LAST_NAME

$ type A.LIS

.
.
.
Command Line Summary:

    SJIS_MOD2_M.SQLMOD /LIST

    /G_FLOAT
    /WARN=(WARNING, DEPRECATED)
    /NOFLAG_NONSTANDARD
    /CONSTRAINT_DEFAULT=DEFERRED
    /NOCONNECT
    /INIT_HANDLES
    /NORESTRICT_INVOKER
    /CHECK_RW
    /ANSI_VIEWS
    /ANSI_DATE
    /ANSI_QUOTING
    /ANSI_PARAMETERS
    /QUERY_ESTIMATES
    Default Character Set: DEC_MCS
    National Character Set: SHIFT_JIS
    Identifiers Character Set: SHIFT_JIS
    Literals Character Set: SHIFT_JIS
    Character Length in Characters
.
.
.

```

Note that the Default Character Set as shown in the listing file has not been set correctly.

A workaround for this problem is to use NAMES ARE in the module header to set the desired character set prior to setting the Default Character Set.

```
.  
. .  
. .  
DECLARE MODULE  
    DIALECT SQL92  
    NAMES ARE SHIFT_JIS  
    DEFAULT CHARACTER SET SHIFT_JIS  
    NATIONAL CHARACTER SET SHIFT_JIS  
    IDENTIFIER CHARACTER SET SHIFT_JIS  
    LITERAL CHARACTER SET SHIFT_JIS  
    DISPLAY CHARACTER SET SHIFT_JIS  
    AUTHORIZATION RDB$DBHANDLE  
    CHARACTER LENGTH CHARACTERS  
. .  
. .
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.3 Oracle RMU Errors Fixed

2.3.1 RMU Extract Not Formatting View Column Expressions Correctly

Bug 1832240

In prior releases of Oracle Rdb, the RMU Extract command did not correctly format VIEW definitions that contained computed expressions in the SELECT clause, such as that shown below.

```
create view V1 (F3) as
  select sum (F3 /
            case (select cast (F1 as integer) from T1
                  where F2 = 'STR_VALUE')
                when 0 then 1
                when 1 then 10
                when 2 then 100
                when 3 then 1000
                when 4 then 10000
                when 5 then 100000
                else 0
            end)
  from T2;
```

This example was extracted below. Note the incorrect formatting of the expression and the missing separating white space. This made the generated definition illegal.

```
create view "V1"
(F3) as
select

sum((C2.F3 / case (select CAST(C3.F1 AS INTEGER) from T1 C3where (C3.F2 =
'STR_VALUE')) when 0 then 1 when 1 then 10 when 2 then 100 when 3 then 1000
when 4 then 10000 when 5 then 100000 else 0end)) from T2 C2;
```

The only workaround for this problem is to manually edit the definition after extracting with RMU Extract or to revert to the original view source.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.3.2 RMU/UNLOAD/AFTER_JOURNAL Fragmented Records Clarification

The RMU /UNLOAD /AFTER_JOURNAL Utility uses additional CPU and memory resources while processing and unloading fragmented records from the after-image journal file. As record fragments are found within a transaction, they are buffered in memory on a "fragment" queue. After all non-fragmented records from the transaction have been output, the fragmented records are reconstructed and output.

Because the fragments are buffered in memory, additional process page file quota may be required when unloading transactions that have a large number of record fragments. Also additional process working set quota may be required in order to limit process page faulting.

2.3.3 RMU/DUMP/BACKUP Did Not Check the VMS BYPASS Privilege

Bug 1966820

The RMU/DUMP/BACKUP command for Oracle Rdb RMU did not check if the user process was granted the VMS BYPASS privilege if the user was not granted the necessary RMU access privileges to the database backup file created by the RMU/BACKUP command. Therefore, the RMU/DUMP/BACKUP command did not execute even though the BYPASS privilege should have allowed the user to execute the command.

The following example shows that even though the BYPASS privilege should have allowed the user to override the lack of RMU privileges to access the backup file, the user was denied access by the RMU/DUMP/BACKUP command.

```
$RMU/DUMP/BACKUP PERSONNEL
%RMU-I-DMPTXT_163, No dump option selected. Performing read check.
%RMU-F-NOPRIVERR, no privileges for attempted operation
%RMU-F-FTL_DUMP, Fatal error for DUMP operation at 30-AUG-2001 16:42:17.96
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.3.4 RMU/BACKUP Invalid Volume 1 Tape Label When Used With COMPAQ SLS

Bug 1969648

The RMU/BACKUP command for Oracle Rdb RMU, when used with COMPAQ SLS, did not detect the case where SLS did not provide a new VOL1 label to replace the VOL1 label that Rdb RMU/BACKUP was about to write to the first tape volume. RMU/BACKUP therefore wrote an 80 character label buffer to the tape that contained invalid characters. This caused an RMU-F-LABELERR when the tape was restored using RMU/RESTORE.

This problem only happens when RMU/BACKUP is run with COMPAQ SLS and when COMPAQ SLS does not modify the 80 character VOL1 label that RMU/BACKUP writes to the first tape volume.

The following example shows that although the RMU/BACKUP with SLS did not show an error, a VMS DUMP command of the BACKUP tape shows an invalid label on the first backup tape volume. Therefore, RMU/RESTORE returns an RMU-F-LABELERR.

Here is an example of a valid RMU/BACKUP tape label on the first tape volume created after this problem was fixed (note that this is just an example and correct labels may vary).

```
$ dump tapedevice:

Dump of device tapedevice: on 29-AUG-2001 11:44:32.94

Block number 1 (00000001), 80 (0050) bytes

20202020 20202020 20202020 20202020 20202020 20203035 30494241 314C4F56

VOL1ABI050                                000000

20202020 20202020 20202020 20202020 20202020 20202020 20202020 20202020

                                000020
```

```
33202020 20202020 20202020 202020203..... 000040
```

Here is an example of an invalid RMU/BACKUP tape label on the first tape volume that has been created by this problem (note that this is just an example and incorrect labels may vary).

```
$ dump tapedevice:
```

```
Dump of device tapedevice: on 29-AUG-2001 ...
```

```
Block number 1 (00000001), 80 (0050) bytes
```

```
00000000 00001F00 00000000 00183390 FFFFFFFF FFFFFFFE 00000000 000119D8
```

```
Ø.....3..... 000000
```

```
00000000 00000000 00000000 007EBFC0 00000000 00183390 00000000 00010DC0
```

```
À.....3.....À¿~..... 000020
```

```
00000000 00000000 00000000 00000D05
```

```
..... 000040
```

Here is an example of the RMU-F-LABELERR returned by RMU/RESTORE.

```
%RMU-F-LABELERR, error in tape label processing on  
tapedevice:[000000]SAMPLE_DB.RBF;  
-RMU-F-NOTANSI, tape is not valid ANSI format  
%RMU-F-FATALERR, fatal error on RESTORE  
%RMU-F-FTL_RSTR, Fatal error for RESTORE operation ...
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.3.5 RMU/ANALYZE/CARDINALITY Fails on Databases With Local Temporary Tables

Bug 2019322

RMU/Analyze/Cardinality, when attempting to process LOCAL temporary tables, generated an error and failed to execute.

```
$ rmu/anal/card sql$database  
%RDMS-E-BAD_CODE, corruption in the query string  
%RMU-F-FATALRDB, Fatal error while accessing Oracle Rdb.  
%RMU-F-FTL_ANA, Fatal error for ANALYZE operation at 27-SEP-2001 13:34:25.79
```

RMU has now been corrected to ignore temporary tables as well as views.

The workaround for this problem is to use the RMU/SHOW OPTIMIZER/STATISTIC=CARD command or the RMU/COLLECT OPTIMIZER_STATISTICS command if RMU/ANALYZE/CARDINALITY/UPDATE was tried.

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.3.6 File Name Not Displayed By RMU /RESTORE for Extend Failure

Bug 1822217

When an RMU /RESTORE operation is unable to extend a storage area, it is possible for the error message displayed to not include the name of the file. This may make it difficult to determine which device has inadequate free space. In the following example note that the name of the file is not displayed.

```
$ RMU /RESTORE ...
%RMU-F-FILACCERR, error extending file
-SYSTEM-W-DEVICEFULL, device full; allocation failure
%RMU-F-FTL_RSTR, Fatal error for RESTORE operation at 17-JUN-2001 03:08:55
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1. RMU /RESTORE now displays the file name, where possible, during a failed file extend operation.

2.3.7 RMU/SHOW STATISTICS Allowed Suspend of Disabled ABS

Previously, the RMU /SHOW STATISTICS Utility allowed the user to suspend AIJ Backup Server (ABS) operations on a node even when the ABS was disabled. This could lead to confusing errors during later manual AIJ backup operations.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. The RMU /SHOW STATISTICS Utility now does not allow the ABS to be suspended when it is not enabled.

2.3.8 RMU/COPY/BLOCKS_PER_PAGE Can Corrupt Copied Database Uniform Areas

Bug 2028181

For the RMU/COPY command, if the "/blocks_per_page" qualifier was not specified for a particular storage area but was for all database storage areas, database corruption of uniform storage areas occurred to the copied database. As documented in the Oracle Rdb RMU Reference Manual for the RMU/COPY command, BLOCKS PER PAGE can only be changed for MIXED storage areas, not UNIFORM storage areas. But when the "/blocks_per_page" qualifier was used for all storage areas, RMU incorrectly bypassed the check for UNIFORM storage areas and attempted to change the BLOCKS PER PAGE setting for UNIFORM as well as MIXED storage areas. This caused the database corruption of the moved copy of the database. Now, the number of BLOCKS PER PAGE will be changed only for MIXED storage areas and a warning message will be output for each UNIFORM storage area that BLOCKS PER PAGE cannot be changed for that area since it is a UNIFORM database storage area.

The following example shows that since /BLOCKS_PER_PAGE=3 was specified for all storage areas in the MF_PERSONNEL database, it caused the database corruption problem for the uniform storage areas in the copied database.

```
$ RMU/COPY/DIR=TMPDIR/ROOT=TMPDIR:MFP1 /NOLOG /BLOCKS_PER_PAGE=3 MF_PERSONNEL
%RMU-W-BADPTLARE, invalid larea for uniform data page 5 in storage area 1
%RMU-W-BADPTLAR2, SPAM larea_dbid: 16385, page larea_dbid: 1
%RMU-W-BADPTLARE, invalid larea for uniform data page 149 in storage area 1

$ RMU/VERIFY/ALL TMPDIR:MFP1
```

```

%RMU-I-BGNROOVER, beginning root verification
%RMU-I-ENDROOVER, completed root verification
%RDB-W-NO_RECORD, access by dbkey failed because dbkey is no longer associated
with a record
-RDMS-F-NODBK, 61:1179:0 does not point to a data record
%RMU-E-ERRRDBREL, error accessing RDB$RELATIONS relation

```

The following example shows that the problem is now fixed.

```

$ RMU/COPY/DIR=TMPDIR/ROOT=TMPDIR:MFP1 /NOLOG /BLOCKS_PER_PAGE=3 MF_PERSONNEL
%RMU-W-UNIFORMBLOCKS, BLOCKS PER PAGE cannot be changed for uniform storage
area RDB$SYSTEM
%RMU-W-UNIFORMBLOCKS, BLOCKS PER PAGE cannot be changed for uniform storage
area MF_PERS_SEGSTR

$ RMU/VER/ALL TMPDIR:MFP1

```

To avoid this problem, specify /BLOCKS_PER_PAGE for each individual storage area in the RMU/COPY command, not as a default for all storage areas.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. A warning message is displayed and the uniform storage area page size is not changed.

2.3.9 DROPPed Storage Area and RMU /VERIFY in Cluster

Bug 1421362

Previously, when a database was opened in a cluster environment, it was possible for the RMU /VERIFY command to be unable to open storage area files when storage areas were moved or dropped on another node in the cluster.

For example, consider the following sequence of events on a two node cluster (consisting of NODE1 and NODE2):

```

Node1$: RMU /OPEN MFP
Node2$: RMU /OPEN MFP
Node1$: SQL$ ALTER DATABASE FILENAME MFP DROP STORAGE AREA U1;
Node2$: RMU /VERIFY MFP
.
.
.
%RMU-F-OPNFILERR, error opening file U1.RDA
%RMU-F-FILNOTFND, file not found
%RMU-E-BDAREAOPN, unable to open file U1.RDA for storage area
%RMU-F-ABORTVER, fatal error encountered; aborting verification

```

This problem has been corrected in Oracle Rdb Release 7.1.0.1. The RMU /VERIFY Utility now correctly detects storage areas that have been dropped or moved.

2.3.10 RMU /VERIFY Checks All Storage Area Files First

Bug 671681

Previously, the RMU /VERIFY command would abort and return a fatal error to the user when a storage area file was unable to be opened (for example, when the storage area file had been deleted). However, other

storage areas were not checked, leading to the possibility that not all problems with missing storage area files were reported.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. The RMU /VERIFY Utility now checks all storage area files and reports problems while opening the files before returning a fatal error. This makes it much easier to know what files must be restored with the RMU /RESTORE command.

2.3.11 RMU/SHOW STATISTICS Multi-Page Report File

Previously, the RMU /SHOW STATISTICS Utility only displayed the first page ("Page: 1 of 1") of multi-page displays. This made it difficult, at times, to find specific information.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. The RMU /SHOW STATISTICS Utility now writes all pages of multi-page displays to the report file. Note that for some screens (storage area information, row cache information, and so on), there can be a significant amount of data written and this can result in a dramatic increase in the size of the report file.

2.3.12 Area Locks Demoted Statistic Not Always Correctly Incremented

Previously, the "locks demoted" statistic for "area" locks was not always correctly incremented. This could occur, for example, when a read-only transaction was started when the previous transaction was a read-write transaction. The "locks promoted" statistic could have been incorrectly incremented in this case. This, in turn, led to potentially confusing results when comparing the "locks promoted" rate with the "locks demoted" rate for "area" locks in the "RMU/SHOW STATISTICS" facility.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. The correct statistic is now incremented when an "area" lock is demoted from one lock mode to a lower mode.

2.3.13 RMU /BACKUP /ONLINE /NOQUIET_POINT Fails

Oracle Rdb Release 7.1.0 introduced a potential regression where the RMU /BACKUP /ONLINE /NOQUIET_POINT command may fail with an incorrect error message indicating that it is unable to write to the root file. This is an example of the incorrect error from the RMU /BACKUP command:

```
$ RMU /BACKUP /ONLINE /NOQUIET_POINT MFP NLA0:MFP
%RMU-F-FILACCERR, error writing file DUA0:[DB]MFP.RDB;1
%RMU-F-FTL_BCK, Fatal error for BACKUP operation ...
```

This problem has been corrected in Oracle Rdb Release 7.1.0.1.

2.4 LogMiner Errors Fixed

2.4.1 LogMiner Compresses Pre-Delete Record Content

Previously, when the Oracle Rdb LogMiner(TM) feature was enabled, the pre-delete record contents were not compressed prior to being journaled. Because of this, it was possible for AIJ files to grow excessively if many large records were being deleted.

This problem has been corrected in Oracle Rdb Release 7.1.0.1. When the Oracle Rdb LogMiner feature is enabled, pre-delete record contents are now correctly compressed. Because of the difference in pre-delete record contents in an AIJ file, it is important that AIJ files created with prior versions of Oracle Rdb be processed with the matching version of the Oracle Rdb LogMiner (RMU /UNLOAD /AFTER_JOURNAL command).

When using the Oracle Rdb LogMiner feature, existing AIJ files should be backed up and processed prior to upgrading to this release of Oracle Rdb.

Failure to use the correct version of the Oracle Rdb LogMiner to process an AIJ file typically results in RMU-W-RECVRDIF warnings when pre-delete record contents are being processed.

LogMiner AIJ files not compatible

When the Oracle Rdb LogMiner(TM) feature is being used, AIJ files from this version of Oracle Rdb are not compatible with the Oracle Rdb LogMiner feature from prior versions of Oracle Rdb. Only the Oracle Rdb LogMiner feature is affected; AIJ recovery is not affected. If the Oracle Rdb LogMiner feature is not enabled for a database, there is no difference in the format or content of an AIJ file.

2.5 Optimizer Problems Fixed in Oracle Rdb Release 7.1.0.

The following Optimizer Bugs were fixed in Oracle Rdb Release 7.1.0 but the release notes were inadvertently left out.

2.5.1 Query Having OR Compound Predicates With Subquery Returns Wrong Results

Bug 1527102

The following query contains the OR of three predicates: one of which is based on the results of a subquery; one of which is a filter predicate of the form column = literal; and one of which is a constant of the form literal = literal. The query should return 1 row.

```
set flags 'strategy,detail';
select t1.hmcnr from t1 t1
  where t1.ean='5410103914978' and
  (t1.shop_class = (select sho.shop_class from r_shop sho
                    where sho.shop='460')
   or t1.shop_class='A'
   or 'XXX'='460');
```

Tables:

```
0 = t1
1 = R_SHOP
```

Cross block of 2 entries

Cross block entry 1

```
Aggregate: (VIA)
Conjunct: 1.SHOP = '460'
Conjunct: 'XXX' = '460'
```

Get Retrieval sequentially of relation 1:R_SHOP

Cross block entry 2

```
Conjunct: (0.ean = '5410103914978') AND ((0.shop_class = {subselect}) OR
(0.shop_class = 'A') OR ('XXX' = '460'))
```

Get Retrieval sequentially of relation 0:t1

HMCNR

```
45281
45134
```

2 rows selected

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

1. A filter predicate is ANDed to an OR compound predicate
2. The OR compound predicate contains a subquery predicate, a couple of filter predicates and a constant predicate

As a workaround, the query works if the constant predicate is removed.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.2 Query Using OR/AND Predicates With EXISTS Clause Returns Wrong Results

Bug 1569972

The following query using AND/OR predicates with an EXISTS clause should return 1 row:

```
set flags 'strategy,detail';

select t1.c1 from t1 t1, t2 t2 where
((t2.c4 = 1 and
  t2.c5 = 5 and
  not exists (select * from t2 t2a
              where t2a.c4 = 4 and t2a.c5 = 5)) or
 (t2.c4 = 4 and t2.c5 = 5))
and t1.c1 = t2.c6
;
Tables:
0 = T1
1 = T2
2 = T2
Cross block of 3 entries
Cross block entry 1
  Conjunct: {subselect} = 0
    Aggregate-F1: (COUNT-ANY)      Index only retrieval of relation 2:T2
    Index name  T2_H [2:2]
    Key: (2.C4 = 4) AND (2.C5 = 5)
Cross block entry 2
  Conjunct: (1.C4 = 1) OR (1.C4 = 4)
  Conjunct: 1.C5 = 5
  Conjunct: {subselect} = 0
  Get      Retrieval by index of relation 1:T2
    Index name  T2_H [(2:2)2] Bool
    Key: ((1.C4 = 1) AND (1.C5 = 5)) OR ((1.C4 = 4) AND (1.C5 = 5))
    Bool: 1.C5 = 5
Cross block entry 3
  Index only retrieval of relation 0:T1
  Index name  T1_H [1:1]
  Key: 0.C1 = 1.C6
0 rows selected
```

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

1. OR parent predicate with AND predicates on each branch
2. One of the OR branches also includes a subquery, such as NOT EXISTS
3. A second AND predicate is appended after the OR parent predicate

As a workaround, the problem can be corrected if you move the second AND predicate to the front of the OR parent predicate, as follows:

```
set flags 'strategy,detail';

select t1.c1 from t1 t1, t2 t2 where
t1.c1 = t2.c6 and
((t2.c4 = 1 and
  t2.c5 = 5 and
  not exists (select * from t2 t2a
              where t2a.c4 = 4 and t2a.c5 = 5)) or
 (t2.c4 = 4 and t2.c5 = 5))
;

```

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.3 Query Using German Collating Sequence Returns Wrong Results

Bug 1530947

The following query, in a database where the German Collating Sequence is used by default, returns wrong results (should return some rows):

```
SELECT p.datum, p.produkt, p.abtlg, p.stelle
FROM v_team_datum p,
     produkte g
where
     p.abtlg=g.abtlg ;
Conjunct
Match
Outer loop
  Sort      Conjunct      Aggregate      Sort      Conjunct
  Leaf#01 BgrOnly PROD_DATEN Card=24063
         BgrNdx1 IDX_PROD_DATEN_SORT [1:1] Fan=8
Inner loop      (zig-zag)
  Conjunct      Get      Retrieval by index of relation PRODUKTE
  Index name    IDX_PRODUKTE_SORT [0:0]
0 rows selected
```

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

1. The query is a simple join between a view and one table, with the join predicate of CHAR data type
2. The optimizer uses a match strategy to join them, where a comparison of the join keys requires the process of encoding the CHAR data type into the German collating sequence

As a workaround, the query works if a view with the same attributes as the table is used instead of the table itself, as in the following example:

```
SELECT p.datum, p.produkt, p.abtlg, p.stelle
FROM v_team_datum p,
     view_produkte g
where
     p.abtlg=g.abtlg ;
Cross block of 2 entries
Cross block entry 1
  Conjunct      Aggregate      Sort      Conjunct
  Leaf#01 BgrOnly PROD_DATEN Card=24063
         BgrNdx1 IDX_PROD_DATEN_SORT [1:1] Fan=8
Cross block entry 2
  Leaf#02 FFirst PRODUKTE Card=25
         BgrNdx1 IDX_PRODUKTE_SORT [3:3] Fan=6
```

The query works because the optimizer applies a cross strategy instead of a match strategy.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.4 Left Outer Join Query Returns Wrong Results When ON Clause Evaluates to False

Bug 1581632

The following left outer join query returns wrong results when the join conditions in the ON clause evaluate to false for all rows:

```

set flags 'strategy,detail';
select tt.employee_id, tt.last_name, jh.job_code
from
  (select e.employee_id, e.last_name
   from degrees d, employees e where
    e.employee_id = '00354'
    and d.employee_id = e.employee_id) as tt
left outer join
job_history jh
  on tt.last_name = '?' and          <----
   jh.job_code = tt.employee_id;    <----

```

Tables:

```

0 = DEGREES
1 = EMPLOYEES
2 = JOB_HISTORY

```

Cross block of 2 entries (Left Outer Join)

Cross block entry 1

Conjunct: "tt.last_name" = '?'

Merge of 1 entries

Merge block entry 1

Cross block of 2 entries

Cross block entry 1

Get Retrieval by index of relation 1:EMPLOYEES

Index name EMPLOYEES_HASH [1:1] Direct lookup

Key: 1.EMPLOYEE_ID = '00354'

Cross block entry 2

Index only retrieval of relation 0:DEGREES

Index name DEG_EMP_ID [1:1]

Key: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID

Cross block entry 2

Conjunct: ("tt.last_name" = '?') AND

(2.JOB_CODE = tt.employee_id)

Get Retrieval by index of relation 2:JOB_HISTORY

Index name JH_EMPLOYEE_ID [0:0]

0 rows selected

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

1. Left outer join query on a subquery and job_history of mf_personnel database
2. ON clause containing two or more predicates, and the ON clause evaluates to false for all rows, for example:

```
"last_name" = '?' and jh.job_code = tt.employee_id
```

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.5 Query With Two IN Clauses on Two Subqueries Returns Wrong Results

Bug 1585429

The following query with two IN clauses on two subqueries with different match keys, returns a count of 0 when it should return a non-0 count:

```

SELECT count(*) FROM t1
WHERE
  subclass_id IN (SELECT DISTINCT subclass_id
                  FROM t2
                  WHERE class_id = 'CAJ_C01#')
AND
  recipe_id IN (SELECT recipe_id
                FROM t3
                WHERE eqp_id = 'CAR-02C'
                )
;
Aggregate      Conjunct
Match
  Outer loop
    Conjunct
      Match
        Outer loop
          Get      Retrieval by index of relation t1
                  Index name t1_ndx [0:0]
        Inner loop (zig-zag)
          Aggregate-F1 Conjunct
          Index only retrieval of relation t2
          Index name t2_ndx [0:0]
        Inner loop (zig-zag)
          Aggregate-F1 Conjunct      Get
          Retrieval by index of relation t3
          Index name t3_ndx [1:1]

          0
1 row selected

```

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

1. Two different IN clauses on two subqueries, with different match keys
2. The query applies a match strategy where the outer leg uses the match key (subclass_id) of another match stream that is different from the other key (recipe_id) of the inner leg without sorting the results of the outer leg using the match key (subclass_id).

Oracle Rdb7 Release 7.0.5 applies a sort node on the outer leg and thus returns the correct results.

As a workaround, use a query outline to change the strategy to cross from match.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.6 Query Having Same SUBSTRINGs Within CASE Expression Returns Wrong Results

Bugs 1489972, 1485656, 975091

The following queries, containing the same SUBSTRING expressions within a CASE expression, return wrong results.

The following example shows two simple queries (from Bug 1485656 and Bug 975091) having the same subexpression (SUBSTRING) appearing more than once within the CASE expression. The query in the case of Bug 1489972 is more complicated and thus omitted here. It contains unions of several subselect queries with nested views and SUBSTRING/CASE expressions.

```

! Bug 1485656
! should return the value 1 for the content of y
! ~Xt: Content of y = 1
!
set FLAGS 'TRACE'
declare :x char(2);
declare :y char(1);
begin
set :x='21';
set :y= case
    when ((substring(:x from 1 for 1)='1') and
          (substring(:x from 2 for 1)='1') )
        then '0'
    else
        (substring(:x from 2 for 1))
    end;
trace 'Content of y = ', :y ;
end;
The output is:
~Xt: Content of y =

! Bug 975091
! should return the value of 295 for the column RESP
!
create table t1 (c1 char(12));
insert into t1 value ( '29500000199');

select substring( c1 from 1 for 3) ress,
       case
         when 'a' = 'c' and (substring(c1 from 1 for 3)) = '295'
           then 'a'
         when 'c' = 'c'
           then (substring(c1 from 1 for 3))
         else ' '
       end resp
  from t1;
RESS  RESP
295
1 row selected

```

The key parts of these queries which contributed to the situation leading to the errors are these:

1. CASE expression contains several similar expressions
2. The expression in the WHEN clause is shared in the same clause of another WHEN clause (in the case of Bug 975091)
3. The expression in the WHEN clause is shared in another part of the CASE statement, such as an ELSE clause (in the case of Bug 1485656)

In the case of Bug 1485656, a workaround is to use an IF instead of a CASE statement to get the correct results:

```

set FLAGS 'TRACE'
declare :x char(2);
declare :y char(1);
begin
set :x='21';
    if ((substring(:x from 1 for 1)='1') and
        (substring(:x from 2 for 1)='1') )
        then
            set :y='0';
        else
            set :y=(substring(:x from 2 for 1));
        end if;

```

```
trace 'Content of y:',:y;
end;
```

Another workaround is to use temporary variables for the substrings.

In the case of Bug 975091, the workaround is to swap the WHEN clauses, as in the following example:

```
select substring( c1 from 1 for 3) ress,
       case
         when 'c' = 'c'
           then (substring(c1 from 1 for 3))           ! <= 1st
         when 'a' = 'c' and (substring(c1 from 1 for 3)) = '295' ! <= 2nd
           then 'a'
         else ' '
       end resp
from t1;
```

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.7 Aggregate Query With Nested MIN Function Returns Wrong Results

Bug 1408892

The following query should return the value of ADMN for min(d1.department_code):

```
create index dept_managerid_code_ndx on departments
(manager_id,department_code);

select min(d1.department_code),
       min((select min (d2.department_code)
            from departments d2
            where d1.manager_id = d2.manager_id AND
                  d2.budget_actual > 0))
       from departments d1;
              NULL          NULL
1 row selected
```

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

1. The subselect query has "where" predicates which cause the query to return 0 rows, e.g. "d2.budget_actual > 0"
2. The subselect query contains an aggregate function, e.g. MIN
3. The subselect query is wrapped inside another aggregate function, e.g. MIN

As a workaround to this problem, the query works if the MIN function is removed from the column 'd2.department_code' in the inner subselect, as seen in the following example.

```
select min(d1.department_code),
       min((select d2.department_code
            from departments d2
            where d1.manager_id = d2.manager_id AND
                  d2.budget_actual > 0))
       from departments d1;
```

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.8 Query with UNION Subselect Returns Wrong Results

Bug 1656974

The following query with UNION subselect should return 0 rows.

```
set flags 'strategy,detail';
select ps.id, ps.kbn, ps.ymd
  from (select ps1.id,
              ps1.kbn,
              '99999999'
        from ps ps1, pm pm
        where pm.id = ps1.id
 union all
 select ps2.id,
        ps2.kbn,
        ps2.end_ymd
        from ps ps2, pm pm
        where pm.id = ps2.id)
      as ps (id, kbn, ymd)
 where ps.id = '021023307' and
        ps.ymd > '12345678' and
        ps.kbn in ('1','2') ;
```

Tables:

- 0 = PS
- 1 = PM
- 2 = PS
- 3 = PM

Merge of 1 entries

Merge block entry 1

Merge of 2 entries

Merge block entry 1

Conjunct: 1.id = 0.ID

Match

Outer loop (zig-zag)

Conjunct: 0.ID = '021023307'

Conjunct: '99999999' > '12345678'

Get Retrieval by index of relation 0:PS

Index name IDX_PS_2 [1:1] Bool

Key: <mapped field> = '021023307'

Bool: '99999999' > '12345678'

Inner loop (zig-zag)

Index only retrieval of relation 1:PM

Index name IDX_PM_0 [0:0]

Merge block entry 2

Conjunct: 3.id = 2.ID

Match

Outer loop (zig-zag)

Conjunct: (2.ID = '021023307') AND (2.end_ymd > '12345678')

AND ((2.kbn = '1') OR (2.kbn = '2'))

Get Retrieval by index of relation 2:PS

Index name IDX_PS_2 [2:1]

Key: (<mapped field> = '021023307') AND (<mapped field> > '12345678')

Inner loop (zig-zag)

Index only retrieval of relation 3:PM

Index name IDX_PM_0 [0:0]

ID	KBN	YMD
021023307	0	99999999

1 row selected

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

1. The query contains a subselect of a UNION, where one of the columns is a literal, e.g. '99999999'.
2. The where clause contains an equality predicate, a GTR predicate, and an IN clause.

As a workaround, the query works if the IN clause is moved before the GTR predicate, as in the following example.

```
set flags 'strategy,detail';
! The following query should return 0 rows
!
select ps.ID, ps.kbn, ps.ymd
  from (select ps1.ID,
              ps1.kbn,
              '99999999'
        from ps ps1, pm pm
        where pm.id = ps1.ID
 union all
 select ps2.id,
        ps2.kbn,
        ps2.end_ymd
        from ps ps2, pm pm
        where pm.id = ps2.id)
      as ps (id, kbn, ymd)
 where ps.id = '021023307' and
        ps.kbn in ('1','2') and          <=== moved
        ps.ymd > '12345678' ;
```

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.9 Query with CONCATENATE in BETWEEN Clause Returns Wrong Results

Bug 1663038

The following query uses the CONCATENATE function in the BETWEEN clause. It should return 3 rows, but it returns only 1 row.

```
SQL> sh tab ORDER;
Information for table ORDER

Columns for table ORDER:
Column Name                Data Type          Domain
-----
ORDER_NO                   CHAR(4)
  Not Null constraint ORDER_NO_NOT_NULL
SHIP_DATE                  CHAR(8)
  Not Null constraint ORDER_NOT_NULL
SHIP_STAT                  CHAR(1)
  Not Null constraint ORDER_NOT_NULL
...etc...

Table constraints for ORDER:
ORDER_NOT_NULL
  Not Null constraint
  Column constraint for ORDER.SHIP_DATE
  Evaluated on COMMIT
Source:
  ORDER.SHIP_DATE NOT null
...etc...
```

```

SQL> sel order_no from customer;
ORDER_NO
1ED0
1j80
1a78
3 rows selected
SQL> sel order_no,ship_date,ship_stat from order;
ORDER_NO  SHIP_DATE  SHIP_STAT
1ED0      20010301   b
1a78      20010228   a
1j80      20010301   a
3 rows selected

```

```

set flags 'strategy,detail';
set flags 'max_stab';
select  a.order_no, a.ship_date, a.ship_stat
from    ORDER a, CUSTOMER b
where   a.order_no = b.order_no and
        ((a.SHIP_DATE || a.SHIP_STAT)
         BETWEEN '20010228a' '20010301d') ;

```

Tables:

```

0 = ORDER
1 = CUSTOMER

```

Cross block of 2 entries

Cross block entry 1

Conjunct:

```

(0.SHIP_DATE > SUBSTRING ('20010228a' FROM 0 FOR 8)) OR
((0.SHIP_DATE = SUBSTRING ('20010228a' FROM 0 FOR 8)) AND
(0.SHIP_STAT >= SUBSTRING ('20010228a' FROM 8)))

```

Conjunct:

```

((0.SHIP_DATE < SUBSTRING ('20010301d' FROM 0 FOR 8)) AND
NOT MISSING (0.SHIP_STAT)) OR
((0.SHIP_DATE = SUBSTRING ('20010301d' FROM 0 FOR 8)) AND
(0.SHIP_STAT <= SUBSTRING ('20010301d' FROM 8)))

```

Get Retrieval by index of relation 0:ORDER

Index name ORDER_UM01 [0:0]

Cross block entry 2

Index only retrieval of relation 1:CUSTOMER

Index name CUSTOMER_UM01 [1:1] Direct lookup

Key: 0.ORDER_NO = 1.ORDER_NO

```

A.ORDER_NO  A.SHIP_DATE  A.SHIP_STAT
1a78        20010228    a

```

1 row selected

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

1. The table columns contain NOT NULL constraints.
2. The query contains a BETWEEN clause with CONCATENATE function on two columns.

As a workaround, the query works if the column constraint ORDER_NOT_NULL is removed from the columns of table ORDER.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.10 ORDER BY Query With GROUP BY on Two Joined Derived Tables Returns Wrong Results

Bug 1694233

The following query with GROUP BY and ORDER BY clauses on two joined derived tables returns the results in the wrong order.

```
set flags 'strategy,detail';

select
  cast (a.name as char(5)) as name,
  a.datum
from (select name, datum,
  cast (count (*) as integer) as count_a
  from a
  group by name, datum) a
join
(select name, datum,
  cast (count (*) as integer) as count_b
  from b
  group by name, datum) b
on a.name = b.name
and a.datum = b.datum
group by a.name, b.name, a.datum, b.datum, count_a
order by name desc, a.datum asc
;

Tables:
  0 = A
  1 = B
Reduce: 0.NAME, 0.DATUM, 1.NAME, 1.DATUM, CAST (<mapped field> AS INT)
Sort: 0.NAME(a), 0.DATUM(a), 1.NAME(a), 1.DATUM(a), CAST (<mapped field> AS INT)
(a)
Cross block of 2 entries
Cross block entry 1
  Merge of 1 entries
    Merge block entry 1
      Aggregate: COUNT (*)
      Sort: 0.NAME(a), 0.DATUM(a)
      Get Retrieval sequentially of relation 0:A
Cross block entry 2
  Merge of 1 entries
    Merge block entry 1
      Aggregate: COUNT (*)
      Sort: 1.NAME(a), 1.DATUM(a)
      Conjoinct: (0.NAME = 1.NAME) AND (0.DATUM = 1.DATUM)
      Get Retrieval sequentially of relation 1:B
A.NAME  A.DATUM
AAAA    1-JAN-2000 00:00:00.00  <=== BBBB should be followed by AAAA
BBBB    1-JAN-2000 00:00:00.00
2 rows selected
```

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

1. The main query contains a GROUP BY clause on the columns of the two joined derived tables with GROUP BY.
2. One of the columns from the derived tables is cast as the same data type.
3. The ORDER BY clause references the cast column but using descending order.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.11 Left Outer Join Query With CONCATENATE Returns Wrong Results

Bug 1680135

The following left OJ query with CONCATENATE should return 1 row but instead returns 0 rows.

```
set flags 'strategy,detail';
SELECT ttt.entity_id,
       ttt.cpty_id,
       ttt.trade_count
FROM (SELECT tt.entity_id,
            tt.cpty_id,
            SUM (tt.trade_count) as trade_count
      FROM (SELECT df.entity_id,
                  df.cpty_id,
                  case
                    when df.deal_status = 'X' then 1 else 0
                  end as trade_count
            from deal_folder df) as tt
      GROUP BY tt.entity_id, tt.cpty_id) as ttt
LEFT OUTER JOIN
contact c ON (c.cpty_id = ttt.cpty_id)
WHERE
  ttt.trade_count <> 0
  and ttt.entity_id || ttt.cpty_id > '' ! <== this is causing problem
;
```

Tables:

0 = DEAL_FOLDER

1 = CONTACT

Conjunct: (<mapped field> <> 0) AND ((0.ENTITY_ID || 0.CPTY_ID) > '') <=(1)

Cross block of 2 entries (Left Outer Join)

Cross block entry 1

Conjunct: <mapped field> <> 0

Merge of 1 entries

Merge block entry 1

Aggregate: SUM (CASE (WHEN (0.DEAL_STATUS = 'X') THEN 1
ELSE 0))

Sort: 0.ENTITY_ID(a), 0.CPTY_ID(a)

Merge of 1 entries

Merge block entry 1

Conjunct: (0.ENTITY_ID || 0.CPTY_ID) > ''

Index only retrieval of relation 0:DEAL_FOLDER

Index name DEAL_FOLDER_MONITOR_IDX [0:0]

Cross block entry 2

Conjunct: (<mapped field> <> 0) AND ((0.ENTITY_ID || 0.CPTY_ID) > '') <=(2)

Conjunct: 1.CPTY_ID = 0.CPTY_ID

Index only retrieval of relation 1:CONTACT

Index name CONTACT_IDX [0:0]

0 rows selected

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

1. The main query is a left outer join between a derived table and a table.
2. The derived table contains a GROUP BY clause on the columns of another derived table with an aggregate function SUM as the output column.
3. The main query has a WHERE predicate containing the CONCATENATE function on two or more columns of the derived table.
4. The main query has another WHERE predicate which references the output column of the aggregate function from the derived table.

As a workaround, the query works if the table 1:CONTACT has some rows or the following CONCATENATE function is replaced by the following predicates:

```
ttt.entity_id || ttt.cpty_id > ''  
  
is replaced by  
  
ttt.entity_id > '' AND ttt.cpty_id > ''
```

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.12 Query With UNION in German Collating Sequence Returns Wrong Results

Bug 1684612

The following query with a UNION clause, in a database where the German Collating Sequence is used by default, returns wrong results (it should return some rows).

```
select d.datum, d.id, d.team  
from teamer d,  
     (select s.datum,s.id, s.team  
      from team_datum s  
      union all  
      select datum, id, team  
      from team_datum  
      ) as s  
where  
     d.datum=s.datum  
     ;  
Tables:  
  0 = teamer  
  1 = team_datum  
  2 = team_datum  
Conjunct: 0.datum = <mapped field>  
Match  
Outer loop  
  Sort: <mapped field>(a)  
  Merge of 1 entries  
    Merge block entry 1  
  Merge of 2 entries  
    Merge block entry 1  
    Get Retrieval sequentially of relation 1:team_datum  
    Merge block entry 2  
    Get Retrieval sequentially of relation 2:team_datum  
Inner loop  
  Temporary relation  
  Sort: <mapped field>(a)  
  Get Retrieval sequentially of relation 0:teamer  
0 rows selected
```

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

1. The query is a simple join between a table and a derived table of subselects unioned together.
2. The join predicate uses CHAR data type.
3. The Optimizer uses a match strategy to join them, where a comparison of the join keys requires the process of encoding the CHAR data type into German collating sequence.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.13 Query With OR Predicate on Aggregate Column Returns Wrong Results

Bugs 1708342 and 1721323

Query #1:

The following query with an OR predicate should return 1 row with T1.STATUS = 3 but returns an extra row with T1.STATUS = 5. This row does not satisfy the condition in the predicate "x.summe is null".

```
set flags 'max_stability';
set flags 'strategy,detail';
select
  t1.id,
  t1.status,
  t1.anzahl_stuecke,
  x.summe
from table1 t1,
  (select sum(anzahl_stuecke) as summe
   from table2 t2
   where t1.id = t2.id ) x
where
  t1.status = 3
  OR
  (t1.status = 5 and x.summe is null) ;
```

Tables:

```
0 = TABLE1
1 = TABLE2
```

Cross block of 2 entries

Cross block entry 1

Conjunct: (0.STATUS = 3) OR (0.STATUS = 5)

Get Retrieval by index of relation 0:TABLE1

Index name XPKTABLE1 [0:0]

Cross block entry 2

Merge of 1 entries

Merge block entry 1

Aggregate: SUM (1.ANZAHL_STUECKE)

Get Retrieval by index of relation 1:TABLE2

Index name XPKTABLE2 [1:1]

Keys: 0.ID = 1.ID

T1.ID	T1.STATUS	T1.ANZAHL_STUECKE	X.SUMME
1	3	10	NULL
2	5	10	10

2 rows selected

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

1. The main query joins a table and a derived table with a column of an aggregate function (e.g. SUM).
2. The WHERE clause contains an OR predicate, where one of the branches references the aggregated column.

As a workaround, the query works if the branches of the OR predicates are swapped, as in the following example.

```

select
  t1.id,
  t1.status,
  t1.anzahl_stuecke,
  x.summe
from table1 t1,
  (select sum(anzahl_stuecke) as summe
   from table2 t2
   where t1.id = t2.id ) x
where
  (t1.status = 5 and x.summe is null)
  OR
  t1.status = 3 ;
Tables:
  0 = TABLE1
  1 = TABLE2
Cross block of 2 entries
Cross block entry 1
  Get      Retrieval by index of relation 0:TABLE1
          Index name  XPKTABLE1 [0:0]
Cross block entry 2
  Conjunct: ((0.STATUS = 5) AND MISSING (var) OR (0.STATUS = 3)
Merge of 1 entries
  Merge block entry 1
  Aggregate: SUM (1.ANZAHL_STUECKE)
  Get      Retrieval by index of relation 1:TABLE2
          Index name  XPKTABLE2 [1:1]
          Keys: 0.ID = 1.ID
          T1.ID      T1.STATUS   T1.ANZAHL_STUECKE      X.SUMME
                1              3                   10              NULL
1 row selected

```

Query #2:

The following query with an OR predicate should return 0 rows.

```

set flags 'max_stability';
set flags 'strategy,detail';
select
  t1.id,
  t1.status,
  t1.anzahl_stuecke,
  x.summe
from table1 t1,
  (select
    sum(anzahl_stuecke) as summe,
    'hello' as Artikel
  from table2 t2
  where t1.id = t2.id ) x
where
  t1.id <> 5 and
  x.Artikel = 'hello should not be found' and
  ((t1.status =3) or
   (t1.status = 5 and (x.summe is NULL)))
);
Tables:
  0 = TABLE1
  1 = TABLE2
Cross block of 2 entries
Cross block entry 1
  Get      Retrieval by index of relation 0:TABLE1
          Index name  XPKTABLE1 [0:0]
          Bool: 0.ID <> 5
Cross block entry 2

```

```

Conjunct: (0.STATUS = 3) OR ((0.STATUS = 5) AND MISSING (var))
Merge of 1 entries
Merge block entry 1
Aggregate: SUM (1.ANZAHL_STUECKE)
Get      Retrieval by index of relation 1:TABLE2
Index name  XPKTABLE2 [1:1]
Keys: 0.ID = 1.ID
Bool: (1.ID <> 5) AND ('hello' = 'hello should not be found')
T1.ID      T1.STATUS      T1.ANZAHL_STUECKE      X.SUMME
   1             3             10             NULL
   2             5             10             NULL
2 rows selected

```

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

1. The main query joins a table and a derived table with the column of an aggregate function (e.g. SUM) and a column of a constant string.
2. The WHERE clause contains an OR predicate, where one of the branches references the aggregate column.
3. The WHERE clause contains additional AND predicates where one of them references the column of a constant string.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.14 Query With Equality Predicate Included in IN Clause Returns Wrong Results

Bug 1727181

The following query with an equality predicate included in the IN clause should find the row.

```

set flags 'strategy,detail';
sel employee_id
  from employees e, departments d
  where
    e.employee_id = d.manager_id and
    d.department_code in ('ADMN', 'ENG', 'MKTG') and
    d.department_code = 'ENG'
;
Tables:
  0 = EMPLOYEES
  1 = DEPARTMENTS
Cross block of 2 entries
Cross block entry 1
Conjunct: (1.DEPARTMENT_CODE = 'ADMN') OR (1.DEPARTMENT_CODE = 'MKTG')
Conjunct: 1.DEPARTMENT_CODE = 'ENG'
Index only retrieval of relation 1:DEPARTMENTS
Index name  DEPT_DEPTCODE_MGRID [1:1]
Keys: 1.DEPARTMENT_CODE = 'ENG'
Cross block entry 2
Index only retrieval of relation 0:EMPLOYEES
Index name  EMP_EMPID_STATUS_CODE [1:1]
Keys: 0.EMPLOYEE_ID = 1.MANAGER_ID
0 rows selected

```


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

1. The query joins two tables using a join predicate.
2. The query has an equality predicate which is also included in the IN clause.

As a workaround, the query works if the equality predicate is moved to the front of the IN clause, as in the following example.

```
set flags 'strategy,detail';
sel employee_id
  from employees e, departments d
  where
    e.employee_id = d.manager_id and
    d.department_code = 'ENG' and                               <== move to front
    d.department_code in ('ADMN', 'ENG', 'MKTG')
```

Tables:

```
0 = EMPLOYEES
1 = DEPARTMENTS
```

Cross block of 2 entries

Cross block entry 1

```
Conjunct: 1.DEPARTMENT_CODE = 'ENG'
Conjunct: (1.DEPARTMENT_CODE = 'ADMN') OR (1.DEPARTMENT_CODE = 'ENG') OR (
  1.DEPARTMENT_CODE = 'MKTG')
Index only retrieval of relation 1:DEPARTMENTS
Index name  DEPT_DEPTCODE_MGRID [1:1]
Keys: 1.DEPARTMENT_CODE = 'ENG'
```

Cross block entry 2

```
Conjunct: 1.DEPARTMENT_CODE = 'ENG'
Index only retrieval of relation 0:EMPLOYEES
Index name  EMP_EMPID_STATUS_CODE [1:1]
Keys: 0.EMPLOYEE_ID = 1.MANAGER_ID
```

E.EMPLOYEE_ID

00471

1 row selected

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.15 Match Strategy on Columns of Different Size, Using Collating Sequence, Returns Wrong Results

Bug 1684643

The following query using match strategy on columns of different size, using German collating sequence, should find the row.

```
select d.datum, d.abtlg, d.team, d.art
  from teamergebnis_kumul d,
     (select m.datum,m.abtlg, m.art, m.team
      from std_team_datum m, prod_kumul_datum v
      where m.datum=v.datum and
           m.abtlg=v.abtlg and
           m.team=v.produkt AND
           m.team='11.3512'
      group by m.datum, m.abtlg, m.art, m.team) AS
     s (datum, abtlg, art, team)
 where d.datum=s.datum and
       d.abtlg=s.abtlg and
       d.team=s.team and
```

```

d.art=s.art and
d.abtlg='465' and d.datum='20001031' and
d.team='11.3512';
Tables:
 0 = TEAMERGEBNIS_KUMUL
 1 = STD_TEAM_DATUM
 2 = PROD_KUMUL_DATUM
Cross block of 2 entries
Cross block entry 1
  Conjunct: 0.TEAM = '11.3512'
  Get      Retrieval by index of relation 0:TEAMERGEBNIS_KUMUL
           Index name  IDX_TEAMERGEBNIS_KUMUL_SORT [3:3]
           Keys: (0.TEAM = '11.3512') AND (0.DATUM = '20001031') AND (0.ABTLG =
                '465')
Cross block entry 2
  Conjunct: 0.ABTLG = 1.ABTLG
  Conjunct: 0.TEAM = 1.TEAM
  Conjunct: 0.ART = 1.ART
  Merge of 1 entries
  Merge block entry 1
  Reduce: 1.TEAM, 1.ABTLG, 1.DATUM, 1.ART
  Sort: 1.TEAM(a), 1.ABTLG(a), 1.DATUM(a), 1.ART(a)
  Conjunct: (1.DATUM = 2.DATUM) AND (1.ABTLG = 2.ABTLG) AND (1.TEAM =
            2.PRODUKT)
  Match
  Outer loop
  Sort: 1.TEAM(a), 1.ABTLG(a), 1.DATUM(a)
  Conjunct: 1.TEAM = '11.3512'
  Get      Retrieval by index of relation 1:STD_TEAM_DATUM
           Index name  IDX_STD_TEAM_DATUM_SORT [2:2]
           Keys: (0.DATUM = 1.DATUM) AND (1.ABTLG = '465')
  Inner loop
  Temporary relation
  Sort: 2.PRODUKT(a), 2.ABTLG(a), 2.DATUM(a)
  Conjunct: 2.PRODUKT = '11.3512'
  Get      Retrieval by index of relation 2:PROD_KUMUL_DATUM
           Index name  IDX_PROD_KUMUL_DATUM_SORT [2:2]
           Keys: (2.DATUM = 0.DATUM) AND (2.ABTLG = '465')
0 rows selected

```

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

1. The main query is a simple join between a table and a derived table of subselect subquery, joining two tables using 3 equality predicates.
2. The join predicate uses columns of CHAR data type but different column size.
3. The optimizer uses a match strategy to join them, where a comparison of the join keys requires the process of encoding the CHAR data type into German collating sequence.

As a workaround, the query works if the match strategy is changed to use cross by using an outline.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.16 Left Outer Join Query With CAST Function on USING Column Bugchecks

Bug 1802653

The following left outer join query with CAST function on USING column bugchecks.

```

select count(*) from
( select p.paketwert from
  ( select
    cast(packet as integer) ! <=== CAST causing bugcheck
    from
    serien inner join sujet using (sujet)
  ) as p (paketwert)
) as astpreis (paketwert)
left outer join
( select t.paketwert from
  ( select
    packet
    from
    serien inner join sujet using (sujet)
  ) as t (paketwert)
) as opt(paketwert)
USING (paketwert) ;

```

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

1. The main query is a left outer join of 2 nested derived tables.
2. The CAST function is placed on the column of USING clause.

There is no known workaround for this problem.

This problem has been corrected in Oracle Rdb Release 7.1.0.

2.5.17 Query Using Constant Values in OR Predicates Returns Wrong Results

Bug 1769447

The following query using constant values in OR predicates should return 3 rows.

```

set flags 'strategy,detail';

SELECT coll FROM
(SELECT
  t2.coll as coll,
  t2.col2 as col2,
  t2.col3 as col3
  from table1 t1, table2 t2
  where t1.coll_id = t2.coll_id
) as
vt (coll, col2, col3)
WHERE
  vt.col3 > 0 AND
  vt.col2 >= 0 AND
  ( vt.coll <= 3 OR 'hostvar' = 'foo' );

```

Tables:

0 = TABLE1

1 = TABLE2

Merge of 1 entries

Merge block entry 1

Conjunct: 0.coll_id = 1.coll_id

Match

Outer loop (zig-zag)

Index only retrieval of relation 0:TABLE1

Index name TABLE1_NDX [0:0]

Inner loop (zig-zag)

```

Conjunct: (1.col3 > 0) AND (1.col2 >= 0)
Get      Retrieval by index of relation 1:TABLE2
Index name TABLE2_NDX [0:0]
Bool: <error: common keyonly boolean no predicates>
COL1
  1
  2
  3
  4
  5
  6
6 rows selected

```

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

1. The query selects from a derived table of a subselect joining 2 tables.
2. The WHERE clause contains 2 AND predicates and 1 OR predicate.
3. The OR predicate contains a branch of constant predicates, such as "1 = 2".

As a workaround, the query works if the constant condition "'hostvar' = 'foo'" is omitted, as in the following example.

```

set flags 'strategy,detail';

SELECT coll from
  (SELECT
    t2.coll1 as coll1,
    t2.col2 as col2,
    t2.col3 as col3
    from table1 t1, table2 t2
    where t1.coll1_id = t2.coll1_id
  ) as
  vt (coll1, col2, col3)
WHERE
  vt.col3 > 0 AND
  vt.col2 >= 0 AND
  ( vt.col1 <= 3
!      OR 'hostvar' = 'foo'          <=== commented out
  );

```

Tables:

```

0 = TABLE1
1 = TABLE2

```

Merge of 1 entries

Merge block entry 1

Conjunct: 0.coll1_id = 1.coll1_id

Match

Outer loop (zig-zag)

Index only retrieval of relation 0:TABLE1

Index name TABLE1_NDX [0:0]

Inner loop (zig-zag)

Conjunct: (1.col3 > 0) AND (1.col2 >= 0) AND (1.coll1 <= 3)

Get Retrieval by index of relation 1:TABLE2

Index name TABLE2_NDX [0:0]

Bool: 1.coll1 <= 3

COL1

```

  1
  2
  3

```

3 rows selected

This problem has been corrected in Oracle Rdb Release 7.1.0.

Chapter 3

Enhancements

3.1 Enhancements Provided in Oracle Rdb Release 7.1.0.1

3.1.1 SQL Now Supports a Native ABS Function

In prior releases of Oracle Rdb, the ABS function was provided by the SQL_FUNCTIONS script. This function was a DOUBLE PRECISION function that allowed values of most data types to be processed.

However, there were some inconsistencies introduced when large BIGINT values were used as rounding errors were introduced since DOUBLE PRECISION supports about 16 digits accuracy compared to the 18 digits supported by BIGINT. In addition, the INTERVAL data type could not be used with the provided function.

With this release, a new conditional function, ABS, conforming to the SQL:1999 database language standard, is now available. The ABS function returns NULL if the passed value expression evaluates to NULL. The datatype of the result is the same as the passed value expression and supports scaled values of these data types: TINYINT, SMALLINT, INTEGER, BIGINT, REAL, FLOAT, DOUBLE PRECISION, INTERVAL, DECIMAL, NUMERIC and NUMBER.

The absolute value function (ABS) returns NULL if the value expression evaluates to NULL. If the value expression evaluates to a value less than zero then that value is negated so that a positive value is returned. Otherwise the value is returned unchanged. For instance, ABS (-1) will return the value 1.

ABS (a) is equivalent to the CASE expression:

```
case
  when a < 0 then - a
  else a
end
```

USAGE NOTES:

- The SQL_FUNCTIONS script still includes the ABS external function definition for those stored definitions (procedures, functions, triggers, views, etc.) or compiled applications that currently use it. However, new references to ABS will use the new builtin conditional expression.
- Applications wishing to continue to use the external function should use delimiters around the ABS function name, as in the following example.

```
SQL> set quoting rules 'SQL92';
SQL> select "ABS" (v) from T;
```

The delimited name will force the function definition to be used.

- Please refer to Appendix G, Oracle Rdb7 SQL Reference Manual, Volume 3 for more information on the SQL_FUNCTIONS script.

Example 1: This example uses the ABS function on an INTERVAL result of a date subtraction.

```
SQL> select
cont>     ABS ((birthday - current_date) year(3))
cont> from employees
cont> order by employee_id
cont> limit to 10 rows;
```

```

054
047
047
064
068
062
044
069
050
074
10 rows selected

```

Example 2: This shows a more complex use of ABS within a statistical function.

```

SQL> -- what is the average time in a job for each employee
SQL> -- exclude anyone on there first job
SQL> select
cont>     employee_id,
cont>     AVG (ABS (EXTRACT (MONTH FROM (job_start - job_end) month (4))))
cont>         as "Average Job" edit using '--,---,--9.99" years"
cont> from JOB_HISTORY
cont> where employee_id < '00200'
cont> group by employee_id
cont> having COUNT (*) > 1;
EMPLOYEE_ID      Average Job
00164            14.00 years
00165            22.67 years
00166            20.00 years
00167            14.50 years
00168            26.33 years
00169            22.67 years
...etc...
00197            26.33 years
00198            37.00 years
00199            35.00 years
30 rows selected
%RDB-I-ELIM_NULL, null value eliminated in set function

```

3.1.2 New DUMP Output Format for LogMiner

A new output format type of "DUMP" has been added to the RMU /UNLOAD /AFTER_JOURNAL command. This output format is intended solely as a debug and informational tool. For each column of a record, the first 200 bytes of data contents are formatted such that binary numeric fields are converted to text and text fields are displayed with periods (.) replacing non-printable characters. NULL columns are indicated with the character string "NULL". The actual data length is indicated for VARCHAR columns.

Example output with the /FORMAT=DUMP qualifier:

```

$ RMU /UNLOAD /AFTER_JOURNAL MYDB.RDB MYDB.AIJBCK /FORMAT=DUMP
  /TABLE=(NAME=ALL_DATATYPES_TBL, OUTPUT=SYS$OUTPUT:)
RDB$LM_ACTION      : M
RDB$LM_RELATION_NAME : ALL_DATATYPES_TBL
RDB$LM_RECORD_TYPE  : 25
RDB$LM_DATA_LEN     : 460
RDB$LM_NBV_LEN      : 66
RDB$LM_DBK          : 46:635:0
RDB$LM_START_TAD    : 21-JUL-2001 15:48:52.6512009
RDB$LM_COMMIT_TAD   : 21-JUL-2001 15:48:53.0586846
RDB$LM_TSN          : 160

```

```

RDB$LM_REC_VER      : 1
TINT                : -123
SINT                : -321
INTEGER             : -212
BINT                : NULL
DECIMAL             : -145
NUMERIC             : NULL
FLOAT               : -1.0000000000000000E+000
DOUBLE_PRECISION    : -2.0000000000000000E+000
CHAR1                : A
CHAR20              : ABCDEFGHIJKLMNOPQRST
VCHAR_COL           : (10) ABCDEFGHIJ

```

Note

The contents and format of the output when the /FORMAT=DUMP qualifier is specified may change in the future.

If needed, the record definition (.RRD) file may be used to determine the actual data type for each field of the table(s) being extracted.

3.1.3 Data and SPAM Prefetch Screens Added to RMU/SHOW STATISTICS

Two new screens have been added to the PIO statistics part of RMU/SHOW STATISTICS. These screens display prefetch statistics (APF and DAPF). In prior versions, the DAPF statistics were displayed on the "Fetch" screens. Those statistics were moved to the new prefetch screens. In addition, APF statistics are now displayed on the new screens as well. An example is provided below:

```

Node: NODE1 (1/1/1)      Oracle Rdb V7.0-62 Perf. Monitor  6-AUG-2001 10:28:10.65
Rate: 3.00 Seconds      PIO Statistics--Data Prefetches  Elapsed: 00:58:17.86
Page: 1 of 1            DEV:[DIR]DB.RDB                Mode: Online

```

```

-----
statistic.....      rate.per.second..... total..... average.....
name.....           max..... cur..... avg..... count..... per.trans....

APF start:success    0          0          0.4         872         1.0
      :failure        0          0          0.0         101         0.1

APF I/O: utilized    0          0          0.4         872         1.0
      : wasted        0          0          0.0          0          0.0

DAPF start:success   0          0          0.0         74          0.0
      :failure        0          0          0.0         62          0.0

DAPF I/O: utilized   0          0          0.0         18          0.0
      : wasted        0          0          0.0         56          0.0

```

The information on these screens may be used to determine the effectiveness of the APF and DAPF features. The individual rows may be interpreted as follows:

- The "APF start:success" statistics shows how many times Oracle Rdb successfully initiated an I/O to prefetch a buffer.
- The "APF start:failure" statistics shows how many times Oracle Rdb attempted to initiate a prefetch but was unable to obtain the necessary buffer lock to proceed.
- The "APF I/O: utilized" statistics shows how many times Oracle Rdb actually used a buffer that was

prefetched.

- The "APF I/O: wasted" statistics shows how many times Oracle Rdb prefetched a buffer but never actually used it.

3.1.4 RMU/SHOW STATISTICS Stall Log Lock Information Optional

Bug 1704232

A new optional keyword "[NO]LOG_STALL_LOCK" has been added to the "/OPTIONS" qualifier of the RMU/SHOW STATISTICS command. When using the /STALL_LOG qualifier to write stall messages to a log file, you can now specify /OPTIONS=NOLOG_STALL_LOCK to prevent lock information from being written to the log file.

The following example shows stall log information first with the lock information and then without the lock information:

```
$ RMU /SHOW STATISTICS /NOINTERACTIVE /STALL_LOG=SYS$OUTPUT: -
  DUA0:[DB]MFP.RDB
Oracle Rdb X7.1-00 Performance Monitor Stall Log
Database DPA500:[RDB_RANDOM.RDB_RANDOM_TST_247]RNDDB.RDB;1
Stall Log created 4-SEP-2001 11:27:03.96
11:27:03.96 0002B8A1:1 11:27:03.67 waiting for record 118:2:2 (PR)
  State... Process.ID Process.name... Lock.ID. Rq Gr Queue "record 118:2:2"
  Blocker: 000220A7 RND_TST_24716 0F019E52 EX Grant
  Waiting: 0002B8A1 RND_TST_24715 4500C313 PR Wait
11:27:03.96 0002B8A8:1 11:27:02.32 waiting for record 101:3:0 (EX)
  State... Process.ID Process.name... Lock.ID. Rq Gr Queue "record 101:3:0"
  Blocker: 000220AD RND_TST_24710 0B00176A PR Grant
  Blocker: 000220A7 RND_TST_24716 52018A3F PR Grant
  Waiting: 0002B8A8 RND_TST_2474 3C00B5AF EX PR Cnvrt
11:27:03.96 0002B89C:1 11:27:00.15 waiting for record 114:4:1 (PR)
  State... Process.ID Process.name... Lock.ID. Rq Gr Queue "record 114:4:1"
  Blocker: 000220A7 RND_TST_24716 180033CC EX Grant
  Waiting: 0002B89C RND_TST_2479 110066BA PR Wait

$ RMU /SHOW STATISTICS /NOINTERACTIVE /STALL_LOG=SYS$OUTPUT: -
  DUA0:[DB]MFP.RDB /OPTIONS=NOLOG_STALL_LOCK
Oracle Rdb X7.1-00 Performance Monitor Stall Log
Database DPA500:[RDB_RANDOM.RDB_RANDOM_TST_247]RNDDB.RDB;1
Stall Log created 4-SEP-2001 11:28:34.68
11:28:34.69 0002B8B8:1 11:28:33.69 waiting for logical area 146 (PR)
11:28:34.69 0002B8A8:1 11:28:32.76 waiting for record 114:4:2 (PR)
11:28:34.69 0002B8B3:1 11:28:33.06 waiting for record 114:4:2 (PR)
11:28:34.69 0002B8B0:1 11:28:31.96 waiting for record 111:7:7 (EX)
```

3.1.5 New Option for the GET DIAGNOSTICS Statement

For Oracle Rdb Release 7.1.0.1, a new option has been added to the GET DIAGNOSTICS statement: IMAGE_NAME.

This keyword requests that the activating image name be returned to the caller. The image name includes the node name from which the attach was started. This might be a node different than that on which the Rdb server is running.

The data is returned to the caller as a VARCHAR (255) value and should be assigned to either a VARCHAR or CHAR data type that supports the ASCII character set.

The following example uses a SQL procedure to fetch the image name for the currently running application (in this case interactive SQL).

```
SQL> set flags 'trace';
SQL> begin
cont> declare :i varchar(512);
cont> get diagnostics :i = image_name;
cont> trace char_length (:i);
cont> trace '' || :i || '';
cont> end;
~Xt: 57
~Xt: "MYNODE:::111$DUA618:[SYS0.SYSCOMMON.][SYSEXE]SQL$71.EXE;1"
```

3.1.6 Alternate Outline Ids

If outlines have not been disabled, Oracle Rdb will search for an appropriate outline for the query it is optimizing, thus allowing some user control of the strategy used for execution of a query.

The OPTIMIZE USING clause may be used to tell the optimizer which outline to use for compilation. If no OPTIMIZE USING clause is present, Rdb uses the query to generate an identifier which it will use to try to locate an appropriate outline.

In many situations, such as when using third party software, it is not possible for the user to provide an outline name for the query and thus the only alternative Rdb had was to try to locate an outline with a matching identifier.

As the identifier is a hashed value that depends on the query structure, small changes in the query, such as different literal values, can change the identifier produced as in the following example.

```
SQL> set flags 'outline';
SQL> select * from employees where employee_id = '1';
-- Rdb Generated Outline : 19-SEP-2001 13:52
create outline QO_8797A75D6D03F6BD_00000000
id '8797A75D6D03F6BDD211A092CE6F3A2C'
mode 0
as (
  query (
-- For loop
    subquery (
      EMPLOYEES 0      access path index      EMP_EMPLOYEE_ID
    )
  )
)
compliance optional      ;
0 rows selected
SQL> select * from employees where employee_id = '9999';
-- Rdb Generated Outline : 19-SEP-2001 13:52
create outline QO_C9F12D27AC5D3163_00000000
id 'C9F12D27AC5D3163907A4329FDC8170A'
mode 0
as (
  query (
-- For loop
    subquery (
      EMPLOYEES 0      access path index      EMP_EMPLOYEE_ID
    )
  )
)
```

```

    )
  )
  compliance optional      ;

```

In this example, the two queries are optimized the same but the differing outline identifiers means that two different outlines would have to be created to control each query.

Oracle Rdb has now been enhanced to allow the optional creation of alternate outline identifiers. In this release, the optimizer discards literal values when producing the identifiers.

A new SET FLAGS attribute has been introduced to allow the control of these alternate identifiers, using either the SQL SET FLAGS statement or the RDMS\$SET_FLAGS logical name.

```
ALTERNATE_OUTLINE_ID(LITERALS)
```

This attribute is not case sensitive and may be abbreviated to:

```
ALT(LIT)
```

The following example uses SET FLAGS to enable alternate query identifiers:

```

SQL> set flags 'alt(LIT), outline';
SQL> select * from employees where employee_id = '1';
-- Rdb Generated Outline : 19-SEP-2001 13:52
create outline QO_847AD7287E247D37_00000000
id '847AD7287E247D37E8E4CC8221FFC12E'
mode 0
as (
  query (
-- For loop
    subquery (
      EMPLOYEES 0      access path index      EMP_EMPLOYEE_ID
    )
  )
)
compliance optional      ;
0 rows selected
SQL> select * from employees where employee_id = '9999';
-- Rdb Generated Outline : 19-SEP-2001 13:52
create outline QO_847AD7287E247D37_00000000
id '847AD7287E247D37E8E4CC8221FFC12E'
mode 0
as (
  query (
-- For loop
    subquery (
      EMPLOYEES 0      access path index      EMP_EMPLOYEE_ID
    )
  )
)
compliance optional      ;
0 rows selected

```

Note that now the two outlines have the same identifier and the user may now store this more generic outline to be used by any similar query where only the literal values differ. For example:

```

SQL> set flags 'alt(lit)';
SQL> create outline o1 from (select * from employees where employee_id = '1');
SQL> set flags 'strat';
SQL> select * from employees where employee_id = '1';
~S: Outline "O1" used

```

```

Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL> select * from employees where employee_id = 'AAAAAA';
~S: Outline "O1" used
Conjunct      Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected

```

Any outline stored for a query without the `ALTERNATE_OUTLINE_ID` flag being set will be created using the full query as in previous versions and will take precedence over any generic outline. For example:

```

SQL> set noflags;
SQL> create outline o1 from (select * from employees where employee_id = '1');
SQL> set flags 'strat';
SQL> select * from employees where employee_id = '1';
~S: Outline "O1" used
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL> select * from employees where employee_id = '9999';
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL> set noflags;
SQL> set flags 'alternate(lit),nooutline';
SQL> create outline o2 from (select * from employees where employee_id = '1');
SQL>
SQL> set flags 'strat';
SQL> select * from employees where employee_id = '1';
~S: Outline "O1" used
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL> select * from employees where employee_id = '9999';
~S: Outline "O2" used
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL>
SQL> set flags 'noalt';
SQL> select * from employees where employee_id = '1';
~S: Outline "O1" used
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL> select * from employees where employee_id = '9999';
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL> drop outline o1;
SQL> set flags 'alt(literals)';
SQL> select * from employees where employee_id = '1';
~S: Outline "O2" used
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected
SQL> select * from employees where employee_id = '9999';
~S: Outline "O2" used
Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [1:1]      Direct lookup
0 rows selected

```

As shown in the previous example, Oracle Rdb will try to locate an outline using the more generic identifier only if the ALTERNATE_OUTLINE_ID flag has been set.

The ALTERNATE_OUTLINE_ID flag is not set by default and must be explicitly set using either SET FLAGS or the RDMS\$SET_FLAGS logical.

This feature is available in Oracle Rdb Release 7.1.0.1.

3.1.7 Field Widths Wider on Row Cache Overview Display

On the "Row Cache Overview" display, the width of the "Searches" column has been increased from 9 to 10 characters to allow a display of up to 4294967295 (after this value, the 32-bit counter wraps back to zero). In addition, the width of the cache name column is tied to the screen width. If the screen is set to be wide enough (over 90 columns), the full cache name will be displayed; normally, only the first 24 characters of the name are displayed.

Additionally, the comparison used when sorting by values on the "Row Cache Overview" display has been modified to be unsigned (rather than signed). This prevents some cases of very large values being sorted in an incorrect order.

3.1.8 FOR Counted Loop Enhancements

In Oracle Rdb Release 7.1, the FOR counted loop was added to SQL. This type of loop increments a declared variable from an initial value to a final value. In the prior release of Rdb, the data type of the variable had to be a numeric data type (TINYINT, SMALLINT, INTEGER, BIGINT, REAL, FLOAT, DOUBLE PRECISION, NUMBER, NUMERIC, or DECIMAL).

The following enhancements have been made for this release:

- The following data types are now also legal for this type of FOR loop.

```
INTERVAL YEAR
INTERVAL MONTH
INTERVAL DAY
INTERVAL HOUR
INTERVAL MINUTE
INTERVAL SECOND
```

If INTERVAL is used, then the initial and final values must be of the same type (i.e. the expressions must have the same data type as the loop variable).

- The data type rules for the initial and final values have been relaxed when the loop variable is numeric. These value expressions can be any compatible numeric data type. For instance, floating point or scaled numeric values can now be used.
- A new optional STEP clause has been added to control the size of the increment between loop iterations. The step size is specified using a numeric value expression.

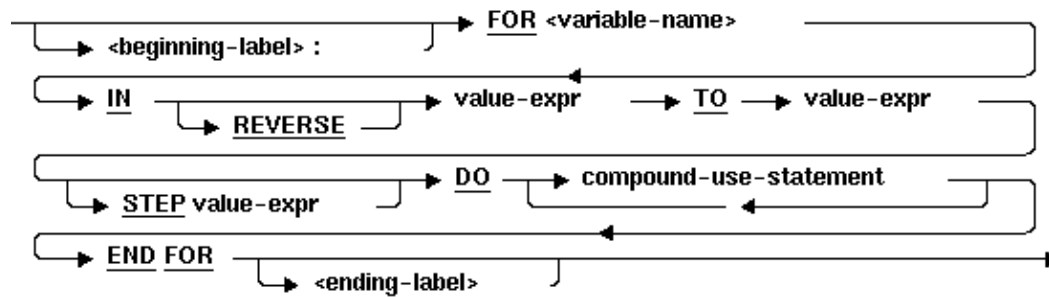
```
SQL> begin
cont> declare :i integer;
cont> for :i in 1 to 20 step 5
cont> do
cont>     trace :i;
cont> end for;
cont> end;
~Xt: 1
~Xt: 6
```

~Xt: 11
~Xt: 16

NOTE: Even if the loop control variable is an INTERVAL type, the STEP must be numeric type. In addition, the value must be greater than zero: use the REVERSE keyword to decrement the loop control variable.

FORMAT

counted-for-statement =



USAGE NOTES

- The FOR loop uses the keyword TO as a separator between the initial and final value expressions. This same keyword is used to separate the field names in an interval qualifier. Therefore, there is an ambiguity possible when an apparently well-formed expression is used.

```
SQL> begin
cont> declare :i interval year;
cont> for :i in interval'1' year to interval'4'year
for :i in interval'1' year to interval'4'year
      ^
%SQL-W-LOOK_FOR_STT, Syntax error, looking for:
%SQL-W-LOOK_FOR_CON,          MONTH,
%SQL-F-LOOK_FOR_FIN,    found INTERVAL instead
```

This occurs because the TO separator is interpreted as part of the INTERVAL literal or expression. Programmers must enclose the initial expression in parentheses to avoid this ambiguity if it ends with an interval qualifier.

- The STEP value expression is evaluated before the loop variable is assigned a value. The value must be greater than zero and never NULL. If these constraints are violated, a runtime error is reported as shown in this simple example.

```
SQL> begin
cont> declare :l, :s integer;
cont>
cont> -- set the step size
cont> set :s = 0;
cont>
cont> for :l in reverse 1 to 10 step :s
cont> do
cont>     trace :l;
cont> end for;
cont> end;
%RDB-E-NOT_VALID, validation on field STEP caused operation to fail
SQL>
```

Example 1: This example shows an INTERVAL type as the loop variable.

```
SQL> begin
cont> declare :i interval year;
cont> for :i in (interval'1' year) to (interval'4'year)
cont> do
cont>     trace :i;
cont> end for;
cont> end;
~Xt: 01
~Xt: 02
~Xt: 03
~Xt: 04
```

Example 2: This example uses a complex expression as the STEP expression.

```
SQL> begin
cont> declare :i interval year;
cont> declare :k interval year = interval'18'year;
cont> declare :j integer = 2;
cont>
cont> for :i in (interval'1' year) to :k/2 step :j*2
cont> do
cont>     trace :i;
cont> end for;
cont> end;
~Xt: 01
~Xt: 05
~Xt: 09
```


3.1.9 Enhancements to SET DISPLAY Statement for Interactive SQL

This release of Oracle Rdb, 7.1.0.1, includes the following enhancements to the SET DISPLAY statement.

- A new NULL STRING clause to change the way NULL values are displayed by interactive SQL.
- A new DEFAULT NULL STRING clause to revert to using the text 'NULL'.
- A new [NO] COMMENT clause to disable or enable the display of comment text by other SHOW commands (e.g. SHOW TABLE).

FORMAT

set-display =



USAGE NOTES

- The width of the displayed column is calculated using the maximum of the length of the column name, the length of the QUERY HEADER, the length of the NULL string and the size of the formatted data.
- The statement SET DISPLAY DEFAULT NULL STRING is equivalent to SET DISPLAY NULL STRING 'NULL'.
- The SET NULL statement has been added for compatibility with Oracle SQL*Plus. SET NULL is a synonym for SET DISPLAY NULL STRING "", and SET NULL 'literal' is equivalent to SET DISPLAY NULL 'literal'.
- SET DISPLAY NULL STRING accepts a string literal, or a declared local variable.
- SHOW DISPLAY now displays the current NULL string.

```
SQL> show display
Output of the query header is enabled
Output of the row counter is enabled
Output using edit strings is enabled
Page length is set to 30 lines
Line length is set to 80 bytes
Display NULL values using "NULL"
```

- The GET ENVIRONMENT statement now includes the NULL_STRING keyword that can be used to save the currently defined text.

Example 1: Replace the NULL values with text to make the output easier to read.

```
SQL> select job_start, job_end,
cont>          (select department_name
cont>          from departments d
cont>          where d.department_code = jh.department_code)
cont> from job_history jh
cont> where employee_id = '00164';
JOB_START      JOB_END
21-Sep-1981    NULL          Board Manufacturing North
5-Jul-1980     20-Sep-1981  Cabinet & Frame Manufacturing
2 rows selected
SQL> set display null string '(still employeeed)'
SQL> select job_start, job_end,
cont>          (select department_name
cont>          from departments d
cont>          where d.department_code = jh.department_code)
cont> from job_history jh
cont> where employee_id = '00164';
JOB_START      JOB_END
21-Sep-1981    (still employeeed)  Board Manufacturing North
5-Jul-1980     20-Sep-1981        Cabinet & Frame Manufacturing
2 rows selected
```

Example 2: Disable the comment display to make the output of SHOW easier to read.

```
SQL> show domain id_dom
ID_DOM          CHAR(5)
Comment:        standard definition of employee id
SQL> set display no comment;
SQL> show domain id_dom
ID_DOM          CHAR(5)
SQL>
```

Example 3: Save the current NULL string using GET ENVIRONMENT and restore after executing a query.

```
SQL> declare :ns varchar(100);
SQL> get environment (session) :ns = NULL_STRING;
```



```

SQL> set null;
SQL> select job_start, job_end,
cont>         (select department_name
cont>         from departments d
cont>         where d.department_code = jh.department_code)
cont> from job_history jh
cont> where employee_id = '00164';
JOB_START      JOB_END
21-Sep-1981
5-Jul-1980     20-Sep-1981   Board Manufacturing North
Cabinet & Frame Manufacturing
2 rows selected
SQL> set display null string :ns;
SQL> select job_start, job_end,
cont>         (select department_name
cont>         from departments d
cont>         where d.department_code = jh.department_code)
cont> from job_history jh
cont> where employee_id = '00164';
JOB_START      JOB_END
21-Sep-1981     NULL           Board Manufacturing North
5-Jul-1980     20-Sep-1981   Cabinet & Frame Manufacturing
2 rows selected

```

3.1.10 New BITSTRING Built In Function

Rdb now supports a BITSTRING function that can be used to extract selected bits from a binary data value. This functionality is primarily intended to query the bit values stored in the RDB\$FLAGS columns in the Rdb system table but can also be used for user data.

BITSTRING accepts numeric and date/time values and processes them as bit arrays. The first (least significant) bit is numbered 1. The most significant bit depends on the data type.

- TINYINT has 8 bits
- SMALLINT has 16 bits
- INTEGER has 32 bits
- BIGINT, DATE, TIME, TIMESTAMP and INTERVAL types have 64 bits

FORMAT

```

BITSTRING  → (  → value-expression
┌──────────┴──────────┐
└── FROM numeric-expression ───┘
└── FOR numeric-expression ───┘
)  →

```

USAGE NOTES

- The numeric expression after the FOR and FROM keywords must be an unscaled numeric value.
- If the numeric expression of the FOR clause is less than or equal to zero then it will be assumed equal to 1.
- If the FOR clause is omitted, it will default to a value that includes all remaining bits of the source value.
- If the FOR clause specifies a larger value than the number of bits remaining in the source then it will only return the remaining bits.

Example: Bit 1 in the RDB\$FLAGS column of RDB\$RELATIONS indicates that the table is a view. This example uses this query to fetch the names of all user defined views in the PERSONNEL database.

```

SQL> select rdb$relation_name
cont> from rdb$relations
cont> where rdb$system_flag = 0 and
cont>      bitstring (rdb$flags from 1 for 1) = 1;
RDB$RELATION_NAME
CURRENT_JOB
CURRENT_SALARY
CURRENT_INFO
3 rows selected
SQL>

```

3.1.11 New SET PAGE LENGTH Command for Interactive SQL

SQL now includes a SET PAGE LENGTH statement to size the page. Currently this is only used by the pagination control in the SQL HELP command.

FORMAT

USAGE NOTES

- The integer value must be a value between 10 and 32767.
- SET PAGE LENGTH can be used to effectively disable the paging performed by help by setting the length to a high value such as 32000.
- The page length is automatically set upon entry to interactive SQL and is based on the OpenVMS terminal setting for this session.
- The SHOW DISPLAY command can be used to view the currently defined page length.

This example uses the SET PAGE LENGTH command to change the pagination length of HELP.

```

SQL> set page length 40;
SQL> show display
Output of the query header is enabled
Output of the row counter is enabled
Output using edit strings is enabled
Page length is set to 40 lines
Line length is set to 80 bytes
Display NULL values using "NULL"

```

3.1.12 New ALTER CONSTRAINT Statement

Oracle Rdb Release 7.1 includes an ALTER CONSTRAINT statement.


FORMAT

```

ALTER CONSTRAINT <constraint-name>
    COMMENT IS '<text-literal>'
    constraint-attributes

```

constraint-attributes =



Note: constraint-attributes are described in the Oracle Rdb New and Changed Features Manual.

USAGE NOTES

- If a constraint attribute is changed, it will only be effective for future references to the table containing that constraint. That is, if a constraint is already active then it will use the previously defined attributes.
- The constraint name can be prefixed with an alias name as in the following example.

```
SQL> alter constraint db1.ALL_UNIQUE
cont>     deferrable initially deferred;
```

This example shows how ALTER CONSTRAINT can be used to change the constraint attributes and add a comment to a constraint.

```
SQL> set dialect 'sql99';
SQL> attach 'file db$:mf_personnel';
SQL>
SQL> create table PERSON
cont>     (last_name char(20)
cont>         constraint MUST_HAVE_LAST_NAME
cont>         not null
cont>         deferrable,
cont>     first_name char(20),
cont>     birthday date
cont>         constraint MUST_BE_IN_PAST
cont>         check (birthday < current_date)
cont>         not deferrable,
cont>     constraint ALL_UNIQUE
cont>         unique (last_name, first_name, birthday)
cont>         deferrable initially immediate
cont> );
SQL>
SQL> show table (constraint) PERSON
Information for table PERSON
```

Table constraints for PERSON:

ALL_UNIQUE

Unique constraint

Null values are considered distinct

Table constraint for PERSON

Evaluated on each VERB

Source:

UNIQUE (last_name, first_name, birthday)

MUST_BE_IN_PAST

Check constraint

Column constraint for PERSON.BIRTHDAY
Evaluated on UPDATE, NOT DEFERRABLE
Source:
CHECK (birthday < current_date)

MUST_HAVE_LAST_NAME
Not Null constraint
Column constraint for PERSON.LAST_NAME
Evaluated on COMMIT
Source:
PERSON.LAST_NAME NOT null

Constraints referencing table PERSON:
No constraints found

```
SQL>
SQL> alter constraint ALL_UNIQUE
cont>     deferrable initially deferred;
SQL>
SQL> alter constraint MUST_HAVE_LAST_NAME
cont>     comment is 'We must assume all persons have a name'
cont>     not deferrable;
SQL>
SQL> alter constraint MUST_BE_IN_PAST
cont>     deferrable initially immediate;
SQL>
SQL> show table (constraint) PERSON
Information for table PERSON
```

Table constraints for PERSON:
ALL_UNIQUE
Unique constraint
Null values are considered distinct
Table constraint for PERSON
Evaluated on COMMIT
Source:
UNIQUE (last_name, first_name, birthday)

MUST_BE_IN_PAST
Check constraint
Column constraint for PERSON.BIRTHDAY
Evaluated on each VERB
Source:
CHECK (birthday < current_date)

MUST_HAVE_LAST_NAME
Not Null constraint
Column constraint for PERSON.LAST_NAME
Evaluated on UPDATE, NOT DEFERRABLE
Comment: We must assume all persons have a name
Source:
PERSON.LAST_NAME NOT null

Constraints referencing table PERSON:
No constraints found


```
SQL>
SQL> commit;
```

3.1.13 DECLARE Variable Now Supports CHECK Constraint


Variables declared within a compound statement (BEGIN...END) can now include a CHECK constraint to prevent out of range assignments to variables.

FORMAT

variable-declaration =



constraint-clause =



USAGE NOTES

- The constraint-clause is applied to all variables listed in DECLARE. The keyword VALUE can be used as a placeholder for the variable name with SQL correctly applying the constraint to all variables.
- Only the NOT DEFERRABLE and INITIALLY IMMEDIATE syntax is supported for variable constraints. This is also the default if no constraint-attributes are specified.
- A runtime error is signaled if the constraint is violated. This error will include the name of the variable.
- When a DEFAULT is not used in the declare statement, the contents of the variable are undefined. Therefore, any variable that uses a CHECK constraint must also provide a DEFAULT clause to ensure that the variable's value is consistent.
- Currently module global variables do not support constraints. This is planned for a future release of Oracle Rdb.

The following example shows the use of a CHECK constraint to prevent illegal values being assigned to control variables for a REPEAT loop. The singleton SELECT will actually return zero to the local variable P which will cause a variable validation to fail.

```
SQL> begin
cont> declare :v integer = 0 check (value is not null);
cont> declare :p integer = 1 check (value is not null and value <> 0);
cont>
cont> repeat
cont>   select count(*) into :p
cont>   from employees
cont>   where employee_id = '00000';
cont>   set :v = :v + :p;
cont> until :v > 1000
cont> end repeat;
cont> end;
%RDB-E-NOT_VALID, validation on field P caused operation to fail
```

3.1.14 RMU/SHOW STATISTICS Active User Stall Messages Sorted by Process ID

The RMU/SHOW STATISTICS "Active User Stall Messages" display now includes the ability to sort the list of database users by process ID (OpenVMS PID). The Config option on the horizontal menu at the bottom of the screen can be used to control how the information is to be sorted. By default, the display is unsorted.

3.1.15 RMU /REPAIR /INITIALIZE ONLY_LAREA_TYPE Keyword

This note was inadvertently left out of the Oracle Rdb Release 7.1.0 Release Notes.

A new ONLY_LAREA_TYPE keyword has been added to the RMU /REPAIR /INITIALIZE qualifier. This keyword, along with the /NOSPAM and /NOABM qualifiers, allows only the logical area "type" field to be updated in the AIP (area inventory pages). Avoiding SPAM page updates significantly improves performance of this operation.

The RMU /UNLOAD /AFTER_JOURNAL and RMU /SHOW STATISTICS commands use the on-disk area inventory pages (AIPs) to determine the appropriate "type" of each logical area. However, this logical area 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 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.
- BTREE – 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 which is used to identify hash buckets. Users cannot explicitly create these types of logical areas. 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 (segmented string; list of byte varying) 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.

The ONLY_LAREA_TYPE keyword can be specified along with the /NOSPAM and /NOABM qualifiers to cause *only* the logical area type field to be updated in the area inventory pages. All other actions specified in the options file are ignored when ONLY_LAREA_TYPE is specified. By updating only the logical area type in the AIP entries and not the SPAM pages, the RMU /REPAIR operation can be considerably faster.

3.1.16 RMU/SHOW STATISTICS Cluster Data Collection Performance Enhancement

The RMU /SHOW STATISTICS Utility has been enhanced to perform "asynchronous" data gathering when statistics are being displayed cluster-wide. Previously, a request for statistics was sent to the remote statistics server and then the response was received synchronously. This was repeated for each node being monitored at each data refresh cycle.

Now, the requests for information are sent to all nodes at once and then the replies are read as they become available. This reduces some of the the delay associated with gathering statistics from multiple nodes in a cluster.

3.1.17 RMU Extract has Enhanced Extract of Conditional Expressions

This release of Oracle Rdb now includes support for the new ABS function by RMU Extract. RMU Extract decodes case expressions into ABS (absolute value) functions.

ABS (a) is equivalent to:

```
CASE
  WHEN a < 0 THEN -a
  ELSE a
END
```

In addition, similar forms of CASE expressions are also converted to ABS.

```
CASE
  WHEN a <= 0 THEN -a
  ELSE a
END
```

and

```
CASE
  WHEN a > 0 THEN a
  ELSE -a
END
```

and

```
CASE
  WHEN a >= 0 THEN a
  ELSE -a
END
```

It is possible that RMU Extract will change existing CASE expressions into this more compact syntax, even if it was not originally coded as an ABS function call.

Chapter 4

Documentation Corrections, Additions and Changes

This chapter provides corrections for documentation errors and omissions.

4.1 Documentation Corrections

4.1.1 DROP INDEX Now an Online Table Operation

The example for the following note was in error in the original documentation.

DROP INDEX can now be used when other users are processing the table on which the index is defined. This requires that the index has previously been disabled with the ALTER INDEX ... MAINTENANCE IS DISABLED statement.

Once maintenance is disabled for an index, it is no longer used by queries on the table. For example, it is not used for retrieval and it is not updated by INSERT, DELETE or UPDATE statements. Therefore, with this release, Rdb has relaxed the requirement of EXCLUSIVE table access for DROP INDEX.

Oracle recommends that the DROP INDEX statement immediately be followed by a COMMIT statement so that all locks on the system metadata be released. Otherwise, access to this and other tables may be stalled waiting for rows locked in the tables RDB\$INDICES, RDB\$INDEX_SEGMENTS, RDB\$STORAGE_MAPS, and RDB\$STORAGE_MAP_AREAS.

This change benefits very large databases (VLDB) which have the need to drop indices stored in MIXED format storage areas on large cardinality tables. These indices may take several hours to erase, which in previous versions required taking the table offline from normal processing until the DROP INDEX completed.

Note that indices stored in UNIFORM format storage areas do not take long to DROP due to optimizations which can be made for UNIFORM areas.

```
-- Disable the index maintenance. This requires exclusive access to the
-- table, but takes a very short time. This should be done during normal
-- offline maintenance
--
set transaction read write;
alter index TRANSACTION_POSTING_INDEX
    maintenance is disabled;
commit;

-- Once disabled the index can be dropped at any time
--
set transaction read write;
drop index TRANSACTION_POSTING_INDEX;

commit;
```

Please note that DROP INDEX will write before image data to the snapshot files (.SNP) if the transaction is started in a mode such as SHARED or PROTECTED. Snapshots can be disabled on the database to avoid the excessive snapshot file I/O during concurrent DROP INDEX operations. Naturally, this may not be possible under normal workloads.

4.2 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

4.3 Online Document Format and Ordering Information

For release 7.1, we are providing documentation in Adobe Acrobat format and HTML format. All of the documentation is available in one or both of these formats on the Oracle Rdb Documentation CD-ROM, part number A90838-01.

In North America, printed documentation is available for sale in the Oracle Store at

<http://oraclestore.oracle.com/>

Customers in Europe, the Middle East, and Africa (EMEA) can purchase documentation from

<http://www.oraclebookshop.com/>

Other customers can contact their Oracle representative to purchase printed documentation.

4.3.1 Documentation in Adobe Acrobat Format

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 documentation available in Adobe Acrobat format is on the Oracle Rdb Documentation CD-ROM. Click on the welcome.pdf file in the top level directory to begin browsing the Oracle Rdb and related products documentation.

Also included with this set of PDF manuals is a master index built using Adobe Acrobat Catalog. You will need Adobe Acrobat Reader with Search capability to utilize this master index.

4.3.2 Documentation in HTML format

The documentation available in HTML format is on the Oracle Rdb Documentation CD-ROM. Click on the introductory page, welcome.htm, in the top level directory to begin viewing the Oracle Rdb and related product documentation.

4.4 Documentation for This Release

The following documents have been updated for this release but are not included on the Oracle Rdb Documentation CD-ROM:

- Oracle Rdb Release Notes
- Before You Install Oracle Rdb letter (cover_letter.txt)
- Oracle Rdb Installation and Configuration Guide (A90407-01)
- Oracle Rdb Oracle SQL/Services Release 7.1.5 Release Notes
- Oracle Rdb Oracle SQL/Services Release 7.1.5 Installation Guide (A90406-01)

These documents can be found in the Documentation directory on the Oracle Rdb Server software CD-ROM. The release notes can also be found in SYS\$HELP after installation.

The following table lists the manuals available on the Oracle Rdb Documentation CD-ROM (part number A90838-01), the part number for each manual, if the manual is available in PDF format, and if the manual is available in HTML format. Select either the welcome.htm or welcome.pdf file in the top level directory to begin viewing the Oracle Rdb and related products documentation.

Table 4-1 Oracle Rdb Documentation

Part Number	Title	PDF Format	HTML Format
A90804-01	Oracle Rdb New and Changed Features for Oracle Rdb, Release 7.1	yes	yes
A90405-01	Oracle Rdb Oracle SQL/Services Server Release 7.1.5 Configuration Guide, Release 7.1	yes	yes
A41741-1	Oracle Rdb7 for OpenVMS Oracle RMU Reference Manual, Release 7.0	yes	yes
A47579-1	Oracle Rdb7 SQL Reference Manual, Volumes 1, 2, and 3, Release 7.0	yes	yes
A40827-1	Oracle Rdb7 Introduction to SQL, Release 7.0	yes	yes
A42867-1	Oracle Rdb7 Guide to SQL Programming, Release 7.0	yes	yes
A41749-1	Oracle Rdb7 Guide to Database Design and Definition, Release 7.0	yes	yes
A41748-1	Oracle Rdb Guide to Database Maintenance, Release 7.0	yes	yes
A41747-1	Oracle Rdb7 Guide to Database Performance and Tuning, Volumes 1 and 2, Release 7.0	yes	yes
A40826-1	Oracle Rdb for OpenVMS Guide to Distributed Transactions, Release 6.1	yes	yes
A41981-1	Oracle Rdb7 Guide to Using Oracle SQL/Services Client API, Release 7.0	yes	yes

Table 4-2 Oracle Rdb Web Agent Documentation

Part Number	Title	PDF Format	HTML Format
n/a		yes	yes

	Oracle Rdb Web Agent Guide to Using Rdb Web Agent, Release 3.0		
--	--	--	--

Table 4–3 SQL*Net for Rdb Documentation

Part Number	Title	PDF Format	HTML Format
A59211–01	Oracle Rdb Guide to SQL*Net for Rdb7, Release 1.0.2	yes	yes
A53248–01	Oracle Rdb A Comparison of SQL Dialects for Oracle and Oracle Rdb, Release 1.0	yes	yes

Table 4–4 Hot Standby Documentation

Part Number	Title	PDF Format	HTML Format
A42860–1	Oracle Rdb and Oracle CODASYL DBMS Guide to Hot Standby Databases, Release 7.0	yes	yes

Table 4–5 Replication Option for Rdb

Part Number	Title	PDF Format	HTML Format
A49000–1	Replication Option for Rdb Handbook, Release 7.0	yes	no
A49001–1	Replication Option for Rdb Installation Guide, Release 7.0	yes	no
n/a	Replication Option for Rdb Release Notes, Release 7.0.0	yes	no
n/a	Replication Option for Rdb Release Notes, Release 7.0.1	yes	no

Table 4–6 Oracle CDD/Repository Documentation

Part Number	Title	PDF Format	HTML Format
A40174–1	Oracle CDD/Repository Using Oracle CDD/Repository on OpenVMS Systems, Release 7.0	yes	no
A31157–1	Oracle CDD/Repository User's Guide Supplement, Release 6.1	yes	no
A24878–1	Oracle CDD/Repository CDO Reference Manual, Release 5.0	yes	no
A24879–2	Oracle CDD/Repository Architecture Manual, Release 5.0	yes	no
A24846–2	Oracle CDD/Repository Callable Interface Manual, Release 5.0	yes	no
A24847–2	Oracle CDD/Repository Information Model Volume 1, Release 5.0	yes	no
A24848–1	Oracle CDD/Repository Information Model Volume 2, Release 5.0	yes	no

Table 4–7 Oracle Expert for Rdb Documentation

Part Number	Title	PDF Format	HTML Format
A32404–1	Oracle Expert for Rdb Getting Started with Oracle Expert for Rdb, Release 3.0	yes	no

A32402-1	Oracle Expert for Rdb Character-Cell Interface Reference Manual, Release 3.0	yes	no
A32403-1	Oracle Expert for Rdb User's Guide, Release 3.0	yes	no
n/a	Oracle Expert for Rdb Power User's Guide, Release 3.0	yes	no

Table 4-8 Oracle Trace for Rdb Documentation

Part Number	Title	PDF Format	HTML Format
A38160-1	Oracle Trace Getting Started, Release 2.2	yes	no
A38159-1	Oracle Trace Collector User's Guide, Release 2.2	yes	no
A38162-1	Oracle Trace Monitor User's Guide, Release 2.2	yes	no
A38161-1	Oracle Trace Reporter User's Guide, Release 2.2	yes	no
A38163-1	Before You Install Oracle Trace Version 2.2 on OpenVMS	yes	no
A38158-1	Oracle Trace Installation Guide, Release 2.2	yes	no

Table 4-9 Oracle Rally Documentation

Part Number	Title	PDF Format	HTML Format
A44323-1	Oracle Rally Introduction to Oracle Rally, Release 7.0	yes	no
A44326-1	Oracle Rally Command Reference Manual, Release 7.0	yes	no
A44325-1	Oracle Rally Using Oracle Rally Applications, Release 7.0	yes	no
A44321-1	Oracle Rally Developing Oracle Rally Applications, Release 7.0	yes	no
A44327-1	Oracle Rally Installing Oracle Rally for OpenVMS Systems, Release 7.0	yes	no
A44328-1	Read Before Installing Oracle Rally Release 7.0 on OpenVMS VAX Systems	yes	no
A44329-1	Read Before Installing Oracle Rally Release 7.0 on OpenVMS Alpha Systems	yes	no

The introductory page, welcome.htm, is located in the top level directory on the Oracle Rdb Documentation CD-ROM. You cannot access the Oracle Rdb Release Notes from this introductory page. See Section 4.4 for information on accessing the release notes.

The following sections provide information about changes or other information that was missing or changed in the Oracle Rdb documentation for release 7.0 and earlier releases.

4.5 Updated Documentation for Oracle Rdb-related Products

Table 4-1 lists the manuals available on the Oracle Rdb Documentation CD-ROM (part number A90838-01). The following are updates to this table:

- Oracle CDD/Repository is currently shipping release 7.0.1 and the release notes, Installing Oracle CDD/Repository (part number A70148-01), and the CDO Reference Manual (part number A70149-01) have been updated along with the CDO help file. You can order these updated books and they are available on the software CD-ROM when you order the 7.0.1 software.
- The Oracle CDD/Repository User's Guide Supplement (part number A31157-1) for release 6.1 is superseded by the Using Oracle CDD/Repository on OpenVMS Systems (part number A40174-1) for release 7.0.

4.6 New and Changed Features in Oracle Rdb Release 7.1

This section provides information about late-breaking new features or information that is missing or changed since the Oracle Rdb New and Changed Features for Oracle Rdb manual was published.

4.6.1 PERSONA is Supported in Oracle SQL/Services

In the "New and Changed Features for Oracle Rdb" Manual under the section "ALTER DATABASE Statement" is a note stating that impersonation is not supported in Oracle SQL/Services. This is incorrect. There was a problem in the first release of Oracle Rdb 7.1 (7.1.0) whereby impersonation through Oracle SQL/Services failed. This problem is resolved in Oracle Rdb Release 7.1.0.1.

4.6.2 NEXTVAL and CURRVAL Pseudocolumns Can Be Delimited Identifiers

The Oracle Rdb New and Changed Features for Oracle Rdb manual describes SEQUENCES but does not mention that the special pseudocolumns NEXTVAL and CURRVAL can be delimited. All uppercase and lowercase variations of these keywords are accepted and assumed to be equivalent to these uppercase keywords.

The following example shows that any case is accepted:

```
SQL> set dialect 'sql92';
SQL> create sequence dept_id;
SQL> select dept_id.nextval from rdb$database;
          1
1 row selected
SQL> select "DEPT_ID".currval from rdb$database;
          1
1 row selected
SQL> select "DEPT_ID"."CURRVAL" from rdb$database;
          1
1 row selected
SQL> select "DEPT_ID"."nextval" from rdb$database;
          2
1 row selected
SQL> select "DEPT_ID"."CuRrVaL" from rdb$database;
          2
1 row selected
```

4.6.3 Only=select_list Qualifier for the RMU Dump After_Journal Command

The Oracle Rdb New and Changed Features for Oracle Rdb manual documents the First=select_list and Last=select_list qualifiers for the RMU Dump After_Journal command. Inadvertently missed was the Only=select_list qualifier.

The First, Last, and Only qualifiers have been added because the Start and End qualifiers are difficult to use since users seldom know, nor can they determine, the AIJ record number in advance of using the RMU Dump After_Journal command.

The select_list clause of these qualifiers consists of a list of one or more of the following keywords:

- TSN=tsn
Specifies the first, last, or specific TSN in the AIJ journal using the standard [n:]m TSN format.
- TID=tid
Specifies the first, last or specific TID in the AIJ journal.
- RECORD=record
Specifies the first or last record in the AIJ journal. This is the same as the existing Start and End qualifiers (which are still supported, but deprecated). This keyword cannot be used with the Only qualifier.
- BLOCK=block#
Specifies the first or last block in the AIJ journal. This keyword cannot be used with the Only qualifier.
- TIME=date_time
Specifies the first or last date/time in the AIJ journal using the standard date/time format. This keyword cannot be used with the Only qualifier.

The First, Last, and Only qualifiers are optional. You may specify any or none of them.

The keywords specified for the First qualifier can differ from the keywords specified for the other qualifiers.

For example, to start the dump from the fifth block of the AIJ journal, you would use the following command:

```
RMU/DUMP/AFTER_JOURNAL /FIRST=(BLOCK=5) MF_PERSONNEL.AIJ
```

To start the dump from block 100 or TSN 52, whichever occurs first, you would use the following command:

```
RMU/DUMP/AFTER_JOURNAL /FIRST=(BLOCK=100,TSN=0:52) MF_PERSONNEL.AIJ
```

When multiple keywords are specified for a qualifier, the first condition being encountered activates the qualifier. In the preceding example, the dump starts when either block 100 or TSN 52 is encountered.

Be careful when searching for TSNs or TIDs as they are not ordered in the AIJ journal. For example, if you want to search for a specific TSN, use the Only qualifier and not the First and Last qualifiers. For example, assume the AIJ journal contains records for TSN 150, 170, and 160 (in that order). If you specify the First=TSN=160 and Last=TSN=160 qualifiers, nothing will be dumped because TSN 170 will match the Last=TSN=160 criteria.

4.7 Oracle Rdb7 and Oracle CODASYL DBMS Guide to Hot Standby Databases

This section provides information that is missing from or changed in V7.0 of the Oracle Rdb7 and Oracle CODASYL DBMS Guide to Hot Standby Databases.

4.7.1 Restrictions Lifted on After-Image Journal Files

The Hot Standby software has been enhanced regarding how it handles after-image journal files. Section 4.2.4 in the Oracle Rdb and Oracle CODASYL DBMS Guide to Hot Standby Databases states the following information:

```
If an after-image journal switchover operation is suspended when
replication operations are occurring, you must back up one or more of
the modified after-image journals to add a new journal file.
```

This restriction has been removed. Now, you can add journal files or use the emergency AIJ feature of Oracle Rdb release 7.0 to automatically add a new journal file. Note the following distinctions between adding an AIJ file and adding an emergency AIJ file:

- You can add an AIJ file to the master database and it will be replicated on the standby database. If replication operations are active, the AIJ file is created on the standby database immediately. If replication operations are not active, the AIJ file is created on the standby database when replication operations are restarted.
- You can add emergency AIJ files anytime. If replication operations are active, the emergency AIJ file is created on the standby database immediately. However, because emergency AIJ files are not journaled, starting replication after you create an emergency AIJ will fail. You cannot start replication operations because the Hot Standby software detects a mismatch in the number of after-image journal files on the master compared to the standby database.

If an emergency AIJ file is created on the master database when replication operations are not active, you must perform a master database backup and then restore the backup on the standby database. Otherwise, an AIJSIGNATURE error results.

4.7.2 Changes to RMU Replicate After_Journal ... Buffer Command

The behavior of the RMU Replicate After_Journal ... Buffers command has been changed. The Buffers qualifier may be used with either the Configure option or the Start option.

When using local buffers, the AIJ Log Roll-forward Server will use a minimum of 4096 buffers. The value provided to the Buffers qualifier will be accepted but ignored if it is less than 4096. In addition, further parameters will be checked and the number of buffers may be increased if the resulting calculations are greater than the number of buffers specified by the Buffers qualifier. If the database is configured to use more than 4096 AIJ Request Blocks (ARBs), then the number of buffers may be increased to the number of ARBs configured for the database. The LRS ensures that there are at least 10 buffers for every possible storage area in the database. Thus if the total number of storage areas (both used and reserved) multiplied by 10 results in a greater number of buffers, then that number will be used.

When global buffers are used, the number of buffers used by the AIJ Log Roll-forward Server is determined as follows:

- If the Buffers qualifier is omitted and the Online qualifier is specified, then the number of buffers will default to the previously configured value, if any, or 256, whichever is larger.
- If the Buffers qualifier is omitted and the Online qualifier is not specified or the Nonline qualifier is specified, then the number of buffers will default to the maximum number of global buffers allowed per user ("USER LIMIT"), or 256, whichever is larger.
- If the Buffers qualifier is specified then that value must be at least 256, and it may not be greater than the maximum number of global buffers allowed per user ("USER LIMIT").

The Buffer qualifier now enforces a minimum of 256 buffers for the AIJ Log Roll-forward Server. The maximum number of buffers allowed is still 524288 buffers.

4.7.3 Unnecessary Command in the Hot Standby Documentation

There is an unnecessary command documented in the Oracle Rdb and Oracle CODASYL DBMS Guide to Hot Standby Databases manual. The documentation (in Section 2.12 "Step 10: Specify the Network Transport Protocol") says that to use TCP/IP as the network protocol, you must issue the following commands:

```
$ CONFIG UCX AIJSERVER OBJECT
$ UCX SET SERVICE RDMAIJSRV
/PORT=n
/USER_NAME=RDMAIJSERVER
/PROCESS_NAME=RDMAIJSERVER
/FILE=SYS$SYSTEM:rdmajserver_ucx.com
/LIMIT=nn
```

The first of these commands (\$ CONFIG UCX AIJSERVER OBJECT) is unnecessary. You can safely disregard the first line when setting up to use TCP/IP with Hot Standby.

The documentation will be corrected in a future release of Oracle Rdb.

4.7.4 Change in the Way RDMAIJ Server is Set Up in UCX

Starting with Oracle Rdb Release 7.0.2.1, the RDMAIJ image became a variant image. Therefore, the information in Section 2.12, "Step 10: Specify the Network Transport Protocol," of the Oracle Rdb7 and Oracle CODASYL DBMS Guide to Hot Standby Databases has become outdated with regard to setting up the RDMAIJSERVER object when using UCX as the network transport protocol. The UCX SET SERVICE command is now similar to the following:

```
$ UCX SET SERVICE RDMAIJ -
/PORT=port_number -
/USER_NAME=RDMAIJ -
/PROCESS_NAME=RDMAIJ -
/FILE=SYS$SYSTEM:RDMAIJSERVER.com -
/LIMIT=limit
```

For Oracle Rdb multiversion, the UCX SET SERVICE command is similar to the following:

```
$ UCX SET SERVICE RDMAIJ70 -
/PORT=port_number -
/USER_NAME=RDMAIJ70 -
/PROCESS_NAME=RDMAIJ70 -
/FILE=SYS$SYSTEM:RDMAIJSERVER70.com -
/LIMIT=limit
```

The installation procedure for Oracle Rdb creates a user named RDMAIJ(nn) and places a file called RDMAIJSERVER(nn).COM in SYS\$SYSTEM. The RMONSTART(nn).COM command procedure will try

to enable a service called RDMAIJ(nn) if UCX is installed and running.

Changing the RDMAIJ server to a variant image does not impact installations using DECNet since the correct DECNet object is created during the Oracle Rdb installation.

4.7.5 CREATE INDEX Operation Supported for Hot Standby

On Page 1–13 of the Oracle Rdb7 and Oracle CODASYL DBMS Guide to Hot Standby Databases, the add new index operation is incorrectly listed as an offline operation not supported by Hot Standby. The CREATE INDEX operation is now fully supported by Hot Standby, as long as the transaction does not span all available AIJ journals, including emergency AIJ journals.

4.8 Oracle Rdb7 for OpenVMS Installation and Configuration Guide

This section provides information that is missing from or changed in V7.0 of the Oracle Rdb7 for OpenVMS Installation and Configuration Guide.

4.8.1 Suggestion to Increase GH_RSRVPGCNT Removed

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

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

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

4.8.2 Prerequisite Software

In addition to the software listed in the Oracle Rdb Installation and Configuration Guide and at the url http://www.oracle.com/rdb/product_info/index.html, note that the MACRO compiler and linker from Compaq Computer Corporation are required software in order to install Oracle Rdb on your OpenVMS Alpha system.

4.8.3 Defining the RDBSERVER Logical Name

Sections 4.3.7.1 and 4.3.7.2 in the Oracle Rdb7 for OpenVMS Installation and Configuration Guide provide the following examples for defining the RDBSERVER logical name: *\$ DEFINE RDBSERVER SYS\$SYSTEM:RDBSERVER70.EXE*

and *\$ DEFINE RDBSERVER SYS\$SYSTEM:RDBSERVER61.EXE*

These definitions are inconsistent with other command procedures that attempt to reference the RDBSERVERxx.EXE image. Below is one example where the RDBSERVER.COM procedure references SYS\$COMMON:<SYSEXE> and SYS\$COMMON:[SYSEXE] rather than SYS\$SYSTEM.

```
$ if .not. -
    ((f$locate ("SYS$COMMON:<SYSEXE>", rdbserver_image) .ne. log_len) .or. -
    (f$locate ("SYS$COMMON:[SYSEXE]", rdbserver_image) .ne. log_len))
$ then
$   say "'rdbserver_image' is not found in SYS$COMMON:<SYSEXE>"
$   say "RDBSERVER logical is 'rdbserver_image'"
$   exit
$ endif
```

In this case, if the logical name were defined as instructed in the Oracle Rdb7 for OpenVMS Installation and Configuration Guide, the image would not be found.

The correct definition of the logical name is as follows: *DEFINE RDBSERVER
SYS\$COMMON:<SYSEXE>RDBSERVER70.EXE*

and *DEFINE RDBSERVER SYS\$COMMON:<SYSEXE>RDBSERVER61.EXE*

4.9 Guide to Database Design and Definition

This section provides information that is missing from or changed in release 7.0 of the Oracle Rdb7 Guide to Database Design and Definition.

4.9.1 Lock Timeout Interval Logical Incorrect

On Page 7–31 of Section 7.4.8 in the Oracle Rdb7 Guide to Database Design and Definition, the RDM\$BIND_LOCK_TIMEOUT logical name is referenced incorrectly. The correct logical name is RDM\$BIND_LOCK_TIMEOUT_INTERVAL.

The Oracle Rdb7 Guide to Database Design and Definition will be corrected in a future release.

4.9.2 Example 4–13 and Example 4–14 Are Incorrect

Example 4–13 showing vertical partitioning, and Example 4–14, showing vertical and horizontal partitioning, are incorrect. They should appear as follows:

Example 4–13:

```
SQL> CREATE STORAGE MAP EMPLOYEES_1_MAP
cont>     FOR EMPLOYEES
cont>     ENABLE COMPRESSION
cont>     STORE COLUMNS (EMPLOYEE_ID, LAST_NAME, FIRST_NAME,
cont>                     MIDDLE_INITIAL, STATUS_CODE)
cont>     DISABLE COMPRESSION
cont>     IN ACTIVE_AREA
cont>     STORE COLUMNS (ADDRESS_DATA_1, ADDRESS_DATA_2, CITY,
cont>                     STATE, POSTAL_CODE)
cont>     IN INACTIVE_AREA
cont>     STORE IN OTHER_AREA;
```

Example 4–14:

```
SQL> CREATE STORAGE MAP EMPLOYEES_1_MAP2
cont>     FOR EMP2
cont>     STORE COLUMNS (EMPLOYEE_ID, LAST_NAME, FIRST_NAME,
cont>                     MIDDLE_INITIAL, STATUS_CODE)
cont>     USING (EMPLOYEE_ID)
cont>     IN ACTIVE_AREA_A WITH LIMIT OF ('00399')
cont>     IN ACTIVE_AREA_B WITH LIMIT OF ('00699')
cont>     OTHERWISE IN ACTIVE_AREA_C
cont>     STORE COLUMNS (ADDRESS_DATA_1, ADDRESS_DATA_2, CITY,
cont>                     STATE, POSTAL_CODE)
cont>     USING (EMPLOYEE_ID)
cont>     IN INACTIVE_AREA_A WITH LIMIT OF ('00399')
cont>     IN INACTIVE_AREA_B WITH LIMIT OF ('00699')
cont>     OTHERWISE IN INACTIVE_AREA_C
cont>     STORE IN OTHER_AREA;
```

4.10 Oracle Rdb7 SQL Reference Manual

This section provides information that is missing from or changed in V7.0 of the Oracle Rdb7 SQL Reference Manual.

4.10.1 Clarification of the DDLDONOTMIX Error Message

The ALTER DATABASE statement performs two classes of functions:

1. Changing the database root structures in the .RDB file
2. Modifying the system metadata in the RDB\$SYSTEM storage area.

The first class of changes do not require a transaction to be active. However, the second class requires that a transaction be active. Oracle Rdb does not currently support the mixing of these two classes of ALTER DATABASE clauses.

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

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

The following clauses may be mixed with each other, but may not appear with other clauses such as ADD STORAGE AREA or ADD CACHE:

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

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

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

4.10.2 Node Specification Allowed on Root FILENAME Clauses

In previous releases of the Oracle Rdb SQL Reference Manual, it was not made clear that a node specification may only be specified for the root FILENAME clause of the ALTER DATABASE, CREATE DATABASE, EXPORT DATABASE, and IMPORT DATABASE statements.

This means that the directory or file specification specified with the following clauses can only be a device, directory, file name, and file type:

- LOCATION clause of the ROW CACHE IS ENABLED, RECOVERY JOURNAL, ADD CACHE, and CREATE CACHE clauses
- SNAPSHOT FILENAME clause
- FILENAME and SNAPSHOT FILENAME clauses of the ADD STORAGE AREA and CREATE STORAGE AREA clauses
- BACKUP FILENAME clause of the JOURNAL IS ENABLED, ADD JOURNAL, and ALTER JOURNAL clauses
- BACKUP SERVER and CACHE FILENAME clauses of the JOURNAL IS ENABLED clause
- FILENAME clause of the ADD JOURNAL clause

Usage notes reflecting this restriction for these clauses will appear in a future release of the Oracle Rdb SQL Reference Manual.

4.10.3 Incorrect Syntax Shown for Routine–Clause of the CREATE MODULE Statement

The Oracle Rdb7 SQL Reference Manual incorrectly showed that a simple–statement could be specified for the routine–clause of the CREATE MODULE statement. You can specify a compound–statement and compound–use–statement for the routine–clause only of the CREATE MODULE statement.

This correction appears in the Oracle Rdb New and Changed Features for Oracle Rdb manual and will appear in a future release of the Oracle Rdb7 SQL Reference Manual.

4.10.4 Omitted SET Statements

The following SET statements and language options were omitted from the Oracle Rdb7 SQL Reference Manual.

4.10.4.1 QUIET COMMIT

The following QUIET COMMIT options were omitted from the documentation:

```
Interactive and dynamic SET QUIET COMMIT statement
SQL
Module Header           QUIET COMMIT option
SQL Module Language     /QUIET_COMMIT and /NOQUIET_COMMIT qualifiers
SQL Precompiler         /SQLOPTIONS=QUIET_COMMIT and
                        /SQLOPTIONS=NOQUIET_COMMIT options
```

These options control the behavior of the COMMIT and ROLLBACK statements in cases where there is no active transaction.

By default, if there is no active transaction, SQL will raise an error when COMMIT or ROLLBACK is executed. This default is retained for backward compatibility for applications that wish to detect the situation. If QUIET COMMIT is set to ON, a COMMIT or ROLLBACK executes successfully when there is no active transaction.

Within a compound statement, the COMMIT and ROLLBACK statements are ignored.

In interactive or dynamic SQL, the SET statement can be used to disable or enable error reporting for COMMIT and ROLLBACK when no transaction is active. The parameter to the SET command is a string

literal or host variable containing the keyword ON or OFF. For example:

```
SQL> COMMIT;
%SQL-F-NO_TXNOUT, No transaction outstanding
SQL> ROLLBACK;
%SQL-F-NO_TXNOUT, No transaction outstanding
SQL> SET QUIET COMMIT 'on';
SQL> ROLLBACK;
SQL> COMMIT;
SQL> SET QUIET COMMIT 'off';
SQL> COMMIT;
%SQL-F-NO_TXNOUT, No transaction outstanding
```

In the SQL module language or precompiler header, the QUIET COMMIT option can be used to disable or enable error reporting for COMMIT and ROLLBACK when no transaction is active. The keyword ON or OFF must be used to enable or disable this feature. The following example enables QUIET COMMIT so that no error is reported if a COMMIT is executed when no transaction is active:

```
MODULE TXN_CONTROL
LANGUAGE BASIC
PARAMETER COLONS
QUIET COMMIT ON

PROCEDURE S_TXN (SQLCODE);
SET TRANSACTION READ WRITE;

PROCEDURE C_TXN (SQLCODE);
COMMIT;
```

4.10.4.2 COMPOUND TRANSACTIONS

The SET COMPOUND TRANSACTIONS statement (for interactive and dynamic SQL) and the module header option, COMPOUND TRANSACTIONS, controls the SQL behavior for starting default transactions for compound statements.

By default, if there is no current transaction, SQL will start a transaction before executing a compound statement or stored procedure. However, this may conflict with the actions within the procedure or may start a transaction for no reason if the procedure body does not perform database access. This default is retained for backward compatibility for applications which may expect a transaction to be started for the procedure.

If COMPOUND TRANSACTIONS is set to EXTERNAL, SQL starts a transaction before executing the procedure. Otherwise, if it is set to INTERNAL, it allows the procedure to start a transaction as required by the procedure execution.

In interactive or dynamic SQL, the following SET command can be used to disable or enable transactions starting by the SQL interface. The parameter to the SET command is a string literal or host variable containing the keyword 'INTERNAL' or 'EXTERNAL'.

```
SQL> SET COMPOUND TRANSACTIONS 'internal';
SQL> CALL START_TXN_AND_COMMIT ();
SQL> SET COMPOUND TRANSACTIONS 'external';
SQL> CALL UPDATE_EMPLOYEES (...);
```

In the SQL module language or precompiler header, the COMPOUND TRANSACTIONS option can be used to disable or enable starting a transaction for procedures. The keyword INTERNAL or EXTERNAL must be used to enable or disable this feature.

```

MODULE TXN_CONTROL
LANGUAGE BASIC
PARAMETER COLONS
COMPOUND TRANSACTIONS INTERNAL

PROCEDURE S_TXN (SQLCODE);
BEGIN
SET TRANSACTION READ WRITE;
END;

PROCEDURE C_TXN (SQLCODE);
BEGIN
COMMIT;
END;

```

4.10.5 Size Limit for Indexes with Keys Using Collating Sequences

When a column is defined with a collating sequence, the index key is specially encoded to incorporate the correct collating information. This special encoding takes more space than keys encoded for ASCII (which is the default when no collating sequence is used). Therefore, the encoded string uses more than the customary one byte per character of space within the index. This is true for all versions of Oracle Rdb which support collating sequences.

For all collating sequences, except Norwegian, the space required is approximately 9 bytes for every 8 characters. Therefore, a CHAR (24) column will require approximately 27 bytes to store. For Norwegian collating sequences, the space required is approximately 10 bytes for every 8 characters.

The space required for encoding the string must be taken into account when calculating the size of an index key against the limit of 255 bytes. Suppose a column defined with a collating sequence of GERMAN was used in an index. The length of that column is limited to a maximum of 225 characters because the key will be encoded in 254 bytes.

The following example demonstrates how a 233 character column, defined with a German collating sequence and included in an index, exceeds the index size limit of 255 bytes, even though the column is defined as less than 255 characters in length.

```

SQL> CREATE DATABASE
cont>     FILENAME 'testdb.rdb'
cont>     COLLATING SEQUENCE GERMAN GERMAN;
SQL> CREATE TABLE employee_info
cont>     (emp_name CHAR (233));
SQL> CREATE INDEX emp_name_idx
cont>     ON employee_info (
cont>     emp_name     ASC)
cont>     TYPE IS SORTED;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-INDTOOBIG, requested index is too big

```

4.10.6 Clarification of SET FLAGS Option DATABASE_PARAMETERS

The Oracle Rdb7 SQL Reference Manual described the option DATABASE_PARAMETERS in table 7–6 in the SET FLAGS section. However, this keyword generates output only during ATTACH to the database

which happens prior to the SET FLAGS statement executing.

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

```
$ define RDMS$SET_FLAGS "database_parameters"
$ sql$
SQL> Attach 'File db$:scratch';
ATTACH #1, Database BLUGUM$DKA300:[SMITHI.DATABASES.V70]SCRATCH.RDB;1
~P Database Parameter Buffer (version=2, len=79)
0000 (00000) RDB$K_DPB_VERSION2
0001 (00001) RDB$K_FACILITY_ALL
0002 (00002) RDB$K_DPB2_IMAGE_NAME "NODE::DISK:[DIR]SQL$70.EXE;1"
0040 (00064) RDB$K_FACILITY_ALL
0041 (00065) RDB$K_DPB2_DBKEY_SCOPE (Transaction)
0045 (00069) RDB$K_FACILITY_ALL
0046 (00070) RDB$K_DPB2_REQUEST_SCOPE (Attach)
004A (00074) RDB$K_FACILITY_RDB_VMS
004B (00075) RDB$K_DPB2_CDD_MAINTAINED (No)
RDMS$BIND_WORK_FILE = "DISK:[DIR]RDMSTTBL$UEOU3LQ0RV2.TMP;" (Visible = 0)
SQL> Exit
DETACH #1
```


4.10.7 Incorrect Syntax for CREATE STORAGE MAP Statement

The main diagram of the CREATE STORAGE MAP statement incorrectly shows the partition-clause as required syntax. The partition-clause is not a required clause.

The partition-clause diagram of the CREATE STORAGE MAP statement incorrectly indicated that the STORE keyword was not repeated. When creating a vertically partitioned table you must repeat the STORE keyword for each partition.

FORMAT

partition-clause =



When creating a vertical record partition, the last STORE clause cannot contain the COLUMNS clause. If you attempt to include the COLUMNS clause on the last STORE clause, you will an error similar to the following:

```
%SQL-F-VRP_ILLEGAL_STO, Storage Map "EMPLOYEES_MAP2" specified STORE COLUMNS
after a STORE
```

The following example shows the correct syntax for creating a storage map with horizontal and vertical partitions:

```
SQL> CREATE STORAGE MAP employees_map2
cont>     FOR employees2
cont> --
cont> -- Store the primary information horizontally partitioned
cont> -- across the areas EMPIDS_LOW, EMPIDS_MID and EMPIDS_OVER.
cont> -- Disable compression because these columns are accessed often.
cont> --
```

```

cont>     STORE
cont>         COLUMNS (employee_id, last_name,
cont>                     first_name, middle_initial)
cont>     VERTICAL PARTITION volatile_columns
cont>         DISABLE COMPRESSION
cont>         USING (employee_id)
cont>             IN empids_low
cont>             (PARTITION id_low)
cont>                 WITH LIMIT OF ('00200')
cont>             IN empids_mid
cont>             (PARTITION id_mid)
cont>                 WITH LIMIT OF ('00400')
cont>         OTHERWISE IN empids_over
cont>             (partition id_ovr)
cont> --
cont> -- Place all the address information in EMP_INFO.
cont> -- Make sure these character columns are compressed.
cont> --
cont>     STORE
cont>         COLUMNS (address_data_1, address_data_2, city, state,
cont>                     postal_code)
cont>         ENABLE COMPRESSION
cont>         IN emp_info
cont> --
cont> -- The remaining columns get written randomly over these area.
cont> --
cont>     STORE
cont>         ENABLE COMPRESSION
cont>         RANDOMLY ACROSS (salary_history, jobs);

```

Refer to Oracle Rdb New and Changed Features for Oracle Rdb for the full syntax of the CREATE STORAGE MAP statement. The Oracle Rdb7 SQL Reference Manual will be corrected in a future release.

4.10.8 Use of SQL_SQLCA Include File Intended for Host Language File

Use of the SQLCA include files such as the SQL_SQLCA.H file for C, are intended for use with the host language files only. That is, only *.C should be including that file. Precompiled files (*.SC files) should use the EXEC SQL INCLUDE SQLCA embedded SQL command in the declaration section of the module. In this way the precompiler can properly define the structure to be used by the related SQL generated code.

Remember that the SQLCA is always scoped at the module level, unlike the SQLCODE or SQLSTATE variables which may be routine specific.

The following example shows this error:

```

#include <stdio.h>
#include <sql_sqlca.h>
struct SQLCA SQLCA;

int main (void)
{
EXEC SQL EXECUTE IMMEDIATE `show version`;
printf ("SQLCODE=%d\n", SQLCA.SQLCODE);
}
$ SQLPRE/CC issues the following error against this program:
%SQL-F-NOSQLCODE, Neither SQLCA, SQLCODE nor SQLSTATE were declared

```

The following example shows correct usage:

```
#include <stdio.h>
#include <sql_sqlca.h>
EXEC SQL INCLUDE SQLCA;

int main (void)
{
EXEC SQL EXECUTE IMMEDIATE `show version`;
printf ("SQLCODE=%d\n", SQLCA.SQLCODE);
}
```

4.10.9 Missing Information on Temporary Tables

The following information was inadvertently omitted from the Oracle Rdb7 SQL Reference Manual. (Should be in the Usage Notes for CREATE TEMPORARY TABLE.)

Data for a temporary table is stored in virtual memory, not in a storage area. For journaling purposes, when changes are made to the data in a temporary table such as updates or deletes, recovery space is required to hold before images of deleted and updated rows. This recovery space also requires virtual memory and may result in having to increase Page File Quota and Virtual Page Count on OpenVMS.

A recommended way to reduce memory usage when using temporary tables is to commit transactions which modify temporary table data as soon as possible. Upon commit the additional copies of data are released and available for reuse by Oracle Rdb. This eliminates extra copies of data and therefore reduces virtual memory usage.

See the Oracle Rdb7 Guide to Database Design and Definition for calculating memory usage for temporary tables.

4.11 Oracle RMU Reference Manual, Release 7.0

This section provides information that is missing from or changed in V7.0 of the Oracle RMU Reference Manual.

4.11.1 RMU Unload After_Journal Null Bit Vector Clarification

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

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

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

```
/* DATATYPES.C */

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

#pragma member_alignment __save
#pragma nomember_alignment

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

typedef struct { /* Null bit vector with one bit for each column */
    unsigned          n_tinyint    :1;
    unsigned          n_smallint   :1;
    unsigned          n_integer    :1;
    unsigned          n_bigint     :1;
    unsigned          n_double     :1;
    unsigned          n_real       :1;
    unsigned          n_fixstr     :1;
    unsigned          n_varstr     :1;
} nbv_t;

struct { /* LogMiner output record structure for table DATATYPES */
    char              rdb$lm_action;
    char              rdb$lm_relation_name [31];
    int               rdb$lm_record_type;
    short             rdb$lm_data_len;
    short             rdb$lm_nbv_len;
    __int64           rdb$lm_dbk;
    __int64           rdb$lm_start_tad;
    __int64           rdb$lm_commit_tad;
    __int64           rdb$lm_tsn;
    short             rdb$lm_record_version;
}
```

```

char          f_tinyint;
short         f_smallint;
int           f_integer;
__int64      f_bigint;
double       f_double;
float        f_real;
char         f_fixstr[10];
short       f_varstr_len; /* length of varchar */
char        f_varstr[10]; /* data of varchar */
nbv_t       nbv;
} lm;

```

```
#pragma member_alignment __restore
```

```
main ()
```

```

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

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

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

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

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

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

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

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

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

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

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

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

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

```



```

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

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

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

        printf ("\n");
    }
}

```

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

```

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

```

4.11.2 New Transaction_Mode Qualifier for Oracle RMU Commands

A new qualifier, `Transaction_Mode`, has been added to the `RMU Copy`, `Move_Area`, `Restore`, and `Restore Only_Root` commands. You can use this qualifier to set the allowable transaction modes for the database root file created by these commands. If you are not creating a root file as part of one of these commands, for example, you are restoring an area, attempting to use this qualifier returns a `CONFLSWIT` error. This qualifier is similar to the `SET TRANSACTION MODE` clause of the `CREATE DATABASE` command in interactive SQL.

The primary use of this qualifier is when you restore a backup file (of the master database) to create a Hot Standby database. Include the `Transaction_Mode` qualifier on the `RMU Restore` command when you create the standby database (prior to starting replication operations). Because only read-only transactions are allowed on the standby database, you should use the `Transaction_Mode=Read_Only` qualifier setting. This setting prevents modifications to the standby database at all times, even when replication operations are not active.

You can specify the following transaction modes for the Transaction_Mode qualifier:

```
All
Current
None
[No]Batch_Update
[No]Read_Only
[No]Exclusive
[No]Exclusive_Read
[No]Exclusive_Write
[No]Protected
[No]Protected_Read
[No]Protected_Write
[No]Shared
[No]Shared_Read
[No]Shared_Write
```

Note that [No] indicates that the value can be negated. For example, the NoExclusive_Write option indicates that exclusive write is not an allowable access mode for this database. If you specify the Shared, Exclusive, or Protected option, Oracle RMU assumes you are referring to both reading and writing in these modes. For example, the Transaction_Mode=Shared option indicates that you want both Shared_Read and Shared_Write as transaction modes. No mode is enabled unless you add that mode to the list or you use the ALL option to enable all modes.

You cannot negate the following three options: All, which enables all transaction modes; None, which disables all transaction modes; and Current, which enables all transaction modes that are set for the source database. If you do not specify the Transaction_Mode qualifier, Oracle RMU uses the transaction modes enabled for the source database.

You can list one qualifier that enables or disables a particular mode followed by another that does the opposite. For example, Transaction_Mode=(NoShared_Write, Shared) is ambiguous because the first value disables Shared_Write access while the second value enables Shared_Write access. Oracle RMU resolves the ambiguities by first enabling all modes that are enabled by the items in the Transaction_Mode list and then disabling those modes that are disabled by items in the Transaction_Mode list. The order of items in the list is irrelevant. In the example discussed, Shared_Read is enabled and Shared_Write is disabled.

The following example shows how to set a newly restored database to allow read-only transactions only. After Oracle RMU executes the command, the database is ready for you to start Hot Standby replication operations.

```
$ RMU/RESTORE/TRANSACTION_MODE=READ_ONLY MF_PERSONNEL.RBF
```

4.11.3 RMU Server After_Journal Stop Command

If database replication is active and you attempt to stop the database AIJ Log Server, Oracle Rdb returns an error. You must stop database replication before attempting to stop the server.

In addition, a new qualifier, Output=filename, has been added to the RMU Server After_Journal Stop command. This optional qualifier allows you to specify the file where the operational log is to be created. The operational log records the transmission and receipt of network messages.

If you do not include a directory specification with the file name, the log file is created in the database root file directory. It is invalid to include a node name as part of the file name specification.

Note that all Hot Standby bugcheck dumps are written to the corresponding bugcheck dump file; bugcheck dumps are not written to the file you specify with the Output qualifier.

4.11.4 Incomplete Description of Protection Qualifier for RMU Backup After_Journal Command

The description of the Protection Qualifier for the RMU Backup After_Journal command is incomplete in the Oracle RMU Reference Manual for Digital UNIX. The complete description is as follows:

The Protection qualifier specifies the system file protection for the backup file produced by the RMU Backup After_Journal command. If you do not specify the Protection qualifier, the default access permissions are `-rw-r-----` for backups to disk or tape.

Tapes do not allow delete or execute access and the superuser account always has both read and write access to tapes. In addition, a more restrictive class accumulates the access rights of the less restrictive classes.

If you specify the Protection qualifier explicitly, the differences in access permissions applied for backups to tape or disk as noted in the preceding paragraph are applied. Thus, if you specify `Protection=(S,O,G:W,W:R)`, the access permissions on tape becomes `rw-rw-r-`.

4.11.5 RMU Extract Command Options Qualifier

A documentation error exists in the description of the `Options=options-list` qualifier of the RMU Extract command. Currently, the documentation states that this qualifier is not applied to output created by the `Items=Volume` qualifier. This is incorrect. Beginning with 6.1 of Oracle Rdb, the behavior of the `Options=options-list` qualifier is applied to output created by the `Items=Volume` qualifier.

4.11.6 RDM\$SNAP_QUIET_POINT Logical is Incorrect

On page 2-72 of the Oracle RMU Reference Manual, the reference to the `RDM$SNAP_QUIET_POINT` logical is incorrect. The correct logical name is `RDM$BIND_SNAP_QUIET_POINT`.

4.11.7 Using Delta Time with RMU Show Statistics Command

Oracle RMU does not support the use of delta time. However, because the OpenVMS platform does, there is a workaround. You can specify delta time using the following syntax with the RMU Show Statistics command:

```
$ RMU/SHOW STATISTICS/OUTPUT=file-spec/UNTIL=" ' ' f$cvtime (" +7:00") ' "
```

The `+7:00` adds 7 hours to the current time.

You can also use `"TOMORROW"` and `"TODAY+n"`.

This information will be added to the description of the `Until` qualifier of the RMU Show Statistics command in a future release of the Oracle RMU Reference Manual.

4.12 Oracle Rdb7 Guide to Database Performance and Tuning

The following section provides corrected, clarified, or omitted information for the Oracle Rdb7 Guide to Database Performance and Tuning manual.

4.12.1 Dynamic OR Optimization Formats

In Table C–2 on Page C–7 of the Oracle Rdb7 Guide to Database Performance and Tuning, the dynamic OR optimization format is incorrectly documented as [l:h...]n. The correct formats for Oracle Rdb Release 7.0 and later are [(l:h)n] and [(l:h,l2:h2)].

4.12.2 Oracle Rdb Logical Names

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

RDM\$BIND_RUJ_ALLOC_BLKCNT Allows you to override the default value of the .ruj file. The block count value can be defined between 0 and 2 billion with a default of 127.

RDM\$BIND_RUJ_EXTEND_BLKCNT Allows you to pre-extend the .ruj files for each process using a database. The block count value can be defined between 0 and 65535 with a default of 127.

4.12.3 Waiting for Client Lock Message

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

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

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

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

The leftmost value seen in the hexadecimal output contains the id of the object. The id is described below for tables and views, routines, modules, storage map areas, and sequences.

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

- For storage map areas, the id represents the physical area id. The "waiting for client lock" message on storage map areas is very rare. This may be raised for databases which have been converted from versions prior to Oracle Rdb 5.1.
- For sequences, the id represents the unique value found in the RDB\$SEQUENCE_ID column of the RDB\$SEQUENCES system relation for the given sequence.

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

Table 4–10 Objects and Their Hexadecimal Type Value

Object	Hexadecimal Value
Tables or views	00000004
Routines	00000006
Modules	00000015
Storage map	0000000E
Sequences	00000019

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

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

```
Process.ID Since..... Stall.reason..... Lock.ID.
46001105:2 10:40:46.38 - waiting for client '.....' 0000001900000000400000055
```

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

```
SQL>select RDB$RELATION_NAME from RDB$RELATIONS where RDB$RELATION_ID = 25;
SQL>select RDB$MODULE_NAME from RDB$MODULES where RDB$MODULE_ID = 12;
SQL>select RDB$ROUTINE_NAME from RDB$ROUTINES where RDB$ROUTINE_ID = 7;
SQL>select RDB$SEQUENCE_NAME from RDB$SEQUENCES where RDB$SEQUENCE_ID = 2;
```

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

For more detailed lock information, perform the following steps:

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

4.12.4 RDMS\$TTB_HASH_SIZE Logical Name

The logical name RDMS\$TTB_HASH_SIZE sets the size of the hash table used for temporary tables. If the logical name is not defined, Oracle Rdb uses a default value of 1249.

If you expect that temporary tables will be large (that is, 10K or more rows), use this logical name to adjust the hash table size to avoid long hash chains. Set the value to approximately 1/4 of the expected maximum number of rows for each temporary table. For example, if a temporary table will be populated with 100,000 rows, define this logical name to be 25000. If there are memory constraints on your system, you should define the logical name to be no higher than this value (1/4 of the expected maximum number of rows).

4.12.5 Error in Updating and Retrieving a Row by Dbkey Example 3–22

Example 3–22 in Section 3.8.3 that shows how to update and retrieve a row by dbkey is incorrect. The example should appear as follows:

```
SQL> ATTACH 'FILENAME MF_PERSONNEL.RDB';
SQL> --
SQL> -- Declare host variables
SQL> --
SQL> DECLARE :hv_row INTEGER;           -- Row counter
SQL> DECLARE :hv_employee_id ID_DOM;    -- EMPLOYEE_ID field
SQL> DECLARE :hv_employee_id_ind SMALLINT; -- Null indicator variable
SQL> --
SQL> DECLARE :hv_dbkey CHAR(8);         -- DBKEY storage
SQL> DECLARE :hv_dbkey_ind SMALLINT;    -- Null indicator variable
SQL> --
SQL> DECLARE :hv_last_name LAST_NAME_DOM;
SQL> DECLARE :hv_new_address_data_1 ADDRESS_DATA_1_DOM;
SQL> --
SQL> SET TRANSACTION READ WRITE;
SQL> BEGIN
cont> --
cont> -- Set the search value for SELECT
cont> --
cont> SET :hv_last_name = 'Ames';
cont> --
cont> -- Set the NEW_ADDRESS_DATA_1 value
cont> --
cont> SET :hv_new_address_data_1 = '100 Broadway Ave.';
cont> END;
SQL> COMMIT;
SQL> --
SQL> SET TRANSACTION READ ONLY;
SQL> BEGIN
cont> SELECT E.EMPLOYEE_ID, E.DBKEY
cont> INTO :hv_employee_id INDICATOR :hv_employee_id_ind,
cont> :hv_dbkey INDICATOR :hv_dbkey_ind
cont> FROM EMPLOYEES E
cont> WHERE E.LAST_NAME = :hv_last_name
cont> LIMIT TO 1 ROW;
cont> --
cont> GET DIAGNOSTICS :hv_row = ROW_COUNT;
cont> END;
SQL> COMMIT;
SQL> --
SQL> SET TRANSACTION READ WRITE RESERVING EMPLOYEES FOR SHARED WRITE;
SQL> BEGIN
cont> IF (:hv_row = 1) THEN
cont> BEGIN
cont> UPDATE EMPLOYEES E
cont> SET E.ADDRESS_DATA_1 = :hv_new_address_data_1
cont> WHERE E.DBKEY = :hv_dbkey;
cont> END;
cont> END IF;
cont> END;
SQL> COMMIT;
SQL> --
SQL> -- Display result of change
SQL> --
SQL> SET TRANSACTION READ ONLY;
SQL> SELECT E.*
cont> FROM EMPLOYEES E
```

```

cont> WHERE E.DBKEY = :hv_dbkey;
EMPLOYEE_ID  LAST_NAME          FIRST_NAME  MIDDLE_INITIAL
ADDRESS_DATA_1  ADDRESS_DATA_2    CITY
STATE  POSTAL_CODE  SEX  BIRTHDAY      STATUS_CODE
00416      Ames          Louie      A
100 Broadway Ave.
NH      03809        M      13-Apr-1941  1

1 row selected
SQL>

```

The new example will appear in a future publication of the Oracle Rdb7 Guide to Database Performance and Tuning manual.

4.12.6 Error in Calculation of Sorted Index in Example 3–46

Example 3–46 in Section 3.9.5.1 shows the output when you use the RMU Analyze Indexes command and specify the Option=Debug qualifier and the DEPARTMENTS_INDEX sorted index.

The description of the example did not include the 8 byte dbkey in the calculation of the sorted index. The complete description is as follows:

The entire index (26 records) is located on pages 2 and 3 in logical area 72 and uses 188 bytes of a possible 430 bytes or the node record is 47 percent full. Note that due to index compression, the node size has decreased in size from 422 bytes to 188 bytes and the percent fullness of the node records has dropped from 98 to 47 percent. Also note that the used/avail value in the summary information at the end of the output does not include the index header and trailer information, which accounts for 32 bytes. This value is shown for each node record in the detailed part of the output. The number of bytes used by the index is calculated as follows: the sort key is 4 bytes plus a null byte for a total of 5 bytes. The prefix is 1 byte and the suffix is 1 byte. The prefix indicates the number of bytes in the preceding key that are the same and the suffix indicates the number of bytes that are different from the preceding key. The dbkey pointer to the row is 8 bytes. There are 26 data rows multiplied by 15 bytes for a total of 390 bytes. The 15 bytes include:

- 7 bytes for the sort key: length + null byte + prefix + suffix
- 8 bytes for the dbkey pointer to the row

Add 32 bytes for index header and trailer information for the index node to the 390 bytes for a total of 422 bytes used. Index compression reduces the number of bytes used to 188 bytes used.

The revised description will appear in a future publication of the Oracle Rdb7 Guide to Database Performance and Tuning manual.

4.12.7 Documentation Error in Section C.7

The Oracle Rdb Guide to Database Performance And Tuning, Volume 2 contains an error in Section C.7 titled Displaying Sort Statistics with the R Flag.

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

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

This is incorrect, the statistics should be described as show below:

- Work File Alloc indicates how much space (in blocks) was allocated in the work files for this sort operation. A zero (0) value indicates that the sort was accomplished completely in memory.

This error will be corrected in a future release of Oracle Rdb Guide to Database Performance And Tuning.

4.12.8 Missing Tables Descriptions for the RDBEXPERT Collection Class

Appendix B in the Oracle Rdb7 Guide to Database Performance and Tuning describes the event-based data tables in the formatted database for the Oracle Rdb PERFORMANCE and RDBEXPERT collection classes. This section describes the missing tables for the RDBEXPERT collection class.

Table 4–11 shows the TRANS_TPB table.

Table 4–11 Columns for Table EPC\$1_221_TRANS_TPB

Column Name	Data Type	Domain
COLLECTION_RECORD_ID	SMALLINT	COLLECTION_RECORD_ID_DOMAIN
IMAGE_RECORD_ID	INTEGER	IMAGE_RECORD_ID_DOMAIN
CONTEXT_NUMBER	INTEGER	CONTEXT_NUMBER_DOMAIN
TIMESTAMP_POINT	DATE VMS	
CLIENT_PC	INTEGER	
STREAM_ID	INTEGER	
TRANS_ID	VARCHAR(16)	
TRANS_ID_STR_ID	INTEGER	STR_ID_DOMAIN
TPB	VARCHAR(127)	
TPB_STR_ID	INTEGER	STR_ID_DOMAIN

Table 4–12 shows the TRANS_TPB_ST table. An index is provided for this table. It is defined with column STR_ID, duplicates are allowed, and the type is sorted.

Table 4–12 Columns for Table EPC\$1_221_TRANS_TPB_ST

Column Name	Data Type	Domain
STR_ID	INTEGER	STR_ID_DOMAIN
SEGMENT_NUMBER	SMALLINT	SEGMENT_NUMBER_DOMAIN
STR_SEGMENT	VARCHAR(128)	

4.12.9 Missing Columns Descriptions for Tables in the Formatted Database

Some of the columns were missing from the tables in Appendix B in the Oracle Rdb7 Guide to Database Performance and Tuning. The complete table definitions are described in this section.

Table 4–13 shows the DATABASE table.

Table 4–13 Columns for Table EPC\$1_221_DATABASE

Column Name	Data Type	Domain
COLLECTION_RECORD_ID	SMALLINT	COLLECTION_RECORD_ID_DOMAIN
IMAGE_RECORD_ID	INTEGER	IMAGE_RECORD_ID_DOMAIN
CONTEXT_NUMBER	INTEGER	CONTEXT_NUMBER_DOMAIN
TIMESTAMP_POINT	DATE VMS	
CLIENT_PC	INTEGER	
STREAM_ID	INTEGER	
DB_NAME	VARCHAR(255)	
DB_NAME_STR_ID	INTEGER	STR_ID_DOMAIN
IMAGE_FILE_NAME	VARCHAR(255)	
IMAGE_FILE_NAME_STR_ID	INTEGER	STR_ID_DOMAIN

Table 4–14 shows the REQUEST_ACTUAL table.

Table 4–14 Columns for Table EPC\$1_221_REQUEST_ACTUAL

Column Name	Data Type	Domain
COLLECTION_RECORD_ID	SMALLINT	COLLECTION_RECORD_ID_DOMAIN
IMAGE_RECORD_ID	INTEGER	IMAGE_RECORD_ID_DOMAIN
CONTEXT_NUMBER	INTEGER	CONTEXT_NUMBER_DOMAIN
TIMESTAMP_START	DATE VMS	
TIMESTAMP_END	DATE VMS	
DBS_READS_START	INTEGER	
DBS_WRITES_START	INTEGER	
RUJ_READS_START	INTEGER	
RUJ_WRITES_START	INTEGER	
AIJ_WRITES_START	INTEGER	
ROOT_READS_START	INTEGER	
ROOT_WRITES_START	INTEGER	
BUFFER_READS_START	INTEGER	
GET_VM_BYTES_START	INTEGER	
FREE_VM_BYTES_START	INTEGER	
LOCK_REQS_START	INTEGER	
REQ_NOT_QUEUED_START	INTEGER	
REQ_STALLS_START	INTEGER	
REQ_DEADLOCKS_START	INTEGER	
PROM_DEADLOCKS_START	INTEGER	
LOCK_RELS_START	INTEGER	
LOCK_STALL_TIME_START	INTEGER	
D_FETCH_RET_START	INTEGER	

D_FETCH_UPD_START	INTEGER
D_LB_ALLOK_START	INTEGER
D_LB_GBNEEDLOCK_START	INTEGER
D_LB_NEEDLOCK_START	INTEGER
D_LB_OLDVER_START	INTEGER
D_GB_NEEDLOCK_START	INTEGER
D_GB_OLDVER_START	INTEGER
D_NOTFOUND_IO_START	INTEGER
D_NOTFOUND_SYN_START	INTEGER
S_FETCH_RET_START	INTEGER
S_FETCH_UPD_START	INTEGER
S_LB_ALLOK_START	INTEGER
S_LB_GBNEEDLOCK_START	INTEGER
S_LB_NEEDLOCK_START	INTEGER
S_LB_OLDVER_START	INTEGER
S_GB_NEEDLOCK_START	INTEGER
S_GB_OLDVER_START	INTEGER
S_NOTFOUND_IO_START	INTEGER
S_NOTFOUND_SYN_START	INTEGER
D_ASYNC_FETCH_START	INTEGER
S_ASYNC_FETCH_START	INTEGER
D_ASYNC_READIO_START	INTEGER
S_ASYNC_READIO_START	INTEGER
AS_READ_STALL_START	INTEGER
AS_BATCH_WRITE_START	INTEGER
AS_WRITE_STALL_START	INTEGER
BIO_START	INTEGER
DIO_START	INTEGER
PAGEFAULTS_START	INTEGER
PAGEFAULT_IO_START	INTEGER
CPU_START	INTEGER
CURRENT_PRIO_START	SMALLINT
VIRTUAL_SIZE_START	INTEGER
WS_SIZE_START	INTEGER
WS_PRIVATE_START	INTEGER
WS_GLOBAL_START	INTEGER
CLIENT_PC_END	INTEGER
STREAM_ID_END	INTEGER
REQ_ID_END	INTEGER
COMP_STATUS_END	INTEGER
REQUEST_OPER_END	INTEGER

TRANS_ID_END	VARCHAR(16)	
TRANS_ID_END_STR_ID	INTEGER	STR_ID_DOMAIN
DBS_READS_END	INTEGER	
DBS_WRITES_END	INTEGER	
RUJ_READS_END	INTEGER	
RUJ_WRITES_END	INTEGER	
AIJ_WRITES_END	INTEGER	
ROOT_READS_END	INTEGER	
ROOT_WRITES_END	INTEGER	
BUFFER_READS_END	INTEGER	
GET_VM_BYTES_END	INTEGER	
FREE_VM_BYTES_END	INTEGER	
LOCK_REQS_END	INTEGER	
REQ_NOT_QUEUED_END	INTEGER	
REQ_STALLS_END	INTEGER	
REQ_DEADLOCKS_END	INTEGER	
PROM_DEADLOCKS_END	INTEGER	
LOCK_RELS_END	INTEGER	
LOCK_STALL_TIME_END	INTEGER	
D_FETCH_RET_END	INTEGER	
D_FETCH_UPD_END	INTEGER	
D_LB_ALLOK_END	INTEGER	
D_LB_GBNEEDLOCK_END	INTEGER	
D_LB_NEEDLOCK_END	INTEGER	
D_LB_OLDVER_END	INTEGER	
D_GB_NEEDLOCK_END	INTEGER	
D_GB_OLDVER_END	INTEGER	
D_NOTFOUND_IO_END	INTEGER	
D_NOTFOUND_SYN_END	INTEGER	
S_FETCH_RET_END	INTEGER	
S_FETCH_UPD_END	INTEGER	
S_LB_ALLOK_END	INTEGER	
S_LB_GBNEEDLOCK_END	INTEGER	
S_LB_NEEDLOCK_END	INTEGER	
S_LB_OLDVER_END	INTEGER	
S_GB_NEEDLOCK_END	INTEGER	
S_GB_OLDVER_END	INTEGER	
S_NOTFOUND_IO_END	INTEGER	
S_NOTFOUND_SYN_END	INTEGER	
D_ASYNC_FETCH_END	INTEGER	
S_ASYNC_FETCH_END	INTEGER	

D_ASYNC_READIO_END	INTEGER
S_ASYNC_READIO_END	INTEGER
AS_READ_STALL_END	INTEGER
AS_BATCH_WRITE_END	INTEGER
AS_WRITE_STALL_END	INTEGER
BIO_END	INTEGER
DIO_END	INTEGER
PAGEFAULTS_END	INTEGER
PAGEFAULT_IO_END	INTEGER
CPU_END	INTEGER
CURRENT_PRIO_END	SMALLINT
VIRTUAL_SIZE_END	INTEGER
WS_SIZE_END	INTEGER
WS_PRIVATE_END	INTEGER
WS_GLOBAL_END	INTEGER

Table 4–15 shows the TRANSACTION table.

Table 4–15 Columns for Table EPC\$1_221_TRANSACTION

Column Name	Data Type	Domain
COLLECTION_RECORD_ID	SMALLINT	COLLECTION_RECORD_ID_DOMAIN
IMAGE_RECORD_ID	INTEGER	IMAGE_RECORD_ID_DOMAIN
CONTEXT_NUMBER	INTEGER	CONTEXT_NUMBER_DOMAIN
TIMESTAMP_START	DATE VMS	
TIMESTAMP_END	DATE VMS	
CLIENT_PC_START	INTEGER	
STREAM_ID_START	INTEGER	
LOCK_MODE_START	INTEGER	
TRANS_ID_START	VARCHAR(16)	
TRANS_ID_START_STR_ID	INTEGER	STR_ID_DOMAIN
GLOBAL_TID_START	VARCHAR(16)	
GLOBAL_TID_START_STR_ID	INTEGER	STR_ID_DOMAIN
DBS_READS_START	INTEGER	
DBS_WRITES_START	INTEGER	
RUJ_READS_START	INTEGER	
RUJ_WRITES_START	INTEGER	
AIJ_WRITES_START	INTEGER	
ROOT_READS_START	INTEGER	
ROOT_WRITES_START	INTEGER	
BUFFER_READS_START	INTEGER	
GET_VM_BYTES_START	INTEGER	

FREE_VM_BYTES_START	INTEGER	
LOCK_REQS_START	INTEGER	
REQ_NOT_QUEUED_START	INTEGER	
REQ_STALLS_START	INTEGER	
REQ_DEADLOCKS_START	INTEGER	
PROM_DEADLOCKS_START	INTEGER	
LOCK_RELS_START	INTEGER	
LOCK_STALL_TIME_START	INTEGER	
D_FETCH_RET_START	INTEGER	
D_FETCH_UPD_START	INTEGER	
D_LB_ALLOK_START	INTEGER	
D_LB_GBNEEDLOCK_START	INTEGER	
D_LB_NEEDLOCK_START	INTEGER	
D_LB_OLDVER_START	INTEGER	
D_GB_NEEDLOCK_START	INTEGER	
D_GB_OLDVER_START	INTEGER	
D_NOTFOUND_IO_START	INTEGER	
D_NOTFOUND_SYN_START	INTEGER	
S_FETCH_RET_START	INTEGER	
S_FETCH_UPD_START	INTEGER	
S_LB_ALLOK_START	INTEGER	
S_LB_GBNEEDLOCK_START	INTEGER	
S_LB_NEEDLOCK_START	INTEGER	
S_LB_OLDVER_START	INTEGER	
S_GB_NEEDLOCK_START	INTEGER	
S_GB_OLDVER_START	INTEGER	
S_NOTFOUND_IO_START	INTEGER	
S_NOTFOUND_SYN_START	INTEGER	
D_ASYNC_FETCH_START	INTEGER	
S_ASYNC_FETCH_START	INTEGER	
D_ASYNC_READIO_START	INTEGER	
S_ASYNC_READIO_START	INTEGER	
AS_READ_STALL_START	INTEGER	
AS_BATCH_WRITE_START	INTEGER	
AS_WRITE_STALL_START	INTEGER	
AREA_ITEMS_START	VARCHAR(128)	
AREA_ITEMS_START_STR_ID	INTEGER	STR_ID_DOMAIN
BIO_START	INTEGER	
DIO_START	INTEGER	
PAGEFAULTS_START	INTEGER	
PAGEFAULT_IO_START	INTEGER	

CPU_START	INTEGER
CURRENT_PRIO_START	SMALLINT
VIRTUAL_SIZE_START	INTEGER
WS_SIZE_START	INTEGER
WS_PRIVATE_START	INTEGER
WS_GLOBAL_START	INTEGER
CROSS_FAC_2_START	INTEGER
CROSS_FAC_3_START	INTEGER
CROSS_FAC_7_START	INTEGER
CROSS_FAC_14_START	INTEGER
DBS_READS_END	INTEGER
DBS_WRITES_END	INTEGER
RUJ_READS_END	INTEGER
RUJ_WRITES_END	INTEGER
AIJ_WRITES_END	INTEGER
ROOT_READS_END	INTEGER
ROOT_WRITES_END	INTEGER
BUFFER_READS_END	INTEGER
GET_VM_BYTES_END	INTEGER
FREE_VM_BYTES_END	INTEGER
LOCK_REQS_END	INTEGER
REQ_NOT_QUEUED_END	INTEGER
REQ_STALLS_END	INTEGER
REQ_DEADLOCKS_END	INTEGER
PROM_DEADLOCKS_END	INTEGER
LOCK_RELS_END	INTEGER
LOCK_STALL_TIME_END	INTEGER
D_FETCH_RET_END	INTEGER
D_FETCH_UPD_END	INTEGER
D_LB_ALLOK_END	INTEGER
D_LB_GBNEEDLOCK_END	INTEGER
D_LB_NEEDLOCK_END	INTEGER
D_LB_OLDVER_END	INTEGER
D_GB_NEEDLOCK_END	INTEGER
D_GB_OLDVER_END	INTEGER
D_NOTFOUND_IO_END	INTEGER
D_NOTFOUND_SYN_END	INTEGER
S_FETCH_RET_END	INTEGER
S_FETCH_UPD_END	INTEGER
S_LB_ALLOK_END	INTEGER
S_LB_GBNEEDLOCK_END	INTEGER

S_LB_NEEDLOCK_END	INTEGER	
S_LB_OLDVER_END	INTEGER	
S_GB_NEEDLOCK_END	INTEGER	
S_GB_OLDVER_END	INTEGER	
S_NOTFOUND_IO_END	INTEGER	
S_NOTFOUND_SYN_END	INTEGER	
D_ASYNC_FETCH_END	INTEGER	
S_ASYNC_FETCH_END	INTEGER	
D_ASYNC_READIO_END	INTEGER	
S_ASYNC_READIO_END	INTEGER	
AS_READ_STALL_END	INTEGER	
AS_BATCH_WRITE_END	INTEGER	
AS_WRITE_STALL_END	INTEGER	
AREA_ITEMS_END	VARCHAR(128)	
AREA_ITEMS_END_STR_ID	INTEGER	STR_ID_DOMAIN
BIO_END	INTEGER	
DIO_END	INTEGER	
PAGEFAULTS_END	INTEGER	
PAGEFAULT_IO_END	INTEGER	
CPU_END	INTEGER	
CURRENT_PRIO_END	SMALLINT	
VIRTUAL_SIZE_END	INTEGER	
WS_SIZE_END	INTEGER	
WS_PRIVATE_END	INTEGER	
WS_GLOBAL_END	INTEGER	
CROSS_FAC_2_END	INTEGER	
CROSS_FAC_3_END	INTEGER	
CROSS_FAC_7_END	INTEGER	
CROSS_FAC_14_END	INTEGER	

Table 4–16 shows the REQUEST_BLR table.

Table 4–16 Columns for Table EPC\$1_221_REQUEST_BLR

Column Name	Data Type	Domain
COLLECTION_RECORD_ID	SMALLINT	COLLECTION_RECORD_ID_DOMAIN
IMAGE_RECORD_ID	INTEGER	IMAGE_RECORD_ID_DOMAIN
CONTEXT_NUMBER	INTEGER	CONTEXT_NUMBER_DOMAIN
TIMESTAMP_POINT	DATE VMS	
CLIENT_PC	INTEGER	
STREAM_ID	INTEGER	
REQ_ID	INTEGER	

TRANS_ID	VARCHAR(16)	
TRANS_ID_STR_ID	INTEGER	STR_ID_DOMAIN
REQUEST_NAME	VARCHAR(31)	
REQUEST_NAME_STR_ID	INTEGER	STR_ID_DOMAIN
REQUEST_TYPE	INTEGER	
BLR	VARCHAR(127)	
BLR_STR_ID	INTEGER	STR_ID_DOMAIN

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

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

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

4.12.11 Using Oracle TRACE Collected Data

The following example shows how the OPTIMIZE AS clause is reflected in the Oracle TRACE database. When a trace collection is started the following SQL commands will record the request names.

```
SQL> attach `file personnel`;
SQL> select last_name, first_name
cont> from employees
cont> optimize as request_one;
.
.
.
SQL> select employee_id
cont> from employees
cont> optimize as request_two;
.
.
.
SQL> select employee_id, city, state
cont> from employees
cont> optimize as request_three;
.
.
.
SQL> select last_name, first_name, employee_id, city, state
cont> from employees
cont> optimize as request_four;
.
.
.
```

Once an Oracle TRACE database has been populated from the collection, a query such as the following can be used to display the request names and types. The type values are described in Table 3–10. The unnamed

queries in this example correspond to the queries executed by interactive SQL to validate the names of the tables and columns referenced in the user-supplied queries.

```
SQL> select REQUEST_NAME, REQUEST_TYPE, TIMESTAMP_POINT
cont> from EPC$1_221_REQUEST_BLR;
REQUEST_NAME          REQUEST_TYPE    TIMESTAMP_POINT
-----
                    1 15-JAN-1997 13:23:27.18
                    1 15-JAN-1997 13:23:27.77
REQUEST_ONE           1 15-JAN-1997 13:23:28.21
REQUEST_TWO           1 15-JAN-1997 13:23:56.55
REQUEST_THREE         1 15-JAN-1997 13:24:57.27
REQUEST_FOUR          1 15-JAN-1997 13:25:25.44
6 rows selected
```

The next example shows the internal query format (BLR) converted to SQL strings after EPC\$EXAMPLES:EPC_BLR_TOSQL_CONVERTER.COM has been run.

```
SQL> SELECT A.REQUEST_NAME, B.SQL_STRING FROM
cont> EPC$1_221_REQUEST_BLR A,
cont> EPC$SQL_QUERIES B
cont> WHERE A.CLIENT_PC = 0 AND A.SQL_ID = B.SQL_ID;
A.REQUEST_NAME
  B.SQL_STRING
REQUEST_ONE
      SELECT C1.LAST_NAME, C1.FIRST_NAME.          FROM EMPLOYEES C1
.
.
.
REQUEST_TWO
      SELECT C1.EMPLOYEE_ID.                      FROM EMPLOYEES C1
.
.
.
REQUEST_THREE
      SELECT C1.EMPLOYEE_ID, C1.CITY, C1.STATE.    FROM EMPLOYEES C1
.
.
.
4 rows selected
```

Table 4–17 shows the Request Types.

Table 4–17 Request Types

Symbolic Name	Value	Comment
RDB_K_REQTYPE_OTHER	0	A query executed internally by Oracle Rdb
RDB_K_REQTYPE_USER_REQUEST	1	A non-stored SQL statement, which includes compound statements
RDB_K_REQTYPE_PROCEDURE	2	A stored procedure
RDB_K_REQTYPE_FUNCTION	3	A stored function
RDB_K_REQTYPE_TRIGGER	4	A trigger action
RDB_K_REQTYPE_CONSTRAINT	5	A table or column constraint

4.12.12 AIP Length Problems in Indexes that Allow Duplicates

When an index allows duplicates, the length stored in the AIP will be 215 bytes, regardless of the actual index node size. Because an index with duplicates can have variable node sizes, the 215-byte size is used as a median length to represent the length of rows in the index's logical area.

When the row size in the AIP is less than the actual row length, it is highly likely that SPAM entries will show space is available on pages when they have insufficient space to store another full size row. This is the most common cause of insert performance problems.

For example, consider a case where an index node size of 430 bytes (a common default value) is used; the page size for the storage area where the index is stored is 2 blocks. After deducting page overhead, the available space on a 2-block page is 982 bytes. Assume that the page in this example is initially empty.

1. A full size (430-byte) index node is stored. As 8 bytes of overhead are associated with each row stored on a page, that leaves $982 - 430 - 8 = 544$ free bytes remaining on the page.
2. A duplicate key entry is made in that index node and thus a duplicate node is created on the same page. An initial duplicate node is 112 bytes long (duplicate nodes can have a variety of sizes depending on when they are created, but for this particular example, 112 bytes is used). Therefore, $544 - 112 - 8 = 424$ free bytes remain on the page.

At this point, 424 bytes are left on the page. That is greater than the 215 bytes that the AIP shows as the row length for the logical area, so the SPAM page shows that the page has space available. However, an attempt to store a full size index node on the page will fail, because the remaining free space (424 bytes) is not enough to store a 430-byte node.

In this case, another candidate page must be selected via the SPAM page, and the process repeats until a page that truly has sufficient free space available is found. In a logical area that contains many duplicate nodes, a significant percentage of the pages in the logical area may fit the scenario just described. When that is the case, and a new full size index node needs to be stored, many pages may need to be read and checked before one is found that can be used to store the row.

It is possible to avoid the preceding scenario by using logical area thresholds. The goal is to set a threshold such that the SPAM page will show a page is full when space is insufficient to store a full size index node.

Using the previous example, here is how to properly set logical area thresholds to prevent excessive pages checked on an index with a 430-byte node size that is stored on a 2-block page. To calculate the proper threshold value to use, you must first determine how full the page can get before no more full size nodes will fit on the page. In this example, a database page can have up to $982 - 430 - 8 = 544$ bytes in use before the page is too full. Therefore, if 544 or fewer bytes are in use, then enough space remains to store another full size node. The threshold is then $544 / 982 = .553971$, or 55%.

In addition, you can determine how full a page must be before a duplicate node of size 112 will no longer fit. In this example, a database page can have up to $982 - 112 - 8 = 862$ bytes in use before the page is too full. Therefore, if 862 or fewer bytes are in use, then enough space remains to store another small duplicates node. The threshold is then $862 / 982 = .8778$, or 88%.

Here is an example of creating an index with the above characteristics:

```
SQL> CREATE INDEX TEST_INDEX ON EMPLOYEES (LAST_NAME)
cont>     STORE IN RDB$SYSTEM
cont>     (THRESHOLD IS (55, 55, 88));
```

These settings mean that any page at over 55% full will not be fetched when inserting a full index node, however, it may be fetched when inserting the smaller duplicates node. When the page is over 88% full then neither a full node nor a duplicate node can be stored, so the page is set as FULL. The lowest setting is not used and so can be set to any value less than or equal to the lowest used threshold.

Note that the compression algorithm used on regular tables that have compression enabled does not apply to index nodes. Index nodes are not compressed like data rows and will always utilize the number of bytes that is specified in the node size. Do not attempt to take into account a compression factor when calculating

thresholds for indexes.

4.12.13 RDM\$BIND_MAX_DBR_COUNT Documentation Clarification

Appendix A in Oracle Rdb7 Guide to Database Performance and Tuning incorrectly describes the use of the RDM\$BIND_MAX_DBR_COUNT logical name.

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

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

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

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

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

4.13 Oracle Rdb7 Guide to SQL Programming

This section provides information that is missing or changed in the Oracle Rdb7 Guide to SQL Programming.

4.13.1 Location of Host Source File Generated by the SQL Precompiler

When the SQL precompiler generates host source files (for example, .c, .pas, or .for) from the precompiler source files, it locates these files based on the Object qualifier in the command given to the SQL precompiler.

The following examples show the location where the host source file is generated.

When the Object qualifier is not specified on the command line, the object and the host source file take the name of the SQL precompiler with the extensions of .obj and .c, respectively. For example:

```
$ sqlpre/cc scc_try_mli_successful.sc
$ dir scc_try_mli_successful.*

Directory MYDISK:[LUND]

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

Total of 3 files.
```

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

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

Directory MYDISK:[LUND]

SCC_TRY_MLI_SUCCESSFUL.SC;2

Total of 1 file.
$ dir myobj.*

Directory MYDISK:[LUND]

MYOBJ.C;1          MYOBJ.OBJ;2

Total of 2 files.

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

Directory MYDISK:[LUND]

SCC_TRY_MLI_SUCCESSFUL.SC;2

Total of 1 file.
$ dir MYDISK:[lund.tmp]scc_try_mli_successful.*

Directory MYDISK:[LUND.TMP]
```

Total of 2 files.

4.13.2 Remote User Authentication

In the Oracle Rdb7 Guide to SQL Programming, Table 15–1 indicates that implicit authorization works from an OpenVMS platform to another OpenVMS platform using TCP/IP. This table is incorrect. Implicit authorization only works using DECnet in this situation.

The Oracle Rdb7 Guide to SQL Programming will be fixed in a future release.

4.13.3 Additional Information About Detached Processes

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

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

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

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

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

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

- Solution 1:

Use the DCL command procedure RUN_PROCEDURE to run the ACCOUNTS application:

RUN_PROCEDURE.COM includes the single line:

```
$ RUN ACCOUNTS_REPORT
```

Then execute this procedure using this command:

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

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

- Solution 2:

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

- ◆ RDMS\$BIND_WORK_FILE

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

For more information on RDMS\$BIND_WORK_FILE and RDMS\$BIND_WORK_VM, see the Oracle Rdb Guide to Database Performance and Tuning.

◆ SORTWORK0, SORTWORK1, and so on

The OpenVMS Sort/Merge utility (SORT/MERGE) attempts to create sort work files in SYSSCRATCH. If the SORTWORK logical names exist, the utility will not require the SYSSCRATCH logical. However, note that not all queries will require sorting, and that some sorts will be completed in memory and so will not necessarily require disk space. If you use the logical RDMS\$BIND_SORT_WORKFILES, you will need to define further SORTWORK logical names as described in the Oracle Rdb Guide to Database Performance and Tuning.

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

4.14 Guide to Using Oracle SQL/Services Client APIs

The following information describes Oracle SQL/Services documentation errors or omissions.

- The Guide to Using Oracle SQL/Services Client APIs does not describe changes to size and format of integer and floating-point data types
Beginning with Oracle SQL/Services V5.1, the size and format of some integer and floating-point data types is changed as follows:

- ◆ Trailing zeros occur in fixed-point numeric data types with SCALE FACTOR.

Trailing zeros are now included after the decimal point up to the number of digits specified by the SCALE FACTOR. In versions of Oracle SQL/Services previous to V5.1, at most one trailing zero was included where the value was a whole number.

The following examples illustrate the changes using a field defined as INTEGER(3):

V5.1 and higher	Versions previous to V5.1
1.000	1.0
23.400	23.4
567.890	567.89

- ◆ Trailing zeros occur in floating-point data types. Trailing zeros are now included in the fraction, and leading zeros are included in the exponent, up to the maximum precision available, for fields assigned the REAL and DOUBLE PRECISION data types.

Data Type	V5.1 and higher	Versions previous to V5.1
REAL	1.2340000E+01	1.234E+1
DOUBLE PRECISION	5.6789000000000000E+001	5.6789E+1

- ◆ Size of TINYINT and REAL data types is changed.

The maximum size of the TINYINT and REAL data types is changed to correctly reflect the precision of the respective data types.

The following table shows the maximum lengths of the data types now and in previous versions:

Data type	V5.1 and higher	Versions previous to V5.1
TINYINT	4	6
REAL	15	24

- The Guide to Using Oracle SQL/Services Client APIs does not describe that the sqlsrv_associate() service returns SQL error code -1028 when connecting to a database service if the user has not been granted the right to attach to the database.

When a user connects to a database service, the sqlsrv_associate() service completes with the SQL error code -1028, SQL_NO_PRIV, if the user has been granted access to the Oracle SQL/Services service, but has not been granted the right to attach to the database. A record of the failure is written to the executor process's log file. Note that the sqlsrv_associate() service completes with the Oracle SQL/Services error code -2034, SQLSRV_GETACCINF if the user has not been granted access to the Oracle SQL/Services service.

4.15 Updates to System Relations

The following sections include updates to system relations that were inadvertently omitted in the SQL Help and Rdb Help files in Release 7.0.

4.15.1 Clarification on Updates to the RDB\$LAST_UPDATED Column for the RDB\$DATABASE System Relation

The ALTER DATABASE statement can be used to change many database attributes, however, only those listed below will cause the RDB\$DATABASE system relation to be changed. The column RDB\$LAST_UPDATED is used to record the date and time when the system relation RDB\$DATABASE is updated and so will change when any of the following clauses are used by ALTER DATABASE.

- CARDINALITY COLLECTION IS { ENABLED | DISABLED }
- DICTIONARY IS [NOT] REQUIRED
- DICTIONARY IS NOT USED
- METADATA CHANGES ARE { ENABLED | DISABLED }
- MULTISCHEMA IS { ON | OFF }
- SECURITY CHECKING IS EXTERNAL (PERSONAL SUPPORT IS { ENABLED | DISABLED })
- SYNONYMS ARE ENABLED
- WORKLOAD COLLECTION IS { ENABLED | DISABLED }

In addition any GRANT and REVOKE statements which use the ON DATABASE clause will cause the RDB\$LAST_UPDATED column to be updated for RDB\$DATABASE.

4.15.2 Missing Descriptions of RDB\$FLAGS

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

Table 4-18 Changed Columns for RDB\$INDICES Table

Column Name	Data Type	Domain Name	Comments
RDB\$FLAGS	integer	RDB\$FLAGS	<p>A bit mask where the bits have the following meaning when set:</p> <ul style="list-style-type: none"> • Bit 0: This index is of type HASHED. • Bit 1: This index uses the MAPPING VALUES clause to compress integer value ranges. • Bit 2: If this is a HASHED index then it is of type ORDERED. If clear this indicates the index if of type SCATTERED. • Bit 3: Reserved for future use. • Bit 4: This index has run length compression enabled (ENABLE COMPRESSION).

			<ul style="list-style-type: none"> • Bit 5: This index is no longer used (MAINTENANCE IS DISABLED). • Bit 6 through 10: Reserved for future use. • Bit 11: This index has duplicates compressed (DUPLICATES ARE COMPRESSED). • Bit 12: This index is of type SORTED RANKED. • Bits 13 through 31: Reserved for future use.
--	--	--	---

Table 4–19 Changed Columns for RDB\$RELATIONS Table

Column Name	Data Type	Domain Name	Comments
RDB\$FLAGS	integer	RDB\$FLAGS	<p>A bit mask where the bits have the following meaning when set:</p> <ul style="list-style-type: none"> • Bit 0: This relation is a view. • Bit 1: This relation is <i>not</i> compressed. • Bit 2: The SQL clause, WITH CHECK OPTION, is used in this view definition. • Bit 3: Indicates a special internal system relation. • Bit 4: This view is not an ANSI updatable view. • Bit 5: This is an imported table in the Distributed Option for Rdb catalog. • Bit 6: This is a passthru table in the Distributed Option for Rdb catalog. • Bit 7: This is a partitioned view in the Distributed Option for Rdb catalog. • Bit 8: This table has compression defined by the storage map. When set Bit 1 in this bit mask is ignored. • Bit 9: This is a temporary table. • Bit 10: When bit 9 is set this is a global temporary table, when clear it indicates a local temporary table. • Bit 11: When bit 9 is set this indicates that the rows in the temporary table should be deleted upon COMMIT. • Bit 12: Reserved for future use. • Bit 13: A table (via a computed by column) or view references a local temporary table. • Bit 14: Reserved for future use. • Bit 15: This is a system table with a

			special storage map. • Bits 16 through 31: Reserved for future use.
--	--	--	--

Table 4–20 Changed Columns for RDB\$STORAGE_MAPS Table

Column Name	Data Type	Domain Name	Comments
RDB\$FLAGS	integer	RDB\$FLAGS	<p>A bit mask where the bits have the following meaning when set:</p> <ul style="list-style-type: none"> • Bit 0: This table or index is mapped to page format MIXED areas. • Bit 1: This partition is <i>not</i> compressed. • Bit 2: This is a strictly partitioned storage map, the partitioning columns become read only for UPDATE. • Bit 3 through 31: Reserved for future use.

4.16 Error Messages

The following subsections further describe or clarify error messages.

4.16.1 Clarification of the DDLDONOTMIX Error Message

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

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

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

The following clauses may be mixed with each other but may not appear with other clauses such as ADD STORAGE AREA or ADD CACHE:

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

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

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

Chapter 5

Known Problems and Restrictions

This chapter describes problems and restrictions relating to Oracle Rdb Release 7.1.0.1, and includes workarounds where appropriate.

5.1 Known Problems and Restrictions in All Interfaces

This section describes known problems and restrictions that affect all interfaces for Release 7.1.0.1.

5.1.1 RDB-E-ARITH_EXCEPT Error From the Rdb Optimizer

Bug 1694309

When using workload statistics, it is possible that a query that joins several tables together can produce a divide by zero error.

The following example shows the result of trying to execute a query that exposes the problem.

```
%RDB-E-ARITH_EXCEPT, truncation of a numeric value at runtime
-SYSTEM-F-HPARITH, high performance arithmetic trap, Imask=00000000,
  Fmask=00000001, summary=04, PC=0000000000F748, PS=0000000B
-SYSTEM-F-FLTDIV, arithmetic trap, floating/decimal divide by zero at
  PC=0000000000F748, PS=0000000B
```

As a side effect of this problem, some queries can be inaccurately costed by the optimizer, which may lead to less than optimal retrieval strategies. The following simple example shows a query where the cardinality is inaccurately calculated from the workload statistics because of this problem.

```
SQL> set flags 'estimates'
SQL> select * from t1, t2 where t1.f1=t2.f1;
Solutions tried 6
Solutions blocks created 4
Created solutions pruned 1
Cost of the chosen solution 1.5162601E+01
Cardinality of chosen solution 0.0000000E+00
~O: Workload statistics used
      T1.F1      T2.F1
      1        1
.
.
.
1000 rows selected
```

The problem can be avoided using any of the following techniques:

- Ensuring that workload data does not have a null factor of exactly 0.0 or 1.0.
- Removing workload statistics.
- Ensuring that the table cardinalities are greater than 1 for all tables in the query.
- Use of the *OLD_COST_MODEL* debug flag.

This problem will be corrected in Oracle Rdb Release 7.1.0.2.

5.1.2 RMU Fails to Perform OPTIMIZER_STATISTICS Actions on Some Databases

Attempts to use RMU/SHOW OPTIMIZER_STATISTICS, RMU/COLLECT OPTIMIZER_STATISTICS, and related commands will fail if the default database character set is not DEC_MCS.

The following example shows the problem for a DEC_KANJI database.

```
$ rmu/show optimizer_statistics DISK1:[TESTING]SAMPLE.RDB
%RDB-F-CONVERT_ERROR, invalid or unsupported data conversion
-RDMS-E-CSETBADCOMPARE, incompatible character sets prohibit the requested
comparison
%RMU-F-FATALRDB, Fatal error while accessing Oracle Rdb.
%RMU-F-FTL_SHOW, Fatal error for SHOW operation at 29-OCT-2001 16:31:20.59
$
$ rmu/collect optimizer_statistics DISK1:[TESTING]SAMPLE.RDB
%RDB-F-CONVERT_ERROR, invalid or unsupported data conversion
-RDMS-E-CSETBADCOMPARE, incompatible character sets prohibit the requested
comparison
%RMU-F-FATALRDB, Fatal error while accessing Oracle Rdb.
%RMU-F-FTL_ANA, Fatal error for ANALYZE operation at 29-OCT-2001 16:31:36.12
```

This problem will be corrected in Oracle Rdb Release 7.1.0.2.

5.1.3 Possible RMU Bugcheck or Failure to Notify Triggering of User Defined Events

The notify or invoke associated with a user defined event in RMU/SHOW STATISTICS may not work or an RMU bugcheck may occur when the user defined event triggers.

This problem will be corrected in Oracle Rdb Release 7.1.0.2.

5.1.4 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
cont>   alter column f2 primary key not deferrable;
SQL> alter table t1
cont>   alter column f1 references t2 not deferrable;
```

When deleting from the PRIMARY KEY table, 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
cont>   alter column f1
cont>   check (f1 in (select * from t2)) not deferrable;
SQL> delete from t2 where f2=1;
Get      Temporary relation      Retrieval by index of relation T2
      Index name  I2 [1:1]
Cross block of 2 entries
Cross block entry 1
      Index only retrieval of relation T1
      Index name  I1 [0:0]
Cross block entry 2
      Conjunct      Aggregate-F1      Conjunct
      Index only retrieval of relation T2
      Index name  I2 [0:0]
%RDB-E-INTEG_FAIL, violation of constraint T1_CHECK1 caused operation to fail
```

The cross block is for the constraint evaluation. This retrieval strategy indicates that to evaluate the constraint, the entire index on table T1 is being scanned and for each key, the entire index in table T2 is being scanned. The behavior can be improved somewhat by using an equality join condition in the select clause of the constraint:

```
SQL> alter table t1
cont>   alter column f1
cont>   check (f1 in (select * from t2 where f2=f1))
cont>   not deferrable;
```

or:

```
SQL> alter table t1
cont>   alter column f1
cont>   check (f1=(select * from t2 where f2=f1))
cont>   not deferrable;
```

In both cases the retrieval strategy will look 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
cont> after insert on t1
cont> when (not exists (select * from t2 where f2=f1))
cont> (error) for each row;
SQL> create trigger t1_update
cont> after update on t1
cont> when (not exists (select * from t2 where f2=f1))
cont> (error) for each row;
SQL> ! A delete trigger is not needed on T1.
SQL> create trigger t2_delete
cont> before delete on t2
cont> when (exists (select * from t1 where f1=f2))
cont> (error) for each row;
SQL> create trigger t2_modify
cont> after update on t2
cont> referencing old as t2o new as t2n
cont> when (exists (select * from t1 where f1=t2o.f2))
cont> (error) for each row;
SQL> ! An insert trigger is not needed on T2.

```

The strategy for a delete on T2 is now:

```

SQL> delete from t2 where f2=1;
Aggregate-F1      Index only retrieval of relation T1
  Index name  I1 [1:1]
Temporary relation      Get      Retrieval by index of relation T2
  Index name  I2 [1:1]
%RDB-E-TRIG_INV_UPD, invalid update; encountered error condition defined for
trigger
-RDMS-E-TRIG_ERROR, trigger T2_DELETE forced an error

```

The trigger strategy is the index only retrieval displayed first. You will note that the index on T1 is used to examine only those rows that may be affected by the delete.

Care must be taken when using this workaround as there are semantic differences in the operation of the triggers, the use of "IN" and "EXISTS", and the use of referential integrity constraints.

This workaround is useful where the form of the constraint is more complex, and cannot be phrased using referential integrity constraints. For example, if the application is such that the value in table T1 may be spaces or NULL to indicate the absence of a value, the above triggers could easily be modified to allow for these semantics.

5.1.5 Using Databases from Releases Earlier Than V5.1

You cannot convert or restore databases earlier than V5.1 directly to V7.1. The RMU Convert command for V7.1 supports conversions from V5.1 through V7.0 only. If you have a V3.0 through V5.0 database, you must convert it to at least V5.1 and then convert it to V7.1. For example, if you have a V4.2 database, convert it first to at least V5.1, then convert the resulting database to V7.1.

If you attempt to convert a database created prior to V5.1 directly to V7.1, Oracle RMU generates an error.

5.1.6 PAGE TRANSFER VIA MEMORY Disabled

Oracle internal testing has revealed that the PAGE TRANSFER VIA MEMORY option for global buffers is not as robust as is needed for the mission critical environments where Oracle Rdb7 is often deployed. This feature has been disabled in release 7.1. Oracle intends to re-enable this feature in a future release.

5.1.7 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.

5.1.8 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.

5.1.9 IMPORT Unable to Import Some View Definitions

View definitions that reference SQL functions, created by the CREATE MODULE statement, cannot be imported by the SQL IMPORT statement. This is because the views are defined before the functions themselves exist.

The following example shows the errors from IMPORT:

```
IMPORTing view TVIEW
%SQL-F-NOVIERES, unable to import view TVIEW
%RDB-E-NO_META_UPDATE, metadata update failed
-RDB-E-OBSOLETE_METADA, request references metadata objects that no
```

```
longer exist
-RDMS-E-RTNNEXTS, routine FORMAT_OUT does not exist in this database
%RDB-E-OBSOLETE_METADA, request references metadata objects that no
longer exist
-RDMS-F-TABNOTDEF, relation TVIEW is not defined in database
```

The following script can be used to demonstrate the problem:

```
create database filename badimp;

create table t (sex char);

create module TFORMAT
  language SQL

  function FORMAT_OUT (:s char)
  returns char(4);
  return (case :s
  when 'F' then 'Female'
  when 'M' then 'Male'
  else NULL
  end);
end module;

create view TVIEW (m_f) as
  select FORMAT_OUT (sex) from t;

commit;

export database filename badimp into exp;
drop database filename badimp;
import database from exp filename badimp;
```

This restriction will be lifted in a future release of Oracle Rdb. Currently the workaround is to save the view definitions and reapply them after the import operation completes.

This restriction does not apply to external functions, created using the CREATE FUNCTION statement, as these database objects are defined before tables and views.

5.1.10 Both Application and Oracle Rdb Using SYS\$HIBER

In application processes that use Oracle Rdb and the \$HIBER system service (possibly through RTL routines such as LIB\$WAIT), the application must ensure that the event being waited for has actually occurred. Oracle Rdb uses \$HIBER/\$WAKE sequences for interprocess communications particularly when the ALS (AIJ Log Server) feature is enabled.

The use of the \$WAKE system service by Oracle Rdb can interfere with other users of \$HIBER (such as the routine LIB\$WAIT) that do not check for event completion, possibly causing a \$HIBER to be unexpectedly resumed without waiting at all.

To avoid these situations, consider altering the application to use a code sequence that avoids continuing without a check for the operation (such as a delay or a timer firing) being complete.

The following pseudo-code shows 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 BEGIN
    LIB$SIGNAL (STAT)
    END
! Hibernate. When the $HIBER completes, check to make
! sure that TIMER_FLAG is set indicating that the wait
! has finished.
WHILE TIMER_FLAG = FALSE
DO BEGIN
    SYS$HIBER()
    END
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 BEGIN
    LIB$SIGNAL (STAT)
    END
END
END

```

The LIB\$K_NOWAKE flag can be specified when using the OpenVMS LIB\$WAIT routine 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.

5.1.11 Bugcheck Dump Files with Exceptions at COSI_CHF_SIGNAL

In certain situations, Oracle Rdb bugcheck dump files indicate an exception at COSI_CHF_SIGNAL. This location is, however, not the address of the actual exception. The actual exception occurred at the previous call frame on the stack (the one listed as the next Saved PC after the exception).

For example, consider the following bugcheck file stack information:

```

$ SEARCH RDSBUGCHK.DMP "EXCEPTION", "SAVED PC", "-F-", "-E-"

***** Exception at 00EFA828 : COSI_CHF_SIGNAL + 00000140
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 00C386F0 : PSIINDEX2JOINSCR + 00000318
Saved PC = 00C0BE6C : PSII2BALANCE + 0000105C
Saved PC = 00C0F4D4 : PSII2INSERTT + 000005CC
Saved PC = 00C10640 : PSII2INSERTTREE + 000001A0
.
.
.

```

In this example, the exception actually occurred at PSIINDEX2JOINSCR offset 00000318. If you have a bugcheck dump with an exception at COSI_CHF_SIGNAL, it is important to note the next "Saved PC" because it is needed when working with Oracle Rdb Worldwide Support.

5.1.12 Read-only Transactions Fetch AIP Pages Too Often

Oracle Rdb read-only transactions fetch Area Inventory Pages (AIP) to ensure that the logical area has not been modified by an exclusive read-write transaction. This check is needed because an exclusive read-write transaction does not write snapshot pages and these pages may be needed by the read-only transaction.

Because AIPs are always stored in the RDB\$SYSTEM area, reading the AIP pages could represent a significant amount of I/O to the RDB\$SYSTEM area for some applications. Setting the RDB\$SYSTEM area to read-only can avoid this problem, but it also prevents other online operations that might be required by the application so it is not a viable workaround in all cases.

This problem has been reduced in Oracle Rdb release 7.0. The AIP entries are now read once and then are not read again unless they need to be. This optimization requires that the carry-over locks feature be enabled (this is the default setting). If carry over locks are not enabled, this optimization is not enabled and the behavior is the same as in previous releases.

5.1.13 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 hotstandby 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.

5.1.14 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 some 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 10. The work files can be individually controlled by the SORTWORKn logical names (where n is from 0 through 9). You can increase the efficiency of sort operations by assigning the location of the temporary sort work files to different disks. These assignments are made by using up to ten logical names, SORTWORK0 through SORTWORK9.

Normally, SORT places work files in the your SYS\$SCRATCH directory. By default, SYS\$SCRATCH is the same device and directory as the SYS\$LOGIN location. Spreading the I/O load over many disks improves efficiency as well as performance by taking advantage of the system resources and helps prevent disk I/O bottlenecks. Specifying that a 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 SYS\$SCRATCH disk device (for example, while Oracle Rdb builds indexes for a very large table). Using the SORTWORK0 through SORTWORK9 logical names can help you avoid this problem.

Note that SORT uses the work files for different sorted runs, and then merges the sorted runs into larger groups. If the source data is mostly sorted, then not every sort work file may need to be accessed. This is a possible source of confusion because even with 10 sort work files, it is possible to exceed the capacity of the first SORT file and the sort operation fails never having accessed the remaining 9 sort work files.

Note that the logical names RDMS\$BIND_WORK_VM and RDMS\$BIND_WORK_FILE do not affect or control the operation of sort. These logical names are used to control other temporary space allocation within Oracle Rdb.

5.1.15 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 5–1 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 2,147,483,647 and the minimum value is 32,768.

5.1.16 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'
cont> order by last_name;
SQL> open emp;
Conjunct      Get      Retrieval by index of relation EMPLOYEES
  Index name  EMP_LAST_NAME [1:0]
SQL> close emp;
```

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
cont> for update of last_name;
SQL> open emp2;
Temporary relation      Conjunct      Get
Retrieval by index of relation EMPLOYEES
  Index name  EMP_LAST_NAME [1:0]
SQL> close emp2;
```

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

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

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

5.2 SQL Known Problems and Restrictions

This section describes known problems and restrictions for the SQL interface for release 7.1.

5.2.1 Interchange File (RBR) Created by Oracle Rdb Release 7.1 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 7.1 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.

5.2.2 Unexpected NO_META_UPDATE Error Generated by DROP MODULE ... CASCADE When Attached by PATHNAME

The SQL DROP MODULE ... CASCADE statement may sometimes generate an unexpected NO_META_UPDATE error. This occurs when the session attaches to a database by PATHNAME. For example:

```
SQL> drop module m1 cascade;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-OBJ_INUSE, object "M1P1" is referenced by M2.M2P1 (usage: Procedure)
-RDMS-E-MODNOTDEL, module "M1" has not been deleted
```

This error occurs because the CASCADE option is ignored because the Oracle CDD/Repository does not support CASCADE. The workaround is to attach by FILENAME and perform the metadata operation.

In a future release of Oracle Rdb, an informational message will be issued describing the downgrade from CASCADE to RESTRICT in such cases.

5.2.3 Problem Exporting and Importing Sequences with ANSI-Style Databases

Exporting and importing sequences defined in an ANSI-style databases may result in an error. An error will occur if a sequence exists in the database with another object imported after the sequence. For example, importing an ANSI-style database which has sequences and modules defined will return an error. For example:

```
%SQL-F-BADCORATT, invalid core attribute 00, 14 in .RBR file
```

This problem will be fixed in a future release of Oracle Rdb.

5.2.4 System Relation Change for International Database Users

Due to an error in creating the RDB\$FIELD_VERSIONS system relation, another system relation, RDB\$STORAGE_MAP_AREAS, cannot be accessed if the session character sets are not set to DEC_MCS.

This problem prevents the new Oracle Rdb GUIs, specifically the Oracle Rdb Schema Manager, from viewing indexes and storage maps from existing Oracle Rdb databases.

The problem can be easily corrected by executing the following SQL statement after attaching to the database:

```
SQL> UPDATE RDB$FIELD_VERSIONS SET RDB$FIELD_SUB_TYPE = 32767
cont> WHERE RDB$FIELD_NAME = 'RDB$AREA_NAME';
```

5.2.5 Single Statement CALL Does Not Support Truncated Parameter List or DEFAULT Keyword

Oracle Rdb now allows the CALL statement in a compound statement to omit trailing IN mode parameters which have had a DEFAULT value defined in the procedure definition. Also supported is the DEFAULT keyword to replace an explicit value for the parameter.

However, the simple CALL statement (used outside a BEGIN END block) is not adaptable in this way and requires a full set of parameters and values. This is because a parameter signature is calculated for this type of CALL statement so that the parameter block passed by the calling routine and used by the called routine match exactly in parameter count and data types.

This is a permanent restriction for the simple CALL statement.

The following example shows that truncated parameter lists are fully supported by the compound use form of the CALL statement, but not by the simple CALL statement.

```
SQL> ATTACH 'FILENAME db$:scratch';
SQL> CREATE MODULE mmm
cont> PROCEDURE mmm_p (IN :a INTEGER DEFAULT 0, IN :b INTEGER DEFAULT 1);
cont> TRACE :a, :b;
cont> END MODULE;
SQL> SET FLAGS 'Trace';
SQL> CALL mmm_p (10,20);
~Xt: 10      20
SQL> CALL mmm_p (10);
%SQL-F-ARGCOUNT, Procedure MMM_P expected 2 parameters, was passed 1
SQL> call mmm_p ();
%SQL-F-ARGCOUNT, Procedure MMM_P expected 2 parameters, was passed 0
SQL> begin
cont> CALL mmm_p (10,20);
cont> CALL mmm_p (10);
cont> call mmm_p ();
cont> END;
~Xt: 10      20
~Xt: 10      1
~Xt: 0       1
```

For maximum flexibility, use the CALL statement inside a compound statement which supports truncated parameter lists, the DEFAULT keyword, and full value expressions for parameter arguments.

5.2.6 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.

5.2.7 Restriction for CREATE STORAGE MAP Statement on Table with Data

Oracle Rdb V7.0 added support that allows a storage map to be added to an existing table that contains data. The Oracle Rdb7 Guide to Database Design and Definition describes this feature and lists restrictions.

Oracle Rdb release 7.1 adds the restriction that the storage map cannot include a WITH LIMIT clause for the storage area. The following example shows the resulting error:

```
SQL> create table MAP_TEST1 (a integer, b char(10));
SQL> create index MAP_TEST1_INDEX on MAP_TEST1 (a);
SQL> insert into MAP_TEST1 (a, b) values (3, 'Third');
1 row inserted
SQL> create storage map MAP_TEST1_MAP for MAP_TEST1
cont> store using (a) in RDB$SYSTEM
cont> with limit of (10); -- cannot use WITH LIMIT clause
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-F-RELNEMPTY, table "MAP_TEST1" has data in it
-RDMS-E-NOCMPLXMAP, can not use complex map for non-empty table
```

5.2.8 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
```

.
.
.

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.

5.2.9 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.

5.3 Oracle RMU Known Problems and Restrictions

This section describes known problems and restrictions for the RMU interface for release 7.1.

5.3.1 RMU/CONVERT Fails to Correctly Define the RDB\$WORKLOAD Table

When a database is converted to Rdb 7.1 and the optional system table RDB\$WORKLOAD is present, Rdb fails to correctly define the metadata for this table, and SQL is unable to see the data type for the RDB\$NULL_FACTOR column.

The collection and utilization of workload data is unaffected by this problem. Only SQL applications are affected.

The following is an example of a database incorrectly converted from Rdb 7.0 to Rdb 7.1:

```
SQL> show table rdb$workload
Information for table RDB$WORKLOAD

Columns for table RDB$WORKLOAD:
Column Name                Data Type                Domain
-----
RDB$CREATED                DATE VMS
RDB$LAST_ALTERED           DATE VMS
RDB$DUPLICITY_FACTOR       BIGINT(7)
RDB$NULL_FACTOR            Data type: 0
RDB$RELATION_ID            INTEGER
RDB$FLAGS                  INTEGER
RDB$FIELD_GROUP            CHAR(31)
RDB$SECURITY_CLASS         CHAR(20)
```

The RDB\$NULL_FACTOR datatype is incorrectly interpreted. This will result in the following problem:

```
SQL> select rdb$null_factor from rdb$workload;
%SQL-F-FLDNOTCRS, Column RDB$NULL_FACTOR was not found in the tables in current
scope
```

A workaround for this problem is to have a sufficiently privileged user execute the following SQL command, commit, and then have applications that use this column DISCONNECT and reattach to the database.

```
SQL> update rdb$relation_fields set rdb$field_source='RDB$SCALED_COUNTER'
cont> where rdb$field_source='RDB$PROBABILITY';
```

This problem will be corrected in Oracle Rdb Release 7.1.0.2.

5.3.2 RMU Convert Fails When Maximum Relation ID is Exceeded

If, when relation IDs are assigned to new system tables during an RMU Convert of an Oracle Rdb V7.0 database to a V7.1 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 V70. 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 V7.1:

```
$rmu/convert mf_personnel
%RMU-I-RMUTXT_000, Executing RMU for Oracle Rdb V7.1-00
Are you satisfied with your backup of
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 and your backup of
  any associated .aij files [N]? Y
%RMU-I-LOGCONVRT, database root converted to current structure level
%RMU-F-RELMAXIDBAD, ROLLING BACK CONVERSION - Relation ID exceeds maximum
  8192 for system table RDB$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 V7.1-00
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 V7.1
%RMU-I-CVTCOMSUC, CONVERT committed for
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 to version V7.1
```

5.3.3 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.
- **BTREE**
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.

5.3.4 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.

5.3.5 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.

5.3.6 Default for RMU CRC Qualifier Changing in Future Release

The default behavior for the Crc qualifier for the following RMU commands is changing in a future release of Oracle Rdb:

- Backup
- Backup After_Journal
- Backup Plan
- Optimize After_Journal

Currently, the default value for the CRC qualifier is:

- Crc=Autodin_II is the default for NRZ/PE (800/1600 bits/inch) tape drives
- Crc=Checksum is the default for GCR (6250 bits/inch) tape drives and for TA78, TA79, and TA81 tape drives
- Nocrc is the default for TA90 (IBM 3480 class) drives

In a future release, the default value for the CRC qualifier will be Crc=Checksum for all tape drives except NRZ/PE (800/1600 bits/inch) tape drives. The default qualifier for the NRZ/PE (800/1600 bits/inch) tape drives will remain Crc=Autodin_II. The Crc=Checksum qualifier verifies the checksum on each buffer of data before it is written to tape or disk. This provides end-to-end error detection for the backup file I/O.

Oracle Corporation recommends that you accept the new behavior, that will be the default in a future release of Oracle Rdb, for your applications. The default behavior prevents you from including corrupt database pages in backup files and optimized .aij files. Without the checksum verifications, corrupt data pages in these files are not detected when the files are restored. The corruptions on the restored page may not be detected until weeks or months after the backup file is created, or it is possible the corruption may not be detected at all.

5.3.7 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.

5.3.8 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.

5.4 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.

5.4.1 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 minimum of 13 blocks and a maximum of 597 blocks. The actual increase depends mostly on the maximum number of users specified for the database.

5.4.2 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.

5.4.3 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.

5.4.4 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.

5.4.5 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 existing page size and the database has been manually opened or users are active, the add operation fails with the following error:

```

%RDB-F-SYS_REQUEST, error from system services request
-RDMS-F-FILACCERR, error opening database root DKA0:[RDB]TEST.RDB;1
-SYSTEM-W-ACCONFLICT, file access conflict

```

You can make this change only when no users are attached to the database and, if the database is set to OPEN IS MANUAL, the database is closed. Several internal Oracle Rdb data structures are based on the minimum page size and these structures cannot be resized if users are attached to the database.

Furthermore, because this particular change is not recorded in the AIJ, any recovery scenario fails. Note also that if you use .ajj files, you must backup the database and restart after-image journaling because this change invalidates the current AIJ recovery.

5.4.6 Support for Single-File Databases to Be Dropped in a Future Release

Oracle Rdb currently supports both single-file and multifile databases on all platforms. However, single-file databases will not be supported in a future release of Oracle Rdb. At that time, Oracle Rdb will provide the means to easily convert single-file databases to multifile databases.

Oracle Rdb recommends that users with single-file databases perform the following actions:

- ◆ Use the Oracle RMU commands, such as Backup and Restore, to make copies, backup, or move single-file databases. Do not use operating system commands to copy, back up, or move databases.
- ◆ Create new databases as multifile databases even though single-file databases are supported.

5.4.7 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.

5.4.8 Network Link Failure Does Not Allow DISCONNECT to Clean Up Transactions

If a program attaches to a database on a remote node and it loses the connection before the COMMIT statement is issued, there is nothing you can do except exit the program and start again.

The problem occurs when a program is connected to a remote database and updates the database, but then just before it commits, the network fails. When the commit executes, SQL shows, as it normally should, that the program has lost the link. Assume that the user waits for a minute or two, then tries the transaction again. The problem is that when the start transaction is issued for the second time, it fails because old information still exists about the previous failed transaction. This occurs even if the user issues a DISCONNECT statement (in V4.1 and earlier, a FINISH statement), which also fails with an RDB-E-IO_ERROR error message.

5.4.9 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.

5.5 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.

5.5.1 SQL Does Not Display Storage Map Definition After Cascading Delete of Storage Area

When you drop a storage area using the CASCADE keyword and that storage area is not the only area to which the storage map refers, the SHOW STORAGE MAP statement no longer shows the placement definition for that storage map.

The following example demonstrates this restriction:

```
SQL> SHOW STORAGE MAP DEGREES_MAP1
      DEGREES_MAP1
For Table:                DEGREES1
Compression is:          ENABLED
Partitioning is:         NOT UPDATABLE
Store clause:            STORE USING (EMPLOYEE_ID)
                        IN DEG_AREA WITH LIMIT OF ('00250')
                        OTHERWISE IN DEG_AREA2

SQL> DISCONNECT DEFAULT;
SQL> -- Drop the storage area, using the CASCADE keyword.
SQL> ALTER DATABASE FILENAME MF_PERSONNEL
cont> DROP STORAGE AREA DEG_AREA CASCADE;
SQL> -- Display the storage map definition.
SQL> ATTACH 'FILENAME MF_PERSONNEL';
SQL> SHOW STORAGE MAP DEGREES_MAP1
DEGREES_MAP1 For Table: DEGREES1
Compression is: ENABLED
Partitioning is: NOT UPDATABLE
```

The other storage area, DEG_AREA2, still exists, even though the SHOW STORAGE MAP statement does not display it.

A workaround is to use the RMU Extract command with the Items=Storage_Map qualifier to see the mapping.

5.5.2 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.)

5.5.3 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.

5.5.4 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 5–2 Elapsed Time for Index Creations

Index Create Job	Elapsed Time
Index1	00:02:22.50
Index2	00:01:57.94
Index3	00:02:06.27
Index4	00:01:34.53
Index5	00:01:51.96
Index6	00:01:27.57
Index7	00:02:34.64
Index8	00:01:40.56
Index9	00:01:34.43
Index10	00:01:47.44
All10	00:03:26.66

5.5.5 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>     lang 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.

5.5.6 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.

[Previous](#) | [Contents](#) | [Contents](#)