

Oracle® Rdb Server on OpenVMS

Release Notes

Release 7.4.1.4

June 2025

ORACLE®

Oracle Rdb Server Release Notes, Release 7.4.1.4 on OpenVMS Industry Standard 64 for Integrity Servers and OpenVMS Alpha

Copyright © 1984, 2025 Oracle and/or its affiliates. All rights reserved.

Oracle Corporation - Worldwide Headquarters, 2300 Oracle Way, Austin, TX 78741, United States

Primary Author: Rdb Engineering and Documentation group

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, Oracle Rdb, Hot Standby, LogMiner for Rdb, Oracle SQL/Services, Oracle CODASYL DBMS, Oracle RMU, Oracle CDD/Repository, Oracle Trace, and Rdb7 are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

Contents

Preface	ix
----------------------	----

1 Installing Oracle Rdb Server Release 7.4.1.4

1.1	Oracle Rdb on OpenVMS Industry Standard 64	1-1
1.2	Requirements	1-1
1.2.1	Ensure No Processes Have RDMSHRP Image Activated	1-2
1.3	Deprecated Features	1-2
1.3.1	Deprecated Support for the PL/I Compiler	1-3
1.4	Intel Itanium Processor 9700 “Kittson” Certified	1-3
1.5	Hardware Emulation Support	1-3
1.6	Maximum OpenVMS Version Check	1-3
1.7	Database Format Changed	1-3
1.8	Using Databases from Releases Earlier than V7.0	1-4
1.9	Changes to Installation Procedure	1-4
1.10	Invoking the VMSINSTAL Procedure	1-5
1.11	Stopping the Installation	1-6
1.12	After Installing Oracle Rdb	1-6
1.13	VMS\$MEM_RESIDENT_USER Rights Identifier Required	1-6

2 Software Errors Fixed in Oracle Rdb Release 7.4.1.4

2.1	Software Errors Fixed That Apply to All Interfaces	2-1
2.1.1	Unexpected Bugcheck Dump When ALTERNATE_OUTLINE_ID Flag is Enabled	2-1
2.1.2	Unexpected BUGCHECK During IMPORT DATABASE Statement	2-2
2.1.3	Unexpected BUGCHECK During Query On a View with Union	2-2
2.2	SQL Errors Fixed	2-3
2.3	RMU Errors Fixed	2-3
2.3.1	Unexpected Failure From RMU Extract When Using TRIM Functions	2-3
2.3.2	RMU Show Statistics Did Not Recognize Options Keywords Longer Than Four Characters	2-3
2.3.3	Unexpected BUGCHECK From RMU Show Statistics	2-4
2.3.4	RMU Recover Returns RMU-F-AIJNORCVR Even When Correct Journal Specified	2-4
2.3.5	RMU Backup Online Bugchecks When Logical Areas Created or Deleted	2-5
2.3.6	RMU Show Statistics Process Monitor Counters Incorrect on RAD Systems	2-6
2.3.7	Unexpected Bugcheck When Using RMU Show Statistics Configure Qualifier	2-6

3 Software Errors Fixed in Oracle Rdb Release 7.4.1.3

3.1	Software Errors Fixed That Apply to All Interfaces	3-1
3.1.1	Accessing Row Cache Fails with CANTMAPSHMEM, NOSUCHSEC Error	3-1
3.2	SQL Errors Fixed	3-1
3.2.1	CAST Function Converted Blank Strings to a Numeric Value 0	3-1
3.2.2	Unexpected SQL-F-BUGCHK When Using SAVEPOINT Statement	3-2
3.2.3	Hash Join Behavior Mismatch	3-2
3.3	RMU Errors Fixed	3-4
3.3.1	Invalid RMU Show Statistics File IO Overview Screen PgDis Field Values	3-4
3.3.2	RMU Verify Validate FLOAT did not Correctly Scan the Target Tables	3-4
3.3.3	Changed Processing of the RMU Load Exception_file Option	3-5
3.3.4	RMU Recover of Duplicated Database Deleted Journal in Original Database	3-6
3.3.5	Stale Entries Written to Statistics Lock Timeout Log	3-6
3.3.6	System Space Buffers not Freed Up After an RMU Close command	3-7
3.3.7	Unexpected Failure of RMU Convert Commit Command	3-8
3.3.8	Unexpected Reset of Logical Area Length by RMU Restore Initialize_tsns	3-9
3.4	RDO, RDBPRE and RDML Errors Fixed	3-10
3.4.1	Improved Error Detection for RDBPRE/BASIC Applications Using PACKED DECIMAL Types	3-10

4 Software Errors Fixed in Oracle Rdb Release 7.4.1.2

4.1	Software Errors Fixed That Apply to All Interfaces	4-1
4.1.1	On Reboot the RMONSTART74.COM Procedure Failed to Restart the RDBSERVER Service	4-1
4.1.2	Unexpected ACCVIO and BUGCHECK Dump for Some Queries on Integrity Systems	4-1
4.1.3	Inserting New Row Fails With RDMS-F-ABMCHNFUL Error	4-2
4.1.4	Error From Disk Write Results In Hang	4-4
4.1.5	Unexpected COSI-F-EXQUOTA Error Using Buffer Objects	4-4
4.2	SQL Errors Fixed	4-5
4.2.1	Incorrect Results From LISTAGG, FIRST_VALUE, and LAST_VALUE	4-5
4.2.2	Unexpected NO_PRIV Errors From DECLARE LOCAL TEMPORARY TABLE Statement	4-5
4.3	RMU Errors Fixed	4-6
4.3.1	Average Values Could be Zeroed on RMU Show Statistics Screens ...	4-6
4.3.2	RMU Recover Fails With RMU-F-CANTSYNCTSNS Using Correct Journal	4-7
4.3.3	SERVER_TRANSPORT Option to RMU Backup Parallel Not Working as Documented	4-8

5 Software Errors Fixed in Oracle Rdb Release 7.4.1.1

5.1	Software Errors Fixed That Apply to All Interfaces	5-1
5.1.1	Journal Backed Up While RMU Unload After Continuous Extracting Journal	5-1
5.1.2	Bugchecks at SEQ\$GET_SEQUENCE + 000001C0 After Incremental Restore	5-1
5.1.3	The SET FLAGS Option NOCOUNT_SCAN Erroneously Disabled Min and Max Optimized Index Lookups	5-2
5.2	SQL Errors Fixed	5-3
5.2.1	Unexpected NULL Returned from MOD and SQRT Functions	5-3
5.2.2	Unexpected RDMS-E-UNQCONFEVAL when Using ALTER CONSTRAINT Statement	5-3
5.2.3	Unexpected Bugcheck When Using More than One COUNT(*) in Query	5-4
5.2.4	Incorrect Handling of the DEFAULT Clause for AUTOMATIC UPDATE AS Columns	5-5
5.3	RMU Errors Fixed	5-5
5.3.1	Unexpected SYSTEM-F-ACCVIO when Using RMU Dump Audit for DACCESS	5-5
5.3.2	RMU SET AUDIT and RMU SHOW AUDIT Now Work on Remote Databases	5-6
5.3.3	RMU Backup After Space_check now Supports Disk Bound Volume Sets	5-7

6 Software Errors Fixed in Oracle Rdb Release 7.4.1.0

6.1	Software Errors Fixed That Apply to All Interfaces	6-1
6.1.1	RDMS-F-BADPAGNUM Exception After Snapshot File Truncated	6-1
6.1.2	SQL/Services V7.4 IVP Fails With RMU-E-NOSQSCLIENT	6-1
6.1.3	Long Running Delete Does Not Advance After-Image Journal Checkpoint	6-2
6.1.4	Long Running Query Blocks Other Queries	6-2
6.2	SQL Errors Fixed	6-3
6.2.1	Unexpected Bugcheck When Using DROP STORAGE AREA ... CASCADE	6-3
6.2.2	Unexpected Error %SQL-F-FLDNOTCRS For Declaring a LIST Cursor	6-3
6.3	RMU Errors Fixed	6-3
6.3.1	Clarification of RMU Unload After_Image Header	6-3
6.3.2	Unexpected RMU Unload After Bugcheck at AIJEXT\$QSORT_RELEASE_REC	6-4

7 Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.4

7.1	Enhancements	7-1
7.1.1	OPCOM Messages Now Always Entered in Oracle Rdb Monitor Log File	7-1
7.1.2	RMU Show Statistics Now Adds Lock ID to STALL_INVOKE Subprocess Parameters List	7-1
7.1.3	Enhanced Integrity Checking Before Page Writes	7-2
7.2	Obsolete Features	7-2

8 Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.3

8.1	Enhancements	8-1
8.1.1	The Lock_timeout Qualifier can now be Used With All RMU Reclaim Qualifiers	8-1
8.1.2	RMU Collect Optimizer_Statistics UNIQUE Index Performance Enhancement	8-2
8.1.3	New Partitions Qualifier for RMU Unload Command	8-3
8.1.4	New PARTITION_LIST Option for the Transaction_type Qualifier	8-4
8.1.5	Oracle Rdb Compilers Now Set the Object Module Severity	8-5
8.1.6	New ALTER DATABASE ... REBUILD SYSTEM INDEX Clause	8-6
8.1.7	New Implementation of the SQL IN Predicate	8-8
8.1.8	Updated Sample Code for RDB_CYPHER	8-9
8.2	Obsolete Features	8-12

9 Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.2

9.1	Enhancements	9-1
9.1.1	Substring Function Changes	9-1
9.1.2	Support for EMPTY_BLOB and EMPTY_CLOB in SQL	9-3
9.1.3	New ANY_VALUE Aggregate Function	9-4
9.1.4	Updated Diagnostics for the RMU REPAIR Command	9-5
9.1.5	Updated Diagnostics for the RMU MOVE_AREA Command	9-6
9.1.6	Hash Join Enhancements and Support	9-6
9.2	Obsolete Features	9-13

10 Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.1

10.1	Enhancements	10-1
10.1.1	Optional Builtin Function RDB\$SIS_ROW_FRAGMENTED	10-1
10.1.2	WITH READ ONLY Clause for CREATE and ALTER VIEW Statements	10-2
10.1.3	Enhanced Conversion of Date/time String Literals	10-3
10.1.4	Support for OVERRIDING Clause in INSERT and REPLACE Statements	10-3
10.1.5	RMU Recover Progress_report Qualifier and Ctrl-T Display	10-5
10.1.6	Additional Options for the RMU Set Database Command	10-8
10.1.7	New Summary_only Qualifier to RMU Dump Audit Command	10-9
10.1.8	New RMU Verify Minimize_conflicts Qualifier	10-9
10.1.9	New OPTION=GENERATED Added to RMU Extract Command	10-10
10.1.10	Changed Behavior for the Noedit_filename Qualifier in RMU Backup After_Journal Command	10-11
10.2	Obsolete Features	10-12
10.2.1	Comma Statement Separator in Trigger Body No Longer Supported	10-12

11 Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.0

11.1	Enhancements	11-1
11.1.1	PCSI Support for Rdb Kit Installation and Deinstallation	11-1
11.1.2	Some Aggregate Functions Inherit Source Column EDIT STRING	11-2
11.1.3	Enhanced LIKE Table Support in CREATE TABLE Statement	11-3

11.1.4	RMU Reclaim Free_pages Qualifier Frees Unused Data Page Clumps	11-6
11.1.5	CREATE DEFAULT AUDIT Command Supports OR REPLACE Clause	11-11
11.1.6	System Privileges Feature	11-12
11.1.7	Database Vault Feature	11-22
11.1.8	SET FLAGS Keyword for Hash Join Feature - HASHING	11-23
11.1.9	JOIN BY HASH Clause in CREATE OUTLINE Statement	11-24
11.1.10	Hash Join Feature	11-25
11.1.11	New ALTER DATABASE ... LOAD ACL IDENTIFIERS Clause	11-27
11.1.12	ALTER TABLE Statement Allows Change to READ ONLY Table	11-29
11.1.13	NULLS FIRST and NULLS LAST Options for ORDER BY Clause ..	11-30
11.1.14	Enhancements to RMU Unload After_Image (LogMiner) Interface ...	11-31
11.1.14.1	XML Option to FORMAT Qualifier	11-31
11.1.14.2	TRIM Option	11-32
11.1.14.3	SYMBOLS Qualifier	11-32
11.1.15	Named Partition Support for RESERVING Clause	11-33
11.2	Obsolete Features	11-34
11.2.1	RMU Backup No Longer Supports HUFFMAN or LZSS Compression, Use ZLIB Instead	11-34

12 Documentation Corrections, Additions and Changes

12.1	Documentation Corrections	12-2
12.1.1	Oracle Rdb Release 7.4 New Features Manual	12-2
12.1.2	Oracle Rdb Timer Queue Limit Increased	12-2
12.1.3	Undocumented /TRANSACTION=EXCLUSIVE Option for RMU Populate_Cache	12-2
12.1.4	Action of DECLARE TRANSACTION When SQL\$DATABASE is Defined	12-3
12.1.5	RDB\$USAGE Field Values	12-3
12.1.6	Updated Documentation for RDMS\$BIND_CODE_OPTIMIZATION Logical	12-5
12.1.7	New RDMS\$BIND_AIJ_BUFFER_POOL_COUNT Logical	12-6
12.1.8	RDMSTT Image Optionally Installed	12-7
12.1.9	Some Optional System Tables Can Be Relocated Using a User Defined Storage Map	12-7
12.1.10	Recovering an Oracle Rdb Database After RMU/RESTORE/ONLY_ROOT	12-9
12.1.11	Oracle Rdb Position on NFS Devices	12-10
12.1.12	RDMS\$BIND_STAREA_EMERGENCY_DIR Logical Name	12-12
12.1.13	RDMS-F-FULLAIJBKUP, Partially-Journalled Changes Made	12-13
12.1.14	Undocumented Hot Standby Logical Names	12-15
12.1.15	Corrections to the EDIT STRING Documentation	12-16
12.1.16	Changes and Improvements to the Rdb Optimizer and Query Compiler	12-17
12.1.17	Sorting Capabilities in Oracle Rdb	12-20
12.1.18	RDMS\$BIND_MAX_DBR_COUNT Documentation Clarification	12-22
12.1.19	Missing Tables Descriptions for the RDBEXPERT Collection Class ...	12-22
12.1.20	Missing Columns Descriptions for Tables in the Formatted Database	12-23
12.2	Optimizer Changes	12-31
12.2.1	New Optimizations for the SQL IN Predicate	12-31

13 Known Problems and Restrictions

13.1	Known Problems and Restrictions in All Interfaces	13-1
13.1.1	OpenVMS TCPIP Issue May Cause Rdb Kit Installation to Fail	13-1
13.1.2	The Format of the RMU Analyze Binary_output Files Can Change	13-2
13.1.3	RMU Verify Key_values May Fail on Some Indices	13-3
13.1.4	REPLACE Statement Fails With Primary Key Constraint Failure When Used on a View	13-3
13.1.5	Possible Incorrect Results When Using Partitioned Descending Indexes	13-4
13.1.6	Application and Oracle Rdb Both Using SYS\$HIBER	13-6
13.1.7	Unexpected RCS Termination	13-7
13.1.8	External Routine Images Linked with PTHREAD\$RTL	13-7
13.1.9	ILINK-E-INVOCRINI Error on IA-64	13-8
13.1.10	SYSTEM-F-INSMEM Fatal Error With SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED in Galaxy Environment	13-8
13.1.11	Oracle Rdb and OpenVMS ODS-5 Volumes	13-9
13.1.12	Optimization of Check Constraints	13-9
13.1.13	Carryover Locks and NOWAIT Transaction Clarification	13-11
13.1.14	Unexpected Results Occur During Read-Only Transactions on a Hot Standby Database	13-12
13.1.15	Row Cache Not Allowed While Hot Standby Replication is Active	13-12
13.1.16	Control of Sort Work Memory Allocation	13-13
13.1.17	The Halloween Problem	13-13
13.1.18	Single Statement LOCK TABLE is Not Supported for SQL Module Language and SQL Precompiler	13-15
13.1.19	Multistatement or Stored Procedures May Cause Hangs	13-15
13.1.20	Use of Oracle Rdb from Shareable Images	13-16
13.1.21	RMU Convert Fails When Maximum Relation ID is Exceeded	13-17
13.1.22	RMU Unload After_journal Requires Accurate AIP Logical Area Information	13-17
13.1.23	Do Not Use HYPERSORT with RMU Optimize After_journal Command	13-18
13.1.24	RMU Backup Operations Should Use Only One Type of Tape Drive	13-19
13.1.25	RMU Verify Reports PGSPAMENT or PGSPMCLST Errors	13-19
13.1.26	Converting Single-File Databases	13-20
13.1.27	Row Caches and Exclusive Access	13-21
13.1.28	Exclusive Access Transactions May Deadlock with RCS Process	13-21
13.1.29	Strict Partitioning May Scan Extra Partitions	13-21
13.1.30	Restriction When Adding Storage Areas with Users Attached to Database	13-22
13.1.31	Multiblock Page Writes May Require Restore Operation	13-22
13.1.32	Replication Option Copy Processes Do Not Process Database Pages Ahead of an Application	13-23
13.1.33	Different Methods of Limiting Returned Rows from Queries	13-23
13.1.34	Suggestions for Optimal Use of SHARED DATA DEFINITION Clause for Parallel Index Creation	13-24
13.1.35	Side Effect When Calling Stored Routines	13-26
13.1.36	Considerations When Using Holdable Cursors	13-27

Examples

9-1	Example 1: Using OPTIMIZE USING	9-10
9-2	Example 2: Using the OPTIMIZE FOR HASH JOIN clause	9-12

Tables

11-1	System Privileges	11-18
12-1	Rdb\$USAGE Field Values	12-4
12-2	Optional System Tables and Their Storage Map Names	12-8
12-3	Hot Standby Logical Names	12-15
12-4	Columns for Table EPC\$1_221_TRANS_TPB	12-23
12-5	Columns for Table EPC\$1_221_TRANS_TPB_ST	12-23
12-6	Columns for Table EPC\$1_221_DATABASE	12-23
12-7	Columns for Table EPC\$1_221_REQUEST_ACTUAL	12-24
12-8	Columns for Table EPC\$1_221_TRANSACTION	12-27
12-9	Columns for Table EPC\$1_221_REQUEST_BLR	12-31
13-1	Sort Memory Logicals	13-13
13-2	Elapsed Time for Index Creations	13-26

Preface

Purpose of This Manual

This manual contains release notes for Oracle Rdb Server Release 7.4.1.4. 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 Server users. Read this manual before you install, upgrade, or use Oracle Rdb Server Release 7.4.1.4.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Reader Comments

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most?

If you find any errors or have any other suggestions for improvement, please indicate the document title, release date, chapter, section, and page number (if available).

Please direct all comments, and corrections to this email address:
infordb_mb@oracle.com.

If you have problems with the software, please contact your local Oracle Support Services.

Online Document Format

You can view the Adobe Acrobat formatted documentation using the Acrobat Reader, which allows anyone to view, navigate, and print documents in the Adobe Portable Document Format (PDF). See <http://www.adobe.com> for information about obtaining a free copy of Acrobat Reader and for information on supported platforms.

The Oracle Rdb documentation in Adobe Acrobat format can be reached from the Oracle Rdb main web page. At the bottom of the page are links to documentation for all of the Rdb Releases. Following is the link for the Oracle Rdb main page:

<https://www.oracle.com/database/technologies/related/rdb.html>

Optional Software

The following optional products may be referenced in these release notes:

- Oracle CDD/Repository
- Oracle Trace
- Oracle Replication Option for Rdb
- Oracle SQL/Services and OCI Services for Oracle Rdb
- Oracle SQL/Services Client API Kits
- Oracle ODBC Driver for Rdb
- Oracle JDBC for Rdb
- Oracle Rdb Developer Tools for Visual Studio
- Oracle Rdb Extension for SQL Developer
- Oracle Rdb Connectivity Manager

References to Products

The Oracle Rdb documentation set to which this manual belongs often refers to the following products by their abbreviated names:

- Oracle Rdb refers to Oracle Rdb Server.
- OpenVMS IA-64 refers to the HPE OpenVMS Industry Standard 64 for Integrity Servers.
- OpenVMS refers to the OpenVMS Alpha and OpenVMS IA-64 operating systems.
- Oracle Rdb refers to Oracle Rdb on OpenVMS Alpha, or OpenVMS IA-64 software.
- The SQL interface to Oracle Rdb is referred to as SQL.
- In Oracle Rdb documentation, the terms release and version (and their abbreviations) are sometimes used interchangeably.
- Oracle CDD/Repository software is referred to as the dictionary, the data dictionary, or the repository.

- Oracle ODBC Driver for Rdb software is referred to as the ODBC driver.
- Oracle Trace on OpenVMS software is referred to as Oracle Trace.
- Oracle Rdb Connectivity Manager is referred to as ORCM.
- Oracle JDBC for Rdb is referred to as Rdb JDBC.
- Hewlett-Packard Enterprise is referred to as HPE.
- VMS Software Inc is referred to as VSI.
- DECnet and DECnet-Plus refer respectively to HPE DECnet for OpenVMS and HPE DECnet-Plus for OpenVMS. DECnet Phase IV is used interchangeably with DECnet in this document.

Document Structure

This manual consists of the following chapters:

Chapter 1	Describes how to install Oracle Rdb Release 7.4.1.4.
Chapter 2	Describes problems corrected in Oracle Rdb Release 7.4.1.4.
Chapter 3	Describes problems corrected in Oracle Rdb Release 7.4.1.3.
Chapter 4	Describes problems corrected in Oracle Rdb Release 7.4.1.2.
Chapter 5	Describes problems corrected in Oracle Rdb Release 7.4.1.1.
Chapter 6	Describes problems corrected in Oracle Rdb Release 7.4.1.0.
Chapter 7	Describes enhancements introduced in Oracle Rdb Release 7.4.1.4.
Chapter 8	Describes enhancements introduced in Oracle Rdb Release 7.4.1.3.
Chapter 9	Describes enhancements introduced in Oracle Rdb Release 7.4.1.2.
Chapter 10	Describes enhancements introduced in Oracle Rdb Release 7.4.1.1.
Chapter 11	Describes enhancements introduced in Oracle Rdb Release 7.4.1.0.
Chapter 12	Provides information not currently available in the Oracle Rdb documentation set.
Chapter 13	Describes problems, restrictions, and workarounds known to exist in Oracle Rdb Release 7.4.1.4

Installing Oracle Rdb Server Release 7.4.1.4

This software update is installed using the OpenVMS VMSINSTAL utility. Please refer to the *OpenVMS System Manager's Manual, Volume 1: Essentials; Preparing Your System to Run VMSINSTAL.COM* for details.

NOTE

Oracle Rdb Release 7.4 kits are full kits. There is no requirement to install any prior release of Oracle Rdb when installing new Rdb Release 7.4 kits.

1.1 Oracle Rdb on OpenVMS Industry Standard 64

The Oracle Rdb product family is available on the OpenVMS Industry Standard 64 platform and the OpenVMS AlphaServer platform. In general, the functionality for one platform is available on the other platform. However, certain differences between the platforms may result in minor capability and functionality differences.

The database format for Oracle Rdb Release 7.4 is the same on both IA-64 and Alpha platforms and databases may be accessed simultaneously from both architectures in a cluster environment. Access to an Oracle Rdb Release 7.4 database from prior Rdb versions (on Alpha, IA-64 or VAX platforms) or from other systems on the network is available via the Oracle Rdb remote database server.

1.2 Requirements

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

- This Oracle Rdb release requires the following OpenVMS environments:
 - OpenVMS Alpha version 8.4 or later.
 - OpenVMS Industry Standard 64 version 8.4 or later.
- Some hardware configurations require the installation of OpenVMS V8.4-1H1 or later from VMS Software Inc. (VSI).
- Oracle strongly recommends that all available OpenVMS patches are installed on all systems prior to installing Oracle Rdb. Contact your HPE or VSI support representative for more information and assistance.
- Oracle Rdb must be shutdown before you install this update kit. That is, the command file SYSSMANAGER:RMONSTOP74.COM should be executed before proceeding with this installation. If you have an OpenVMS cluster, you must shutdown the Rdb Release 7.4 monitor on all nodes in the cluster before proceeding.

- After executing RMONSTOP74.COM, no process on any system in the cluster should have any existing RDMSHRP74.EXE image activated. See Section 1.2.1 for additional information.
- The installation requires approximately 400,000 blocks for OpenVMS Alpha systems.
- The installation requires approximately 600,000 blocks for OpenVMS IA-64 systems.

1.2.1 Ensure No Processes Have RDMSHRP Image Activated

The Oracle Rdb installation procedure checks to make sure that the Oracle Rdb Monitor (RDMMON) process is not running. However, it is also important to make sure that there are no processes on the cluster that share the system disk that have image activated a prior Rdb 7.4 version RDMSHRP image. Such processes may not be currently attached to a database but may do so in the future and could cause problems by using an older RDMSHRP image with a later Rdb installation.

The following command procedure (referenced as FIND_RDMSHRP74_PROC.COM below) can be used on each cluster node that shares the system device to determine if there are any processes that have activated the RDMSHRP74.EXE image. This procedure should be executed by a privileged account after RMONSTOP74 has been run. Any processes that have RDMSHRP74.EXE activated at this point should be terminated prior to starting the Rdb installation procedure.

```
$ RDB$TMP = "SYS$SCRATCH:RDB$TMP" + F$GETJPI(0,"MASTER_PID") + ".TMP"
$ DEFINE /NOLOG /USER RDB$TMP 'RDB$TMP'
$ ANALYZE /SYSTEM
    SET OUTPUT RDB$TMP
    SHOW PROCESS /CHANNELS ALL
    EXIT
$ SEARCH /OUTPUT='RDB$TMP' 'RDB$TMP';-1 RDMSHRP74.EXE,"PID:"
$ SEARCH 'RDB$TMP' RDMSHRP74.EXE /WINDOW=(1,0)
$ DELETE /NOLOG 'RDB$TMP';*
```

In the following example, the process 25E5A8E9 named "FOO\$SERVER" has the image RDMSHRP74.EXE activated even after RMONSTOP74.COM has been executed and this process is terminated prior to starting the Rdb installation procedure:

```
$ @SYS$MANAGER:RMONSTOP74.COM
.
.
.
$ @FIND_RDMSHRP74_PROC.COM
OpenVMS system analyzer
Process index: 00E9   Name: FOO$SERVER   Extended PID: 25E5A8E9
0290 7FEA0500 8A5A50C0 $1$DGA1:[VMS$COMMON.SYSLIB]RDMSHRP74.EXE;6(section file)
$ STOP/IDENTIFICATION=25E5A8E9
```

1.3 Deprecated Features

1.3.1 Deprecated Support for the PL/I Compiler

Oracle no longer supports the PL/I compiler on OpenVMS Alpha systems. Any existing support (SQL\$PRE /PLI qualifier, and LANGUAGE PLI in the SQL Module language) will be provided "AS IS". However, Oracle Rdb no longer plans to enhance or test this functionality.

1.4 Intel Itanium Processor 9700 “Kittson” Certified

For this release of Oracle Rdb on HPE Integrity servers, the Intel Itanium Processor 9700 series, code named “Kittson”, is the newest processor for which Rdb is certified. Please note that OpenVMS V8.4-2L1 or later is required for this class of processors.

1.5 Hardware Emulation Support

Stromasys and AVTware offer Alpha emulator products that prolong the use of OpenVMS Alpha applications. These products include Charon-AXP from Stromasys and vtAlpha from AVTware. The Charon-AXP and vtAlpha products emulate complete Alpha systems allowing OpenVMS applications, layered products, tools, and middleware to run unmodified. For more details on these options, see our web page about these platforms:

<https://www.oracle.com/database/technologies/related/alpha-vax-emulators.html>

1.6 Maximum OpenVMS Version Check

OpenVMS Version 8.4-x is the maximum supported version of OpenVMS for this release of Oracle Rdb.

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

1.7 Database Format Changed

The Oracle Rdb on-disk database format is 740 as shown in the following example.

```
$ RMU/DUMP/HEADER databasename
...
  Oracle Rdb structure level is 74.0
...
```

An RMU/CONVERT operation is required for databases created by or accessed by Oracle Rdb V7.0, V7.1, V7.2 or V7.3 to be accessed with Rdb Release 7.4.

Prior to upgrading to Oracle Rdb Release 7.4 and prior to converting an existing database to Oracle Rdb Release 7.4 format, Oracle strongly recommends that you perform a full database verification (with the RMU /VERIFY /ALL command) along with a full database backup (with the RMU /BACKUP command) to ensure a valid and protected database copy.

1.8 Using Databases from Releases Earlier than V7.0

The RMU Convert command for Oracle Rdb V7.4 supports conversions from Rdb V7.0, V7.1, V7.2 and V7.3 format databases only. This restriction also applies to the RMU Restore command which implicitly uses RMU Convert during the restore operation.

```
$ RMU/RESTORE/NOCD/DIRECTORY=SYS$DISK:[] [-]TESTING.RBF
%RMU-I-AIJRSTAVL, 0 after-image journals available for use
%RMU-I-AIJISOFF, after-image journaling has been disabled
%RMU-I-LOGCONVRT, database root converted to current structure level
%RMU-S-CVTDBSUC, database DEV1:[DOCUMENTATION.T]TESTING.RDB;1 successfully
converted from version V7.0 to V7.4
%RMU-I-CVTCOMSUC, CONVERT committed for DEV1:[DOCUMENTATION.T]TESTING.RDB;1
to version V7.4
%RMU-W-USERECCOM, Use the RMU Recover command. The journals are not available.
```

If you have an Oracle Rdb V3.0 through V6.1 format database or database backup, you must convert it to at least Oracle Rdb V7.0 format and then convert it to Oracle Rdb V7.4 format.

For example, if you have a V4.2 format database, you must convert it first to at least Oracle Rdb V7.0 format, then convert it to Oracle Rdb V7.4 format. This might be achieved by using Rdb V6.1 in an OpenVMS Alpha multiversion environment to convert from V4.2 to V6.1. Then installing Rdb V7.1 in that multiversion environment to convert V6.1 to V7.1. Then installing Rdb V7.4 in that multiversion environment to convert V7.1 to V7.4.

Note

This multi-step conversion will require you to use RMU/CONVERT/COMMIT so rolling back to Rdb V4.2 from V7.4 will not be possible. Oracle recommends that full database backups be performed prior to starting conversions to later versions.

Alternately, use the SQL EXPORT DATABASE statement to save the database structure and recreate it using the SQL IMPORT DATABASE statement. This allows a single step migration but will take more I/O and time because the database is completely rebuilt.

If you attempt to convert or restore a database format that is prior to Oracle Rdb V7.0 format directly to Oracle Rdb V7.4 format, Oracle RMU generates an error.

```
$ rmu/convert scratch/noconfirm
%RMU-I-RMUTXT 000, Executing RMU for Oracle Rdb V7.4-140 on OpenVMS Alpha V8.4
%RMU-F-CVRTUNS, The minimum database version that can be converted is version 70.
%RMU-F-FTL_CNV, Fatal error for CONVERT operation at 10-JUN-2023 15:48:04.61
```

1.9 Changes to Installation Procedure

Bug 30622247

Starting with Oracle Rdb V7.4.1 the installation procedure was enhanced to support mixed case passwords and the VSI Enhanced Password Management software.

These changes affect the installation of Oracle Rdb when it is installed on a system that does not have one or more of the required Rdb accounts. These accounts are RDB\$REMOTE74 (support account for remote access), RDMALJ74 (support account for Hot Standby and the Log Shipping Services) and RDMSTT74 (support account for cluster-wide RMU Show Statistics).

Typically, when upgrading the Rdb installation to a point release, these accounts will already exist and in that case will not be re-created.

The following changes have been made to the Oracle Rdb installation procedures.

- If the installer account has the UAF flag PWDMIX set, then Oracle Rdb no longer forces the generated password (provided by OpenVMS VMSINSTAL.COM) to upper case. This means that the created accounts (RDB\$REMOTE74, RDMALJ74 and RDMSTT74) will also have the UAF flag PWDMIX set and use mixed case passwords.

If the VSI (VMS Software, Inc.) Enhanced Password Management software is installed, then it is likely that other (numeric, mixed case, and special) characters will also be generated. Oracle Rdb now supports the use of such enhanced passwords. However, if the generated password contains either single (') or double (") quotes or exclamation point (!), then it will be rejected and Rdb will request an alternate password.

Note

A revised version of the OpenVMS VMSINSTAL.COM procedure (located in SYSS\$UPDATE) is required to compliment the changes made to the Oracle Rdb installation procedure. Please ensure that the VSI OpenVMS UTILITY V1.0 patch kit has been installed prior to installing Oracle Rdb.

- If the installer account does not have the UAF flag PWDMIX set, then the password will be upper case as in previous releases.

Please refer to VMS Software, Inc. for details and availability of the VSI Enhanced Password Management software for VSI OpenVMS Integrity Version V8.4-2L1 and V8.4-2L3, and VSI OpenVMS Alpha Versions V8.4-2L2 and V8.4-2L1.

1.10 Invoking the VMSINSTAL Procedure

The installation procedure is the same for Oracle Rdb Server on OpenVMS Alpha and Oracle Rdb Server on OpenVMS IA-64.

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

- To install the Oracle Rdb Server on OpenVMS IA-64 kit

```
@SYSS$UPDATE:VMSINSTAL RDBV74140IM074 device-name
```

- To install the Oracle Rdb Server on OpenVMS Alpha kit

```
@SYSS$UPDATE:VMSINSTAL RDBV74140AM074 device-name
```

device-name

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

```
DKA400:[RDB.KIT]
```

1.11 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.

1.12 After Installing Oracle Rdb

This update provides a new Oracle TRACE facility definition for Oracle Rdb. Any Oracle TRACE selections that reference Oracle Rdb will need to be redefined to reflect the new facility version number for the updated Oracle Rdb facility definition, "RDBVMSV7.4".

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=RDBVMSV7.4 -  
_ $ /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.

1.13 VMS\$MEM_RESIDENT_USER Rights Identifier Required

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

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

Software Errors Fixed in Oracle Rdb Release 7.4.1.4

This chapter describes software errors that are fixed by Oracle Rdb Release 7.4.1.4.

2.1 Software Errors Fixed That Apply to All Interfaces

2.1.1 Unexpected Bugcheck Dump When ALTERNATE_OUTLINE_ID Flag is Enabled

It was reported that using the new native SUBSTR (introduced in Rdb 7.4.1.2) and also enabling the flag ALTERNATE_OUTLINE_ID would lead to a BUGCHKDMP error and a failed query.

The following example shows the problem.

```
SQL> set flags 'alternate_outline_id(literals)';
SQL>
SQL> select
cont>     EMPLOYEE_ID,
cont>     substr(DEPARTMENT_CODE, 1, 1),
cont>     any_value (SUPERVISOR_ID)
cont> from JOB_HISTORY
cont> group by
cont>     EMPLOYEE_ID,
cont>     substr(DEPARTMENT_CODE, 1, 1),
cont>     SUPERVISOR_ID
cont> ;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file USER2:[TESTING]RDSBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
virtual address=0000000000A84000, PC=FFFFFFFF80531990, PS=0000001B
SQL>
```

This problem can also occur when using the SUBSTRING ANSI/ISO SQL function when the dialect is set to ORACLE LEVEL3.

This problem can be avoided by not enabling the flag ALTERNATE_OUTLINE_ID(LITERALS), or by coding the query to use SUBSTRING (... FROM ... FOR ...) in a dialect other than ORACLE LEVEL3.

This problem has been corrected in Oracle Rdb Release V7.4.1.4.

2.1.2 Unexpected BUGCHECK During IMPORT DATABASE Statement

Bug 36044573

In prior releases of Oracle Rdb it was possible on rare occasions for a query to generate a BUGCHECK dump during processing. It was reported when an existing stored procedure was being recreated by an IMPORT DATABASE statement, but may occur from an ALTER MODULE or CREATE MODULE statement when defining a stored procedure.

The following example shows the bugcheck footprint

```
***** Exception at 80C816F1 : RDMSHRP74\rdms_exe_interp
Module EXEINTERP + 00046AB1; line 92482
%SYSTEM-F-ACCVIO, access violation, reason mask=04,
virtual address=0000000000000006, PC=FFFFFFFF80C816F1, PS=0000000B
Saved PC = 817967D0 : RDMSHRP74\RDMS$$SETUP_RSS$NDX_RL_MERGE
Module RDMS$PREEXEDYN + 00007BF0; line 10320
Saved PC = 815D75E0 : RDMSHRP74\RDMS$$SETUP_RSS$NDX
Module RDMS$PREEXE + 00029370; line 18490
Saved PC = 81792490 : RDMSHRP74\RDMS$$CREATE_RLEAF
Module RDMS$PREEXEDYN + 000038B0; line 9009
Saved PC = 815D86D0 : RDMSHRP74\RDMS$$SETUP_RSS$NDX
Module RDMS$PREEXE + 0002A460; line 18899
Saved PC = 81591220 : RDMSHRP74\RDMS$$CREATE_RETRIEVAL
Module RDMS$PREEXE_CREATE + 0000F8B0; line 8425
```

This problem occurred when a constant expression calling a stored function used by queries within the stored procedure were evaluated during the optimization step. This should not have been attempted during the CREATE or ALTER MODULE statements.

This problem has been corrected in Oracle Rdb Release V7.4.1.4.

2.1.3 Unexpected BUGCHECK During Query On a View with Union

Bug 36044573

In prior releases of Oracle Rdb it was possible, on rare occasions, for a query to generate a BUGCHECK dump during processing.

This problem occurred when query was performed on a view and caused a zig-zag strategy to be applied to one leg of a UNION statement. The zig-zag code was not correctly using the join logical column(s) from the query.

The following example shows the bugcheck footprint

```
***** Exception at 80F6CA20 : RDMSHRP74\FIND_SNDX_SEGMENT_DSC
Module RDMS$PREEXE_CREATE + 00026860; line 15208
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 80F6BEF0 : RDMSHRP74\GEN_UNMAKE_ZIGZAG_IKEY
Module RDMS$PREEXE_CREATE + 00025D30; line 15035
Saved PC = 80F60400 : RDMSHRP74\RDMS$$CREATE_RSS$MTCH
Module RDMS$PREEXE_CREATE + 0001A240; line 11840
Saved PC = 80F09A00 : RDMSHRP74\CSOL_TO_RSS
Module RDMS$PREEXE + 00004250; line 6672
.
.
.
```

This problem has been corrected in Oracle Rdb Release V7.4.1.4.

2.2 SQL Errors Fixed

None for this release.

2.3 RMU Errors Fixed

2.3.1 Unexpected Failure From RMU Extract When Using TRIM Functions

Bug 37555173

Oracle Rdb Server release V7.3.3 introduced enhancements to the TRIM builtin function, and also new related functions BTRIM, LTRIM, and RTRIM. View, triggers, constraints, and so on created using these new builtin functions may cause the RMU Extract command to fail with an error similar to:

```
%RMU-F-BLRINV, internal error - BLR string 83 for . is invalid
-RDMS-E-BAD_CODE, corruption in the query string
%RMU-F-FTL_RMU, Fatal error for RMU operation at 28-JAN-2025
```

This problem is caused by references to these TRIM functions within a CASE, DECODE, NULLIF, or COALESCE function.

This problem has been corrected in Oracle Rdb Release V7.4.1.4.

2.3.2 RMU Show Statistics Did Not Recognize Options Keywords Longer Than Four Characters

Bug 37482605

The RMU Show Statistics utility would ignore any option specified in the /OPTIONS qualifier if the option required more than four characters to be considered unique. For example, even though /OPTIONS=LOG_STALL_ALARM is specified in the following command to prevent short duration stalls from being logged, they were logged anyway.

```
$ RMU/SHOW STATISTICS -
  /STALL_LOG=STALLS.LOG -
  /ALARM=300 -
  /OPTION=LOG_STALL_ALARM -
SQL$DATABASE
```

The following options would be ignored by RMU Show Statistics.

- LOG_STALL_LOCK
- NOLOG_STALL_LOCK
- LOG_STALL_ALARM
- NOLOG_STALL_ALARM

Any option that requires more than four characters is now properly parsed.

This problem has been corrected in Oracle Rdb Release V7.4.1.4.

2.3.3 Unexpected BUGCHECK From RMU Show Statistics

It was reported that RMU Show Statistics would bugcheck when displaying some Row Cache statistics. This problem only occurs in Oracle Rdb Server release V7.4.1.3.

The screens affected are:

- Row Cache Computations
- Row Cache Overview

The following example shows the problem.

```
Node: NODEA1 (1/1/1) Oracle Rdb V7.4-130 Perf. Monitor    19-NOV-2024 10:17:24.77
Rate: 3.00 Seconds Row Cache Overview (Unsorted)          Elapsed: 7 00:11:04.68
Page: 1 of 1      USERS3:[TESTING.DATABASE]ADMINDB.RDB;1  Mode: Online
-----
Cache.Name..... #Searches Hit% Full% #Inserts #Wrap #Slots Len
.
.
.
-----
Config Exit Help Menu >next_page <prev_page ]next_obj [prev_obj Options Pause Re
%COSI-F-BUGCHECK, internal consistency failure
%RMU-F-FATALOSI, Fatal error from the Operating System Interface.
%RMU-I-BUGCHKDMP, generating bugcheck dump file USERS2:[BUGCHECK]RMUBUGCHK.DMP;
```

This problem has been corrected in Oracle Rdb Release V7.4.1.4.

2.3.4 RMU Recover Returns RMU-F-AIJNORCVR Even When Correct Journal Specified

Bug 18028302

When journals were backed up using the Noquiet_point qualifier, it was possible for the RMU Recover operation to request an older journal than necessary when a recovery operation was initiated.

In the following example, the RMU Restore operation requests journal sequence number 3 as the first journal needed for the RMU Recover command. However, RMU Recover gives an error when that journal is specified and requests an older journal, sequence number 2.

```
$ RMU/RESTORE/NOLOG/NOCD/NORECOVER TEST
%RMU-I-AIJRECFUL, Recovery of the entire database starts with AIJ file sequence 3
...
$ RMU/RECOVER/LOG TEST_BACKUP_3
%RMU-I-LOGRECDB, recovering database file DISK:[DIRECTORY]TEST.RDB;1
%RMU-F-AIJNORCVR, recovery must start with journal sequence 2
```

Retrying the operation with journal sequence number 2 succeeds, although no transactions from that journal are actually applied to the restored database.

RMU Recover was using incorrect criteria when determining which journal was needed to begin recovery. The journal requested by RMU Restore was the correct journal. RMU Recover now uses the journal specified by RMU Restore when determining the journal needed to begin recovery of the database.

This problem has been corrected in Oracle Rdb release V7.4.1.4.

2.3.5 RMU Backup Online Bugchecks When Logical Areas Created or Deleted

Bugs 8420535, 2296903, 17029010

When logical areas (tables, indices, etc.) were being created or deleted while an online backup was active, RMU could bugcheck with the following exception.

```
***** Exception at 80CA53C0 : RMU74\PIOFETCH$WITHIN_DB + 00000CD0
%RMU-F-CANTREADDBS, error reading pages 10:4294967295-4294967295
-RMU-F-BADPAGNUM, page 4294967295 is out of valid range (1:100) for physical area 10
```

The above failure can be demonstrated with the following simple test.

In one session, execute the following SQL statement.

```
SQL> attach 'filename MF_PERSONNEL';
SQL> drop table DEGREES cascade;
```

While the SQL session still has an active transaction, from another session, attempt to backup the database.

```
$ RMU/BACKUP/ONLINE/NOQUIET MF_PERSONNEL MF_PERSONNEL
%RMU-I-BUGCHKDMP, generating bugcheck dump file DEV:[DIR]RMUBUGCHK.DMP;
```

This problem is due to how RMU Backup Online reconstructs a page from the snapshot backing storage. RMU uses information from the Area Inventory Pages (AIP), which defines how logical areas are represented in the database. RMU reads the AIP pages when the backup starts, and relies on the information to not change for the duration of the backup operation. If the AIP changes, RMU may fail when attempting to build a consistent page for the backup.

The documentation for the RMU Backup command does suggest that logical area operations, such as CREATE TABLE or CREATE INDEX, not be done while a backup is active, but RMU did not enforce that restriction.

To avoid this problem refrain from doing any logical area operations while an online backup is active.

Now, when RMU Backup Online initializes, it locks the AIP and keeps it locked for the duration of the backup. That prevents any AIP changes that could cause snapshot page related errors.

A new message is now displayed when an RMU Backup Online command is issued.

```
%RMU-I-AIPLOCK, pausing logical area creation and deletion at 7-NOV-2024 16:35:13.13
```

If there is a transaction active that created or deleted a logical area, RMU Backup Online will stall waiting for that transaction to complete. An example of the stall message displayed by RMU Show Statistics is shown below.

```
Node: NODE1 (1/1/16)   Oracle Rdb V7.4-130 Perf. Monitor   7-NOV-2024 09:17:20.12
Rate: 3.00 Seconds           Stall Messages           Elapsed: 00:03:33.36
Page: 1 of 1                DEV:[DIR]MF_PERSONNEL           Mode: Online
-----
Process.ID Since..... T Stall.reason.....Lock.ID.
5B84AC3A:1u 7-NOV-2024 09:19:36.1516290 R waiting for AIP larea synch (EX)
```

This problem has been corrected in Oracle Rdb release V7.4.1.4.

2.3.6 RMU Show Statistics Process Monitor Counters Incorrect on RAD Systems

When monitoring database statistics using the RMU Show Statistics utility the displayed counters could be incorrect when the Process Monitor option was selected from the tools “!” menu. The numbers were only incorrect if the system CPU supports the RAD/NUMA memory feature. When process monitoring was selected, the counters displayed would represent all users, not just the selected user, as expected.

RMU Show Statistics has been corrected to only show counters unique to the selected process. There is no way to work around this problem.

This problem has been corrected in Oracle Rdb release V7.4.1.4.

2.3.7 Unexpected Bugcheck When Using RMU Show Statistics Configure Qualifier

Bug 32562913

In prior releases of Oracle Rdb the RMU Show Statistics command could bugcheck when accessing a database with ROW CACHE enabled.

The following example shows the reported problem:

```
$ RMU/SHOW STAT SQL$DATABASE -  
  /CONFIG=REPORT.CFG/NOINTER/WRITE_REPORT_DELAY=5 -  
  /REPORT_SCREEN=(INCLUDE="checkpoint*")  
%COSI-F-BUGCHECK, internal consistency failure  
%RMU-F-FATALOSI, Fatal error from the Operating System Interface.  
%RMU-I-BUGCHKDMP, generating bugcheck dump file USERS2:[BUGCHECK]RMUBUGCHK.DMP;
```

The bugcheck dump would have a footprint similar to the following:

```
***** Exception at 808A1530 : RMU74\DI0CCH$FREE_RCCB + 00000040  
%SYSTEM-F-ACCvio, access violation, reason mask=00,  
  virtual address=0000000000000080, PC=FFFFFFFF808A1530, PS=0000001B  
Saved PC = 80A8E020 : RMU74\KUTDIS$WRITE_REPORT + 000012D0  
Saved PC = 80A24500 : RMU74\KUTDIS$NEW_NODE_ASTX + 00000C20  
Saved PC = 8091D360 : RMU74\KUTDIS$DISPATCH + 00000160  
Saved PC = 8090C1E0 : RMU74\KUT$DISPLAY + 000045D0  
Saved PC = 807B8450 : RMU74\RMU$DISPLAY + 00004C40  
Saved PC = 8001B400 : RMU74\RMU$SHOW + 00001A10  
Saved PC = 803CF180 : RMU74\RMU_DISPATCH + 00002C90  
Saved PC = 803CBD40 : RMU74\RMU_STARTUP + 00000920  
Saved PC = 80000050 : RMU74\RMU$MAIN + 00000050  
Saved PC = 884D4E80 : PTHREAD$RTL\thdBase  
  Module THD_THREAD + 00005C00; line 245262  
Saved PC = 8848A6C0 : PTHREAD$RTL\pthread_main  
  Module THD_INIT + 000006C0; line 245046  
Saved PC = 7AD97480 : Image DCL + 0007D480
```

This problem has been corrected in Oracle Rdb Release V7.4.1.4.

Software Errors Fixed in Oracle Rdb Release 7.4.1.3

This chapter describes software errors that are fixed by Oracle Rdb Release 7.4.1.3.

3.1 Software Errors Fixed That Apply to All Interfaces

3.1.1 Accessing Row Cache Fails with CANTMAPSHMEM, NOSUCHSEC Error

Bug 36453739

It was possible for applications accessing a row cache to fail with the following error when attempting to initialize the row cache.

```
%RDMS-F-CANTMAPSHMEM, error mapping to shared memory "RDM..."  
-SYSTEM-W-NOSUCHSEC, no such (global) section
```

If the row cache server (RCS) encountered the error the database was shutdown.

This problem could occur in rare circumstances where an attempt to create the shared memory section for cache would encounter an error. After the failure, internal data structures would show that the cache exists when in fact it does not. Subsequent attempts to access that cache would result in the error shown above.

This problem has been corrected in Oracle Rdb Release 7.4.1.3. The cache shared memory section must be successfully created before setting the indicator that shows the cache exists.

3.2 SQL Errors Fixed

3.2.1 CAST Function Converted Blank Strings to a Numeric Value 0

Bug 32389302

In prior releases of Oracle Rdb the CAST function converted blank strings (or zero length strings) to a numeric value 0. This was expected behavior for Oracle Rdb but not for the ANSI/ISO Database SQL Language standard, nor for Oracle Database compatibility.

The following example shows the behavior.

```
SQL> select cast('          ' as bigint) from rdb$database;
0
1 row selected
SQL>
```

This has been the behavior of Oracle Rdb since the first version and is possibly expected by current customers. Therefore, the CAST function is only changed for ORACLE LEVEL3 and SQL2011 dialects so that it is conformant with the behavior expected by Oracle Database applications and the ANSI/ISO SQL Database Standard.

```
SQL> select cast('          ' as bigint) from rdb$database;
%RDB-E-ARITH EXCEPT, truncation of a numeric value at runtime
-COSI-F-INVNBDS, invalid numeric byte data string
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.4.1.3.

3.2.2 Unexpected SQL-F-BUGCHK When Using SAVEPOINT Statement

Bug 35469620

Attempts to use the SAVEPOINT statement in an EXEC SQL statement (SQL Precompiler) or within a PROCEDURE (SQL Module Language) resulted in a SQL-F-BUGCHK error.

```
$ sql$mod test_mod.sqlmod
%SQL-F-BUGCHK, There has been a fatal error.
Please contact your Oracle support representative. gem$gendml - 2
```

This problem has been corrected in Oracle Rdb Release 7.4.1.3. Both the SQL Module Language and SQL Precompiler correctly handle the SAVEPOINT statement.

3.2.3 Hash Join Behavior Mismatch

Bug 35760888

It has been observed that a query may request to use HASH JOIN but that strategy is not selected. If the SET FLAGS command is used to specify 'STRATEGY' the displayed message is confusing. For example:

```
SQL> select e.employee_id, (sh.salary_end - sh.salary_start) month (3)
cont> from employees e, salary_history sh
cont> where e.employee_id = sh.employee_id
cont> and e.employee_id <= '00164'
cont> and sh.salary_end is not null
cont> optimize for HASH JOIN;
~Sh: Hash Join #4S inner rejected-EMP_LAST_NAME and SH_EMPLOYEE_ID column 0 lengths differ
Tables:
  0 = EMPLOYEES
  1 = SALARY_HISTORY
Cross block of 2 entries  Q1
Cross block entry 1
Index only retrieval of relation 0:EMPLOYEES
Index name EMP_EMPLOYEE_ID [0:1]
Keys: 0.EMPLOYEE_ID <= '00164'
```

```

Cross block entry 2
  Leaf#01 FFirst 1:SALARY_HISTORY Card=729
    Bool: (0.EMPLOYEE_ID = 1.EMPLOYEE_ID) AND NOT MISSING (1.SALARY_END)
    BgrNdx1 SH_EMPLOYEE_ID [1:1] Fan=17
    Keys: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID
    Bool: 1.EMPLOYEE_ID <= '00164'
E.EMPLOYEE_ID
00164          008
00164          006
00164          016
3 rows selected

```

In this release of Oracle Rdb, if HASH JOIN is requested via the clause OPTIMIZE FOR HASH JOIN but the optimizer determines that HASH JOIN is not possible (probably caused by an earlier choice of a CROSS JOIN solution) it will now display this new informational message. For example:

```

SQL> select e.employee_id, (sh.salary_end - sh.salary_start) month (3)
cont> from employees e, salary_history sh
cont> where e.employee_id = sh.employee_id
cont> and e.employee_id <= '00164'
cont> and sh.salary_end is not null
cont> optimize for HASH JOIN;
~S: Hash Join for Q1 requested but was not possible when solution #7S is Cross join
Tables:
  0 = EMPLOYEES
  1 = SALARY_HISTORY
Cross block of 2 entries  Q1
Cross block entry 1
  Index only retrieval of relation 0:EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [0:1]
  Keys: 0.EMPLOYEE_ID <= '00164'
Cross block entry 2
  Leaf#01 FFirst 1:SALARY_HISTORY Card=729
    Bool: (0.EMPLOYEE_ID = 1.EMPLOYEE_ID) AND NOT MISSING (1.SALARY_END)
    BgrNdx1 SH_EMPLOYEE_ID [1:1] Fan=17
    Keys: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID
    Bool: 1.EMPLOYEE_ID <= '00164'
E.EMPLOYEE_ID
00164          008
00164          006
00164          016
3 rows selected

```

If HASH JOIN is not requested via the clause OPTIMIZE FOR HASH JOIN but hash join is enabled by using the SQL SET FLAGS statement with 'HASHING(JOINS)', and hash join is rejected, it will print out the following ~S strategy message:

```

SQL> set flags 'hashing(joins)';
SQL> select e.employee_id, (sh.salary_end - sh.salary_start) month (3)
cont> from employees e, salary_history sh
cont> where e.employee_id = sh.employee_id
cont> and e.employee_id <= '00164'
cont> and sh.salary_end is not null;
~S: Hash Join #7S for Q1 tried but rejected
Tables:
  0 = EMPLOYEES
  1 = SALARY_HISTORY
Cross block of 2 entries  Q1
Cross block entry 1
  Index only retrieval of relation 0:EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [0:1]
  Keys: 0.EMPLOYEE_ID <= '00164'

```

```

Cross block entry 2
Leaf#01 FFirst 1:SALARY_HISTORY Card=729
Bool: (0.EMPLOYEE_ID = 1.EMPLOYEE_ID) AND NOT MISSING (1.SALARY_END)
BgrNdx1 SH EMPLOYEE_ID [1:1] Fan=17
Keys: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID
Bool: 1.EMPLOYEE_ID <= '00164'
E.EMPLOYEE_ID
00164          008
00164          006
00164          016
3 rows selected
SQL> set flags 'nohashing(joins)';

```

This problem has been corrected in Oracle Rdb Release 7.4.1.3.

3.3 RMU Errors Fixed

3.3.1 Invalid RMU Show Statistics File IO Overview Screen PgDis Field Values

When the Oracle Rdb RMU Show Statistics command was executed to display the **File IO Overview** screen, and the "PgDis" - pages discarded - field numeric values were too large to fit in this fixed length 4 digit field, the value "****k" was displayed, where the asterisks followed by the "k" indicate that the numeric value could not fit in this field, even though the number was scaled by 1000.

This scaling problem has been fixed. The value will be scaled to fit on the screen. If a value exceeds 10000 it will be scaled by 1000 and displayed with a "k" suffix, and so on. An "m" suffix indicates scaling by 1000^2 , a "g" suffix indicates scaling by 1000^3 , and a "t" suffix indicates scaling by 1000^4 .

This problem has been corrected in Oracle Rdb Release 7.4.1.3. Note that if "****k" was displayed then that by itself indicates excessive pages discarded and investigation is encouraged.

3.3.2 RMU Verify Validate FLOAT did not Correctly Scan the Target Tables

Bug 35220492

RMU Verify Validate FLOAT was not processing the table completely and therefore not displaying found bad floating point values.

The following example shows the expected output when bad values are found.

```

$          RMU/VERIFY/VALIDATE=FLOAT/LOG TESTING_DB
%RMU-I-BGNROOVER, beginning root verification
%RMU-I-ENDROOVER, completed root verification
%RMU-I-DBBOUND, bound to database
          "USERS1:[TESTING.VALIDATE_FLOAT]TESTING_DB.RDB;1"
%RMU-I-OPENAREA, opened storage area
          USERS1:[TESTING.VALIDATE_FLOAT]TESTING_DB.RDB;1 for protected retrieval
%RMU-I-BGNAIPVER, beginning AIP pages verification
%RMU-I-ENDAIPVER, completed AIP pages verification
%RMU-I-BGNABMSPM, beginning ABM pages verification
%RMU-I-ENDABMSPM, completed ABM pages verification
%RMU-I-CLOSAREAS, releasing protected retrieval lock on all storage areas

```

```
%RMU-I-BGNVALDAT,      Beginning of table data validation 28-MAR-2023 15:26:21.01
%RMU-I-TABCOLVLD, Scanning table FLOATS_TAB which has 2 floating point columns.
%RMU-I-INVDATF,      Invalid F floating data 01438000 in field F_FLOAT
      of table FLOATS_TAB at dbkey 57:498:2
%RMU-I-INVDATG,      Invalid G floating data 0000000001438000 in field G_FLOAT
      of table FLOATS_TAB at dbkey 57:498:2
%RMU-I-ENDVALDAT,      End of data validation at 28-MAR-2023 15:26:21.04
%RMU-S-ENDVERIFY, elapsed time for verification :    0 00:00:00.13
$
```

The provided DBKEY can be used to update the rows such as shown in this example.

```
SQL> select * from FLOATS_TAB where dbkey = _dbkey'57:498:2';
      ROW_NUM      F_FLOAT      G_FLOAT
%RDB-E-ARITH_EXCEPT, truncation of a numeric value at runtime
-COSI-F-FLTINV, invalid floating conversion
SQL> update FLOATS_TAB set F_FLOAT = null, G_FLOAT = null
cont> where dbkey = _dbkey'57:498:2';
      3
1 row updated
SQL> select * from FLOATS_TAB;
      ROW_NUM      F_FLOAT      G_FLOAT
      1      1.0000000E+00      1.0000000000000000E+000
      2      2.0000000E+00      2.0000000000000000E+000
      3      NULL      NULL
3 rows selected
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.4.1.3.

3.3.3 Changed Processing of the RMU Load Exception_file Option

In prior releases of Oracle Rdb the RMU Load command would not open the Exception_file until the first exception was detected. Unfortunately, if the Exception_file was improperly specified this state may not be detected until RMU had already performed significant I/O. In particular, during a parallel load this may cause some of the executors to abort during the load because they cannot record exceptions.

Reasons for the an erroneous Exception_file specification may include:

- The file specification is incorrect OpenVMS file naming format
In the case of normal execution the specification captured from the DCL command line might never be used and therefore no validation was performed.
- The file specification is no longer valid
This might be due to changes in logical name translations, names of devices, or names of directories which appear in the specification captured from the command line.
- A version limit was exceeded
If exception file already exist but has reach 32767 then no further files can be created. This would be corrected by purging, those versions and renaming to version ;1.
- The target device is not mounted.

- and so on.

With this release of Oracle Rdb the RMU Load command now attempts to validate the file specification defined by the `Exception_file` option on the command line or in the plan file before beginning the load so that these problems can be resolved. In the case of Parallel Backup the actual file name used includes the executor id number and may not be fully validated.

This change appears in Oracle Rdb Release 7.4.1.3.

3.3.4 RMU Recover of Duplicated Database Deleted Journal in Original Database

Bug 35469591

When duplicating a database using the RMU Restore and RMU Recover commands it was possible for a journal in the original database to be deleted when executing the RMU Restore command against the duplicate database. This problem could occur when a journal was dropped in the original database, creating a journal entry that denotes the journal deletion, and then a new journal was created in the original database with the same name as the deleted journal. When the journal containing the journal deletion was applied to a separate database, the RMU Recover command would delete the journal used by the original database.

The following sequence demonstrates the problem. In the original database, delete and then recreate a journal.

1. Backup the original database
2. Recreate a journal in the original database


```
ALTER DATABASE FILE TEST DROP JOURNAL AIJ2;
ALTER DATABASE FILE TEST ADD JOURNAL AIJ2 FILENAME TEST_2;
```

Restore the database backup to a new location and then apply the journals created for the original database. When the RMU Recover command is executed, it will delete the journal AIJ2 that is being used by the original database.

A similar problem was addressed in a previous release of Oracle Rdb. If a journal was created in the original database and the journal containing the “create journal” entry was applied to a duplicate database, the journal in the original database could be overwritten by the RMU Recover command. A new qualifier `/NOAFTER` was created to prevent RMU Recover from applying the journal create entries. However, that qualifier did not prevent the application of “delete journal” entries. In this release, the `/NOAFTER` qualifier will prevent the application of both create and delete journal entries.

This problem has been corrected in Oracle Rdb Release 7.4.1.3.

3.3.5 Stale Entries Written to Statistics Lock Timeout Log

Bug 34957220

If the RMU Show Statistics command is used to create a lock timeout logfile (when the `Lock_timeout_log` qualifier is specified) the logfile could contain entries from events that occurred before the utility was invoked.

In the following output, note the timestamp for the log creation “4-JAN-2023 00:10:04.34” is newer than the entry in the log with a timestamp of “2-JAN-2023 21:25:04.50”.

```
Oracle Rdb V7.4-110 Performance Monitor Lock Timeout Log
Database DEV: [DIR]TEST.RDB;1
Lock Timeout Log created 4-JAN-2023 00:10:04.34
00078E7B:1 4-JAN-2023 16:25:04.5290650 - waiting for record 520:2370514:0 (EX)
00079A6A:1 2-JAN-2023 21:25:04.5008225 - waiting for record 520:2354672:0 (EX)
```

This problem was caused by incorrect initialization of the lock timeout data structure used by RMU Show Statistics.

The logging data structure is now properly initialized. Only timeouts that occur after the RMU Show Statistics utility is invoked will be logged.

This problem has been corrected in Oracle Rdb Release 7.4.1.3.

3.3.6 System Space Buffers not Freed Up After an RMU Close command

Bug 13590034

In prior releases of Oracle Rdb a database that has SYSTEM SPACE BUFFERS buffers enabled (SHARED MEMORY IS SYSTEM), never freed the memory that was allocated when the database was closed. For example, after some sequence of database open and close the following error would be reported.

```
$ RMU/OPEN USERS1:[TESTER.DATABASE]TEST.RDB
%RDMS-F-CANTOPENDB, database could not be opened as requested
-RDMS-F-CANTCREGBL, error creating and mapping database global section
-SYSTEM-F-INSFSPTS, insufficient SPTEs available
%RMU-W-CANTOCDB, Error encountered while opening or closing database file
USERS1:[TESTER.DATABASE]TEST.RDB;1
```

In some rare cases a OpenVMS system crash could occur.

The following example shows how to enable SYSTEM SPACE BUFFERS.

```
$ RMU/SET SHARED_MEMORY/TYPE=SYSTEM TEST
%RMU-I-MODIFIED, Shared Memory state modified
%RMU-W-DOFULLBCK, full database backup should be done to ensure future recovery
$
```

Once system space buffers are enabled, a dump of the header will indicate that the database will be mapped into system space. Below is an abbreviated section of the Database Parameters that indicate that the database will be mapped in system space. For example:

```

$ RMU/DUMP/HEADER TEST
.
.
.
Database Parameters:
  Root filename is "USERS1:[TESTER.DATABASE]TEST.RDB;1"
  Created at 19-SEP-2023 09:07:46.42
  Physical database created at 19-SEP-2023 09:07:46.42
  Oracle Rdb structure level is 74.0
  Maximum user count is 50
  Maximum node count is 16
  Database open mode is AUTOMATIC
  Database close mode is AUTOMATIC
  Database will be mapped in system space
.
.
.
$

```

Conversely, SYSTEM SPACE BUFFERS can be toggled back to using process space by the qualifier on the RMU Set command. For example:

```

$ RMU/SET SHARED_MEMORY/TYPE=PROCESS TEST
%RMU-I-MODIFIED, Shared Memory state modified
%RMU-W-DOFULLBCK, full database backup should be done to ensure future recovery
$

```

If you are experiencing this type of error, the work around would be to revert the database back to being mapped into process space, followed by a system reboot to regain usage of all of the available system page table entries.

This problem has been corrected in Oracle Rdb Release 7.4.1.3. Any system space buffers that were left allocated by a previous version of Oracle Rdb will still require a system reboot for reclamation. However, with this release open and close of the database will correctly return system memory to a free memory pool for future use.

3.3.7 Unexpected Failure of RMU Convert Commit Command

Bug 35970239

If a database was converted to Oracle Rdb V7.3 using first **RMU Convert NoCommit** and later completed with **RMU Convert Commit** it was possible that Rdb V7.3 left stale metadata as part of the RDB\$INDICES table. This stale metadata was inadvertently used by an RMU Convert Commit step during an Oracle Rdb V7.4 convert strategy to cleanup obsolete metadata. However, that cleanup was already performed as part of the Rdb V7.3 conversion.

This problem has various symptoms such as a failed RMU Convert command (as shown in the example below), errors failing to ATTACH to the database after a successful RMU Convert, or possibly errors selecting data from some tables.

```

$ rmu/convert/noconfirm/commit ABC
%RMU-I-RMUTXT_000, Executing RMU for Oracle Rdb V7.4-120 on OpenVMS IA64 V8.4-2L3
%RMU-I-LOGCONVRT, database root converted to current structure level
%RMU-I-CVTCOMSUC, CONVERT committed for USERS2:[TESTER]ABC.RDB;1 to version V7.4
%RDB-E-NO_RECORD, access by dbkey failed because dbkey is no longer associated with a record
-RDMS-F-NODBK, 1:193:14 does not point to a data record
%RDB-E-OPEN_TRANS, detach failed: 1 transaction still active
$

```

Other related failures include RMU Convert failing with this error

```
%RDB-E-BAD_DB FORMAT, RDMS_BUILD_DB does not reference a database known to Rdb
-RMS-E-FNF, file not found
%RMU-F-FATALRDB, Fatal error while accessing Oracle Rdb
```

Or failing attempts to attach to the converted database

```
SQL> attach 'filename USERS2:[TESTER]XYZ.RDB';
%SQL-F-ERRATDEC, Error attaching to database USERS2:[TESTER]XYZ.RDB;1
-RDB-E-NO RECORD, access by dbkey failed because dbkey is no longer associated with a record
-RDMS-F-NÖDBK, 1:641:7 does not point to a data record
```

This problem has been corrected in Oracle Rdb Release 7.4.1.3. The RMU Convert Commit command now detects and ignores the stale metadata from Rdb V7.3.

- Any Oracle Rdb databases currently in the first stage (*RMU Convert NoCommit*) should **not** be committed until after this update is applied.
- Any database that was completely converted using the aforementioned RMU Convert steps via Rdb V7.3 may exhibit these symptoms. If so please contact Oracle Support for assistance.
- Databases converted from Rdb V7.0, V7.1, or V7.2 directly to Rdb V7.4 are not affected by this problem.
- Databases converted from Rdb V7.3 to Rdb V7.4 using a single step **RMU Convert Commit** are not affected by this problem.

3.3.8 Unexpected Reset of Logical Area Length by RMU Restore Initialize_tsns

In prior releases of Oracle Rdb the RMU Restore Initialize_tsns command would erroneously reset the logical area details to zero.

The following example shows that the logical area length was changed to zero. This occurs for system and user tables, and indices.

```
$! Initialize TSN
$ rmu/restore-
  /only_root-
  /nolog-
  /initialize_tsns-
  /dir=new_database -
  /after=new_database:pers.aij -
  mf_pers_back
$
$ rmu/show aip /brief /type=sorted new_database:mf_personnel
*-----
* Logical Area Name          LArea PArea   Len Type
*-----
RDB$NDX_SEG_NAM_FLD_POS_NDX      12     1     0 SORTED INDEX
RDB$NDX_NDX_NAME_NDX             13     1     0 SORTED INDEX
.
.
.
```

This problem has been corrected in Oracle Rdb Release 7.4.1.3. RMU now correctly initializes the TSN details on the AIP (area inventory) pages.

3.4 RDO, RDBPRE and RDML Errors Fixed

3.4.1 Improved Error Detection for RDBPRE/BASIC Applications Using PACKED DECIMAL Types

In prior releases of Oracle Rdb the runtime support for RDBPRE BASIC applications was not correctly reporting conversion errors. When an application uses PACKED DECIMAL data types in BASIC as sources or targets for Rdb data then calls are made to conversion routines in the supplied object library SYS\$LIBRARY:RDBVMSLIB74.OLB.

This library has been updated in this release with improved error detection when an overflow could occur when converting from/to DECIMAL types.

The following example shows an attempt to fetch the SALARY_AMOUNT (signed longword scale -2) from SALARY_HISTORY into a DECIMAL(5,1) local variable.

```
program TEST
&rdb&  invoke database filename "personnel"
        declare decimal (5,1) number1

! Really expect this to fail because 5,1 is too small for the salary-amount
print "Decimal (5,1)"
&rdb&  for r in salary_history with r.employee_id = '00173'
&rdb&      sorted by r.salary_start
&rdb&      get
&rdb&          number1 = r.salary_amount
&rdb&      end_get
&rdb&      print "number1 is ", number1
&rdb&  end_for
&rdb&  commit
&rdb&  finish
end
```

The new behavior is to report the conversion error at runtime. The target of the GET should be a DECIMAL type allowing larger values.

```
$ RUN PACKED_DECIMAL_TEST
Decimal (5,1)
%RDB-F-CONVERT_ERROR, invalid or unsupported data conversion
-LIB-F-DECOVF, decimal overflow
$
```

Oracle encourages customers with RDBPRE /BASIC applications to relink them with the latest runtime support for RDBPRE BASIC. The LINK command implicitly locates the Oracle Rdb runtime support for RDBPRE BASIC.

This problem has been corrected in Oracle Rdb Release 7.4.1.3.

Software Errors Fixed in Oracle Rdb Release 7.4.1.2

This chapter describes software errors that are fixed by Oracle Rdb Release 7.4.1.2.

4.1 Software Errors Fixed That Apply to All Interfaces

4.1.1 On Reboot the RMONSTART74.COM Procedure Failed to Restart the RDBSERVER Service

Bug 32834189

The Oracle Rdb procedure SYS\$STARTUP:RMONSTART74.COM should enable the RDBSERVER service upon execution. This is particularly important after a system reboot. Unfortunately this action was erroneously omitted from the Oracle Rdb V7.4.1 and V7.4.1.1 releases.

A workaround is to add this command to the SYSTEM startup command procedure:

```
$ PIPE TCPIP ENABLE SERVICE RDBSERVER > nla0: 2> nla0:
```

This should be done immediately after the line that starts Oracle Rdb

```
$ @SYS$STARTUP:RMONSTART74
$ PIPE TCPIP ENABLE SERVICE RDBSERVER > nla0: 2> nla0:
$
```

This problem has been corrected in Oracle Rdb Release 7.4.1.2. A revised version of SYS\$LIBRARY:SQL\$STARTUP.COM which is ultimately called by RMONSTART74.COM will again include this TCP/IP command.

4.1.2 Unexpected ACCVIO and BUGCHECK Dump for Some Queries on Integrity Systems

Bug 32931643

In prior releases of Oracle Rdb running on Integrity servers it was possible that large and complex queries or procedures might fail with an ACCVIO and have a bugcheck foot print similar to the following:

```
***** Exception at 000000008246D4C1 : symbol not found
%SYSTEM-F-ACCVIO, access violation, reason mask=04, virtual
address=000000000200850, PC=FFFFFFFF8246D4C1, PS=0000000B
```

```

Saved PC = 0000000081189030 : RDMSHRP73\RDMS$$EXE_LEAF +
0000C170
Saved PC = 00000000810E5E70 : RDMSHRP73\RDMS$$EXE_NEXT +
000073C0
Saved PC = 00000000810E76A0 : RDMSHRP73\RDMS$$EXE_NEXT +
00008BF0
Saved PC = 00000000811C4B40 : RDMSHRP73\RDMS$$C_EXE_NEXT +
00000070
Saved PC = 000000008097FE40 : RDMSHRP73\RDMS_EXE_INTERP +
0000F3E0
Saved PC = 00000000811BDEF0 : RDMSHRP73\JSB_REQX + 00000120
Saved PC = 00000000811BCDD0 : RDMSHRP73\RDMS$TOP_START_
REQUEST + 00001A00
Saved PC = 0000000081CCA1A0 : Image RDMSHRP73 + 000001A0

```

A temporary workaround would be to disable the code optimization for the affected query or application. The affected applications will likely run more slowly since optimization is disabled. This workaround can be removed once this update (for V7.4.1.2 or later) has been installed.

Use this command to see how to disable the code optimization for Integrity systems.

```
$ RMU /SHOW LOGICAL_NAME RDMS$BIND_CODE_OPTIMIZATION /DESCRIPTION
```

This problem has been corrected in Oracle Rdb Release 7.4.1.2.

4.1.3 Inserting New Row Fails With RDMS-F-ABMCHNFUL Error

In rare circumstances, when attempting to insert a new row into a table, an RDMS-F-ABMCHNFUL error could be returned.

Note

In older releases, that is prior to release 7.2.5.3, the error RDMS-F-STAREAFUL error was returned, even if it was possible for the physical storage area to be extended.

Oracle Rdb uses an internal data structure called an Area Bitmap (ABM) to represent all the pages in a storage area that could hold rows for a specific table or index (logical area). In prior releases, that structure has a finite number of entries. Once all of the bitmap was utilized, it was not possible to insert any more data into a logical area, regardless of whether or not there was space available in the storage area. Any attempt to insert additional data into the logical area would fail with the exceptions described above. The logical area might be the table itself, or one of the indices defined on that table.

For most applications, it is very unusual to encounter this limitation. Typically, this error would only be returned for storage areas that have a small page size, and are used by tables that have a large number of rows. Smaller pages provide less room for a bitmap, which increases the likelihood that the bitmap would become full.

To avoid this problem, don't use small page sizes. That is, use the default value or larger page size when creating a storage area.

This problem has been corrected in Oracle Rdb Release 7.4.1.2. Now, when inserting new rows, if an ABM is completely full, Oracle Rdb will add another clump of ABM pages to the ABM chain instead of returning an error.

Note

A clump is defined as the number of pages that will fit wholly within the BUFFER as defined by the database BUFFER SIZE. This clump value is set for each storage area when it is created. Subsequent changes of the database BUFFER SIZE do not change this invariant clump size for the storage area.

Guidance and Warnings for the changed After Image Journals

Please note that this new feature adds several requirements:

- The new mechanism for extending the area bit map (ABM) requires new after image journal information to define this action and allow RMU Recover to successfully rollforward inserts that might also cause the ABM to be extended.

Therefore, all after image journals created with Rdb V7.4.1.2 have a new AIJ Version Number (741.1). This prevents older Oracle Rdb V7.4 releases from using those journals. Use RMU /Dump /AFTER_IMAGE to view the changed header.

```
$ RMU/DUMP/AFTER/END=1 TEST$:TEST_AIJ_BACKUP_0.AIJ
*-----
* Oracle Rdb V7.4-120                                dd-mm-yy hh:mm:ss.xxxx
*
* Dump of After Image Journal
*   Filename: DISK:[DIRECTORY]TEST_AIJ_BACKUP_0.AIJ;VERSION
*
*-----
1/1          TYPE=0, LENGTH=510, TAD=dd-mm-yy hh:mm:ss.xxxx, CSM=00
Database DISK:[DIRECTORY]TEST.RDB;VERSION
Database timestamp is dd-mm-yy hh:mm:ss.xxxx
Facility is "RDMSAIJ ", Version is 741.1
Database version is 74.0
AIJ Sequence Number is 0
.
.
.
```

Environments that use Hot Standby, LogMiner or RMU /Recover on other systems must ensure that those systems are upgraded to this Rdb release to handle the new after image journal version.

- Older versions of the after image journals (created using Rdb V7.4.1, or Rdb V7.4.1.1) can be still be applied to this Oracle Rdb release. Those AIJ Version Numbers would be 741.0.
- If for any reason the database administrator wishes to downgrade from Rdb V7.4.1.2 to a prior Rdb V7.4 release then new after image journals must be created and a full database backup created against which possible RMU Restore and RMU Recover operations can be applied.

- Attempts to recover a database running Rdb V7.4.1 or Rdb V7.4.1.1 using the new After Image journals will be prevented:

```
$ RMU/RECOVER/NOLOG/ROOT=TEST$:TEST TEST$:TEST AIJ_BACKUP_0
%RMU-F-BADAIJVER, after-image journal version is incompatible with the runtime system
%RMU-F-FTL_RCV, Fatal error for RECOVER operation at dd-mmm-yyyy hh:mm:ss.xxxx
$
```

4.1.4 Error From Disk Write Results In Hang

Bug 34156672

An I/O intensive application would sometimes hang in Oracle Rdb in LEF state if an error was encountered when writing data page buffers to disk. Examination of the process stack revealed that an exception occurred, and Oracle Rdb was waiting for a SYSSYNCH system call to return.

Oracle Rdb writes one or more modified pages to disk concurrently in a single batch. In the unusual situation where no write requests succeeded, Oracle Rdb would still wait for the write requests to complete, even though none were started.

This problem has been corrected in Oracle Rdb Release 7.4.1.2. If no disk writes were successfully initiated Oracle Rdb will no longer attempt to wait for the writes to complete.

4.1.5 Unexpected COSI-F-EXQUOTA Error Using Buffer Objects

Bug 34156672

It was possible for a *COSI-F-EXQUOTA* error to be signaled when the BUFFER OBJECT feature was enabled and many data page updates were being written to disk concurrently. An error message like the following may be returned, or a hang might occur.

```
%RDB-F-IO_ERROR, input or output error
-RDMS-F-CANTWRITEDBS, error writing pages 1:2-3
-COSI-F-EXQUOTA, exceeded quota
```

Oracle Rdb was previously reserving enough BUFFER OBJECT resources to have an asynchronous write active on each data page buffer in the buffer pool. For example, if the database user was allocated 20 buffers, then a maximum of 20 asynchronous buffer object writes were allowed. However, there are scenarios where more than 20 concurrent page buffer writes are possible.

To avoid this problem, the database administrator can disable data page buffer objects. This is controlled using an RMU command or by changing the controlling logical.

```
$ RMU /SET BUFFER_OBJECT /DISABLE=(PAGE) MF_PERSONNEL
$ DEASSIGN RDM$BIND_PAGE_BUFOBJ_ENABLED
```

or

```
$ DEFINE RDM$BIND_PAGE_BUFOBJ_ENABLED 0
```

This problem has been corrected in Oracle Rdb Release 7.4.1.2. Oracle Rdb now allocates sufficient resources to allow the maximum possible concurrent database page buffer object writes.

4.2 SQL Errors Fixed

4.2.1 Incorrect Results From LISTAGG, FIRST_VALUE, and LAST_VALUE

In prior releases of Oracle Rdb it was possible that SQL would incorrectly generate single valued aggregate functions if they contained ORDER BY, FETCH FIRST ... ROWS ONLY (LIMIT TO ... ROWS), and OFFSET ROWS clauses. This problem affected only the functions LISTAGG, FIRST_VALUE, and LAST_VALUE used in a singleton select expression.

The following example shows the problem. This query wants the last value of LAST_NAME from the top 10 EMPLOYEES after ordering by FIRST_NAME but instead processes all the EMPLOYEES.

```
SQL> set flags 'trace';
SQL> begin
cont> declare :last_last char(20);
cont> set :last_last = (select last_value (last_name)
cont>                      from employees
cont>                      fetch first 10 rows only
cont>                      order by first_name);
cont> trace :last_last;
cont> end;
~Xt: Ziemke
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.4.1.2. Affected applications written using SQLSPRE or SQL\$MOD will need to be recompiled and relinked to correct this problem. Applications that use Dynamic SQL, such as ODBC, JDBC, SQL/Services and OCI Services for Rdb need no changes once upgrading to this release.

4.2.2 Unexpected NO_PRIV Errors From DECLARE LOCAL TEMPORARY TABLE Statement

Bug 34314418

In prior releases of Oracle Rdb V7.4 the CREATE ANY TEMPORARY TABLE privilege was erroneously checked for the DECLARE LOCAL TEMPORARY TABLE statement.

The following example shows the reported problem when a non-privileged user attempted to declare a temporary table.

```

declare local temporary table module.EMPLOYEES (
    EMPLOYEE_ID
        CHAR (5),
    LAST_NAME
        CHAR (14),
    FIRST_NAME
        CHAR (10),
    MIDDLE_INITIAL
        CHAR (1),
    ADDRESS_DATA_1
        CHAR (25),
    ADDRESS_DATA_2
        CHAR (20),
    CITY
        CHAR (20),
    STATE
        CHAR (2),
    POSTAL_CODE
        CHAR (5),
    SEX
        CHAR (1),
    BIRTHDAY
        DATE VMS,
    STATUS_CODE
        CHAR (1))
on commit preserve rows
comment is
    'personal information about each employee';
%RDB-E-NO_META_UPDATE, metadata update failed
-RDB-E-NO_PRIV, privilege denied by database facility

```

A work around for this issue is for the database administrator to grant CREATE ANY TEMPORARY TABLE privilege to the affected users.

Note

Oracle Corporation recommends that after Rdb V7.4.1.2 (or subsequent version) is installed that any such workaround be revoked as it will no longer be necessary.

This problem has been corrected in Oracle Rdb Release 7.4.1.2. The CREATE ANY TEMPORARY TABLE privilege is not required for the DECLARE LOCAL TEMPORARY TABLE statement. However, you must have REFERENCES or SELECT privilege on the table specified by any LIKE table clause used in the definition.

4.3 RMU Errors Fixed

4.3.1 Average Values Could be Zeroed on RMU Show Statistics Screens

It was reported that the "Summary IO Statistics", when displayed by RMU Show Statistics, would show the average values as zeroes. This occurred if the numeric values used to calculate the moving average values had become very large; such as when statistics have been collected over a long time period.

This problem has been corrected by increasing the numerical precision used to calculate these moving average values.

This problem has been corrected in Oracle Rdb Release 7.4.1.2.

4.3.2 RMU Recover Fails With RMU-F-CANTSYNCTSNS Using Correct Journal

Bug 33906596

After restoring a database, RMU Recover would sometimes fail with the following error, even though the correct journal was specified in the command.

```
%RMU-F-BACKUPNOAIJ, After Image Journaling was enabled after
the database was backed up or has since been disabled and reinitialized
-RMU-F-CANTSYNCTSNS, Last committed TSN 610 in the after image
journal file exceeds last committed TSN 609 in the database root
```

When database prestarted transactions were enabled (the default), there was a small timing window where the database backup could be assigned a transaction sequence number (TSN) that is less than the TSN stored in the active after-image journal. Attempts to use that journal to recover the restored database would fail with RMU-F-CANTSYNCTSNS.

The following scenario could lead to this problem.

1. There is an active user in the database that has committed an update and is sitting idle. By default, that user has a prestarted transaction active, even if no explicit SET TRANSACTION was executed. The prestarted transaction is assigned a TSN at the time that the previous transaction commits.
2. An after-image backup is executed. The newly activated journal is assigned the next database TSN. Executing an RMU Dump AFTER_JOURNAL command on the journal open record displays this value as "Last Commit TSN is".
3. The idle user executes another transaction and commits. Its older TSN becomes the last committed TSN for the database.
4. An online database backup is performed. The backup will record the last committed database TSN. This is displayed as "Latest full backup transaction sequence number is" in the RMU Dump HEADER output.
5. A journal backup is performed for the journal that was active at the time of the database backup.
6. The database is restored.
7. An RMU Recover command is executed, attempting to apply the backup of the journal that was active at the time the database full backup was performed.

When an RMU Recover command is started, RMU verifies that the specified journal is applicable to the database. RMU compares the last commit TSN in the journal with the latest backup TSN in the restored database. If the database TSN is greater than or equal to the journal's TSN then the journal is applied. Otherwise, an RMU-F-CANTSYNCTSNS error is issued. In the scenario described above, due to an old, prestarted transaction sequence number, this test would fail, even though the journal should have been applicable.

To workaroud this problem, provide the prior AIJ backup as well as the expected journal to the RMU Recover command. In the following example, the expected journal is AIJ_BACKUP_1. For the command to succeed, the prior journal, AIJ_BACKUP_0 is specified along with AIJ_BACKUP_1.

```
$ RMU/RECOVER/LOG AIJ_BACKUP_0, AIJ_BACKUP_1
```

This problem will not occur if prestarted transactions are disabled.

```
SQL> ALTER DATABASE
cont> FILENAME MF_PERSONNEL
cont> PRESTARTED_TRANSACTIONS ARE DISABLED
cont> ;
```

Note that disabling prestarted transactions will likely have a noticeable impact on performance in high transaction rate applications.

This problem has been corrected in Oracle Rdb Release 7.4.1.2. Now, when a quiet-point after-image journal backup is executed, all idle processes that have a prestarted transaction active are forced to obtain a new TSN when the next transaction is explicitly started.

4.3.3 SERVER_TRANSPORT Option to RMU Backup Parallel Not Working as Documented

Bug 32929697

The *Oracle RMU Reference Manual* (and DCL HELP) describes the option SERVER_TRANSPORT as allowing the choice between DECNET or TCPIP. Unfortunately, this option is accepted but does not change the default setting.

By default when RMU Backup Parallel is executed it will attempt to first use DECnet transport and then re-try using TCP/IP upon failure. If /PARALLEL=(..., SERVER_TRANSPORT=TCPIP, ...) is used then the fallback from a failed DECnet attempt will be avoided. The failure might occur because DECnet isn't installed, not active, or the current user does not have DECnet proxy definitions for the chosen cluster nodes. See the OpenVMS HELP for the AUTHORIZE utility ADD /PROXY command.

As an alternative to using the RMU Backup Parallel SERVER_TRANSPORT option you may use either the logical name RMU_NETWORK_TRANSPORT_TYPE or SQL_NETWORK_TRANSPORT_TYPE. Either logical name must be defined prior to running the RMU Backup Parallel Execute command or the RMU Backup Plan Execute command. These logical names expect an uppercase value for DECNET or TCPIP.

```
$ define /user RMU_NETWORK_TRANSPORT_TYPE "TCPIP"
$ rmu /backup -
    /parallel=(executor_count=3,nodes=(ALPHA,BETA,GAMMA)) -
    /execute -
    ...etc... -
    ADMIN_DB -
    ADMIN_BCK_JULY010
```

The setting for the SERVER_TRANSPORT option is now written to the generated Plan file when the List_Plan qualifier is used. However, if SERVER_TRANSPORT is not specified in the Parallel qualifier options then no SERVER_TRANSPORT setting is written to the Plan file. Therefore, if you plan to use either logical name then those logical names need to be defined prior to the RMU Backup .. Execute command.

This problem has been corrected in Oracle Rdb Release 7.4.1.2. RMU Backup now correctly supports the RMU Backup Parallel SERVER_TRANSPORT option, and writes it to the plan file when using the List_Plan qualifier. The translation of the RMU_NETWORK_TRANSPORT_TYPE and SQL_NETWORK_TRANSPORT_TYPE logical names will only be used if the option SERVER_TRANSPORT is not specified.

Software Errors Fixed in Oracle Rdb Release 7.4.1.1

This chapter describes software errors that are fixed by Oracle Rdb Release 7.4.1.1.

5.1 Software Errors Fixed That Apply to All Interfaces

5.1.1 Journal Backed Up While RMU Unload After Continuous Extracting Journal

Bug 18601419

In rare circumstances it was possible for a journal backup to start executing a backup on the database's current journal at the same time that a journal extraction (RMU/UNLOAD/AFTER/CONTINUOUS) was starting on the same journal. That could sometimes result in the extraction process not extracting all of the contents of the journal. No error would be displayed by the extraction process and it would continue normally.

To avoid this issue, do not allow journal backups to run until the journal extraction process has completed its initialization.

This problem has been corrected in Oracle Rdb Release 7.4.1.1. RMU/UNLOAD/AFTER/CONTINUOUS has been changed to complete its initial journal checkpoint sequence before allowing journal backups to proceed.

5.1.2 Bugchecks at SEQ\$GET_SEQUENCE + 000001C0 After Incremental Restore

Bug 31826798

After restoring a database and applying an incremental backup, attempts to fetch a sequence number from the database failed with the following exception.

```
***** Exception at 0000000081E65140 : RDMSHRP73\SEQ$GET_SEQUENCE + 000001C0
%COSI-F-BUGCHECK, internal consistency failure
Saved PC = 000000008158DAA0 : RDMSHRP73\RDMS$$SEQUENCE_GET_NEXT + 000000C0
```

This problem would only occur when additional journals were added to the database after a full backup was performed and before an incremental backup was executed.

1. RMU/BACKUP
2. RMU/SET AFTER /RESERVE=
3. RMU/BACKUP/INCREMENTAL

Various data structures are stored in an Oracle Rdb database root (.RDB) file, including the sequence numbers for a database and the list of journal files. If any of those structures change in size then the contents of the root file must be shuffled to accommodate the new size. The offsets for each of those data structures are saved at the beginning of the root file.

In this situation, after new journals were added to the database, the sequence number structures were moved. The new offset for the sequence number structure was backed up in the incremental backup. However, when the database incremental backup was applied to the restored database, RMU did not update the database root file with the new offsets for the various root file data structures. Subsequent attempts to access those structures could fail since some structures might not be at the expected location. Using the wrong offset would cause unrelated data to be read which could lead to various failures.

To avoid this problem, always issue a full backup when changing the following values.

- Number of journals
- Number of storage areas
- Number of row caches
- Number of sequences

This problem has been corrected in Oracle Rdb Release 7.4.1.1. When RMU processes an incremental restore, it now properly updates the data structure offsets in the database root file.

5.1.3 The SET FLAGS Option NOCOUNT_SCAN Erroneously Disabled Min and Max Optimized Index Lookups

On rare occasions, Oracle support may request that you use the Rdb database system flag "NOCOUNT_SCAN" to disable COUNT SCAN optimizations as a workaround for a problem you may be experiencing.

In prior releases of Oracle Rdb, the use of this flag also incorrectly suppressed the MIN and MAX key lookup optimizations carried out by the optimizer.

The following example shows this problem; selections no longer carry out "Min key lookup" or "Max key lookup" after disabling COUNT SCAN:

```
select max(employee_id) from employees;
Tables:
  0 = EMPLOYEES
Aggregate: 0:MAX (0.EMPLOYEE_ID) Q2
Index only retrieval of relation 0:EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [0:0]      Max key lookup

00471
1 row selected

! Now disable COUNT_SCAN
!
set flags 'nocount_scan';

select max(employee_id) from employees;
Tables:
  0 = EMPLOYEES
Aggregate: 0:MAX (0.EMPLOYEE_ID) Q2
Index only retrieval of relation 0:EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [0:0]
```



```

00471
1 row selected
SQL>

```

These problems have been corrected in Oracle Rdb Release 7.4.1.1. The NOCOUNT_SCAN keyword now only disables the COUNT (Index counts, Index counts lookup) and COUNT-DISTINCT (Index distinct counts, Index distinct lookup) optimizations when requested.

5.2 SQL Errors Fixed

5.2.1 Unexpected NULL Returned from MOD and SQRT Functions

Bug 32038549

In prior releases of Oracle Rdb, the MOD and SQRT functions erroneously returned NULL once these functions accepted a NULL parameter value.

```

SQL> select * from x order by id;
      ID      I
      1      123
      2      NULL
      3      456
3 rows selected
SQL> select id,i,mod(i,100) from x order by id;
      ID      I
      1      123
      2      NULL
      3      456
3 rows selected
SQL>

```

As a workaround the MOD and SQRT definitions in SQL_FUNCTIONS74.SQL could be used. Ensure that the optional SQL functions are defined in the database using SYSSLIBRARY:SQL_FUNCTIONS74.SQL. In such cases these functions must be executed using a quoted name to avoid activating the builtin functions.

```

SQL> set quoting rules 'sql99';
SQL> select id,i,"MOD"(i,100) from x order by id;
      ID      I
      1      123      2.300000000000000E+001
      2      NULL      NULL
      3      456      5.600000000000000E+001
3 rows selected
SQL>

```

This problem has been corrected in Oracle Rdb Release 7.4.1.1.

5.2.2 Unexpected RDMS-E-UNQCONFVAL when Using ALTER CONSTRAINT Statement

Bug 31914348

In prior releases of Oracle Rdb, attempts to use the ALTER CONSTRAINT statement to change a FOREIGN KEY constraint to NOT DEFERRABLE would result in an error. The following example shows this problem.

```

SQL> create table T1 (
cont>     C1
cont>         INTEGER,
cont>     C2
cont>         INTEGER,
cont>     constraint PK1
cont>         primary key (C1)
cont>         initially deferred deferrable);
SQL>
SQL> create table T2 (
cont>     C1
cont>         INTEGER,
cont>     constraint FK1
cont>         foreign key (C1)
cont>         references T1 (C1)
cont>         initially deferred deferrable);
SQL>
SQL> alter constraint PK1 not deferrable;
SQL>
SQL> alter constraint FK1 not deferrable;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDMS-E-UNQCONFVAL, the constraint "PK1" referenced by "FK1" has
    a conflicting evaluation time attribute
-RDMS-E-CONNOTCHG, constraint "FK1" has not been changed
SQL>

```

This problem has been corrected in Oracle Rdb Release 7.4.1.1. Such an ALTER CONSTRAINT statement is legal and no longer produces an error.

When SET DIALECT 'SQL2011' or 'ORACLE LEVEL3' is in effect, the ANSI and ISO SQL Database Language rules apply and an attempt to alter the PRIMARY KEY or UNIQUE constraint referenced by a FOREIGN KEY constraint will result in an error. For other dialects a warning will be issued.

5.2.3 Unexpected Bugcheck When Using More than One COUNT(*) in Query

Bug 31680491 and 31740191

In prior releases of Oracle Rdb, the following query unexpectedly bugchecks.

```

SQL> select
cont>     200-count(*) as v1,
cont>     count(*)      as v2
cont> from employees;
%RDMS-I-BUGCHKDMP, generating bugcheck dump file DISK_USER2:[USER]RDSBUGCHK.DMP;
%SYSTEM-F-ACCVIO, access violation, reason mask=00,
    virtual address=0000000000000028, PC=0000000000C25F14, PS=00000001B
SQL>

```

This occurs within the multi-aggregate feature introduced in Oracle Rdb V7.3.2.1. The optimizer is trying to match MAX(expr), MIN(expr), COUNT(expr), and COUNT(*) to perform optimized lookups.

A workaround is to include another aggregate function to the query to provide the optimizer with a source index for the query.

```

SQL> set flags 'strategy,detail(2)';
SQL>
SQL> select
cont>      200-count(*) as v1,
cont>      count(*)      as v2,
cont>      max(employee_id)
cont> from employees;
Tables:
      0 = EMPLOYEES
Aggregate: 0:COUNT (*) Q2
           1:COUNT (*) Q2
           2:MAX (0.EMPLOYEE_ID) Q2
Index only retrieval of relation 0:EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [0:0]      Index counts
Index only retrieval of relation 0:EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [0:0]      Index counts
Index only retrieval of relation 0:EMPLOYEES
  Index name  EMP_EMPLOYEE_ID [0:0]      Max key lookup
           V1                          V2
           100                         100      00471

1 row selected
SQL>

```

These problems have been corrected in Oracle Rdb Release 7.4.1.1.

5.2.4 Incorrect Handling of the DEFAULT Clause for AUTOMATIC UPDATE AS Columns

In prior releases of Oracle Rdb the DEFAULT keyword could be used to assign a value to any generated column; IDENTITY, AUTOMATIC AS, or GENERATED AS column. However, if the column was defined by the AUTOMATIC UPDATE AS clause and also included the DEFAULT clause in the definition the incorrect value was returned for the DEFAULT keyword during the INSERT or REPLACE statements. Rdb erroneously used the computed expression for the column instead of the DEFAULT. For AUTOMATIC UPDATE AS the computed expression should only be stored during an UPDATE statement.

This problem has been corrected in Oracle Rdb Release 7.4.1.1. Oracle Rdb now correctly handles the DEFAULT keywords for AUTOMATIC UPDATE AS columns. If there is no DEFAULT defined for the column then the result will be to assign NULL to the column.

5.3 RMU Errors Fixed

5.3.1 Unexpected SYSTEM-F-ACCVIO when Using RMU Dump Audit for DACCESS

Bug 32288393

In prior releases of Oracle Rdb, the RMU Dump Audit command may encounter an unhandled error and report a SYSTEM-F-ACCVIO exception when processing records in the AUDIT journal.

The following example shows the problem.

```

$ rmu/dump/audit -
  arg_database -
    RMU_AJ -
      /type=(DACCESS) -
      /format=LIST -
      /output=audit_dump.txt
%SYSTEM-F-ACCPIO, access violation, reason mask=00, virtual address=0000000000000002,
  PC=FFFFFFFF804D8F80, PS=0000001B
%RMU-F-FATALOSI, Fatal error from the Operating System Interface.
%RMU-F-FTL_RMU, Fatal error for RMU operation at 15-DEC-2020 09:43:45.04
$

```

This problem only affects the /FORMAT=LIST display (which is the default). As a workaround the qualifier /FORMAT=XML may be used to display the audit records, or the data can be successfully unloaded using RMU/UNLOAD/AUDIT command.

This problem has been corrected in Oracle Rdb Release 7.4.1.1.

5.3.2 RMU SET AUDIT and RMU SHOW AUDIT Now Work on Remote Databases

Bug 17859939

With this release of Oracle Rdb, the RMU Set AUDIT and RMU Show Audit commands will operate on remote databases that are specified using the OpenVMS DECnet file specification syntax. In prior releases various operations were performed expecting the database to be located on the current node. These restrictions have been removed.

The following example shows that the RMU privileges are still checked for the remote database. In addition the database must grant SELECT access to the remote user to allow an ATTACH to occur.

```

$ rmu/show audit remnod"j_jones":db$:manufacturing
%RMU-F-NOPRIVERR, no privileges for attempted operation
-COSI-F-NOPRIV, no privilege for attempted operation
%RMU-F-FTL_SHOW, Fatal error for SHOW operation at 11-JAN-2021 20:07:38.00

```

The database administrator should ensure that the user is granted RMUSSECURITY on the database to allow the RMU/SHOW AUDIT and RMU/SET AUDIT commands to execute.

```

$ rmu/set priv db$:manufacturing/acl=(ident=j_jones,access=rmu$security)

```

Now the RMU Show AUDIT command will execute.

```

$ rmu/show audit remnod"j_jones":db$:manufacturing
Security auditing STOPPED for:
  PROTECTION (disabled)
  RMU (disabled)
  DBVAULT (disabled)
  AUDIT (enabled)
  DACCESS (disabled)

Security alarms STOPPED for:
  PROTECTION (disabled)
  RMU (disabled)
  DBVAULT (disabled)
  AUDIT (enabled)
  DACCESS (disabled)

```

Note that identifiers (users and roles) will be stored according to the definitions on the remote node. So in the following example J_JONES and SQLNET4RDB must be defined on the remote node.

```
$ rmu/set audit -  
    /enable=ident=(j_jones,sqlnet4rdb) remnod"j_jones":db$manufacturing  
$ rmu/show audit remnod"j_jones":db$manufacturing/ident  
Enabled identifiers:  
    (IDENTIFIER=[MAN,J_JONES])  
    (IDENTIFIER=SQLNET4RDB)  
$
```

This problem has been corrected in Oracle Rdb Release 7.4.1.1.

Note

For most RMU operations the database can only be accessed locally due to the requirements to access local servers, and storage media. However, some small subset of RMU commands, including RMU Show AUDIT, RMU Set AUDIT, and RMU Extract are expected to successfully access remote databases.

5.3.3 RMU Backup After Space_check now Supports Disk Bound Volume Sets

Bug 32320241

In prior releases of Oracle Rdb, the RMU Backup After_journal SPACE_CHECK qualifier only checked the free space available on the first volume of a bound volume disk set.

The SPACE_CHECK qualifier is the default for the RMU Backup After_journal command, therefore, use /NOSPACE_CHECK as a workaround if the target device of RMU Backup is a bound volume set.

This problem has been corrected in Oracle Rdb Release 7.4.1.1. RMU now checks the total free space available on all members of a bound volume set.

Software Errors Fixed in Oracle Rdb Release 7.4.1.0

This chapter describes software errors that are fixed by Oracle Rdb Release 7.4.1.0.

6.1 Software Errors Fixed That Apply to All Interfaces

6.1.1 RDMS-F-BADPAGNUM Exception After Snapshot File Truncated

Bug 11073409

A bugcheck could occur in the following, rare, scenario:

1. A user updates a table but does not commit the transaction.
2. Another user truncates the snapshot file for the area updated by the first user.
3. Another user starts a read only transaction and attempts to read the table updated by the first user. The query fails and a bugcheck dump is produced.

The bugcheck contains a RDMS-F-BADPAGNUM exception, and the area noted in the error is the snapshot area. For example:

```
***** Exception at 0000000081FF7FC0 : RDMSHRP74\PIOFETCH$WITHIN_DB + 00000CB0
%RDMS-F-CANTREADDBS, error reading pages 4:2-2
-RDMS-F-BADPAGNUM, page 2 is out of valid range (1:1) for physical area 4
```

To avoid this problem, make sure there are no active update transactions accessing an area while a snapshot file truncate operation is active.

This error has been corrected in Oracle Rdb Release 7.4.1.0. All active transactions must now end before a snapshot file may be truncated. Once there is no database transaction activity, the snapshot file is truncated and then new transactions are allowed.

Note that attempting to truncate a snapshot file while there are idle transactions in the database will cause all database activity to stop. Care should be taken when using this feature. Ensure that there are no idle transactions before attempting to truncate a snapshot file.

6.1.2 SQL/Services V7.4 IVP Fails With RMU-E-NOSQCLIENT

Bug 30895210

The Oracle SQL/Services Release 7.4 IVP fails with the following error on a system where Oracle SQL/Services Release 7.3 has never been installed and Oracle Rdb 7.3 or 7.4 is being executed.

```
$ RMU/EXECUTE/COMMAND "RMU/SHOW VERSION"
%RMU-F-ERREXCCMD, Error executing command "RMU/SHOW VERSION".
-RMU-E-NOSQSLCLIENT, Cannot find image "SQLSRV_LIBCLIENT"
%RMU-F-FTL_RMU, Fatal error for RMU operation at 12-FEB-2020 11:03:20.19
```

RMU was attempting to activate the Oracle SQL/Services image SYSS\$LIBRARY:SQLSRV_LIBCLIENT73.EXE, rather than using the logical SQLSRV_LIBCLIENT to activate the desired image. Also, each release's Oracle SQL/Services startup and shutdown procedure was defining the logical SQLSRV_LIBCLIENT to the version specific image. That resulted in the logical being defined for whichever release of Oracle SQL/Services was last started up.

The SQLSRV_LIBCLIENT logical should be defined as a system logical in the system startup procedure, SYSS\$STARTUP:SYSTARTUP_VMS.COM, to point to the Oracle SQL/Services SQLSRV_LIBCLIENTnn image, where nn is the release of Oracle SQL/Services where the RMU_DISP dispatcher and RMU_SERVICE service are executing. These can only be running for one release of Oracle SQL/Services since they only execute using a predefined port that cannot be shared among different releases.

For example:

```
$ DEFINE/SYSTEM SQLSRV_LIBCLIENT SYSS$LIBRARY:SQLSRV_LIBCLIENT74.EXE
```

This problem has been corrected in Oracle Rdb Release 7.4.1.0 and Oracle SQL/Services 7.4. RMU will activate the image defined by the SQLSRV_LIBCLIENT logical and Oracle SQL/Services procedures will no longer define or deassign the logical.

6.1.3 Long Running Delete Does Not Advance After-Image Journal Checkpoint

Bug 6323701

If a database attach executes a data manipulation operation, like a delete, that takes a very long time to complete, it might not move its checkpoint location. In extreme circumstances, this could lead to all after-image journals filling without the possibility of a backup. This kind of problem is often referred to as a “long verb” issue.

To avoid this problem, try to break up large data manipulation operations into smaller ones.

This problem has been corrected in Oracle Rdb Release 7.4.1.0. Oracle Rdb now periodically checks to see if it has items on its work queue waiting to be executed. If there is a waiting request, it will be processed and then the current operation will continue. This should allow checkpoint locations to be advanced as the journals fill.

6.1.4 Long Running Query Blocks Other Queries

Bug 12972389

If a database user executes a long running query and it is holding a lock needed by another user, it might not release that lock for a long period of time, causing other users to wait. This kind of problem is often referred to as a “long verb” issue.

This problem has been corrected in Oracle Rdb Release 7.4.1.0. Oracle Rdb now periodically checks to see if it has items on its work queue waiting to be executed. If there is a waiting request, it will be processed and then the current operation will continue. This should prevent most long running queries from blocking other users.

6.2 SQL Errors Fixed

6.2.1 Unexpected Bugcheck When Using DROP STORAGE AREA ... CASCADE

Bug 30815275

In prior releases of Oracle Rdb, it was possible for ALTER DATABASE to bugcheck if one of the storage areas was missing. Specifically, when using ALTER DATABASE ... DROP STORAGE AREA ... CASCADE for an area that is referenced by an index or a table's storage map, this bugcheck could occur.

This problem has been corrected in Oracle Rdb Release 7.4.1.0. Oracle Rdb has been corrected and this is now the expected output from such a command under these conditions.

```
alter database
  filename mf_personnel
  drop storage area sample_data cascade
;
%RDB-F-SYS_REQUEST, error from system services request
-RDMS-F-FILACCERR, error opening storage area file
DISK: [DIRECTORY]SAMPLE_DATA.RDA;1
-RMS-E-FNF, file not found
```

6.2.2 Unexpected Error %SQL-F-FLDNOTCRS For Declaring a LIST Cursor

Bug 31028255

In previous releases of Oracle Rdb, the DECLARE CURSOR statement for a LIST cursor may generate an unexpected error as shown in the following example:

```
SQL> declare curs_dyn_t cursor for
cont>   select rr.resume res1, r2.resume res2
cont>     from resumes rr, resumes r2
cont>     where rr.employee_id = r2.employee_id;
SQL>
SQL> open curs_dyn_t;
SQL> fetch curs_dyn_t;
  RES1                                RES2
          1:720:1                      1:720:1
SQL>
SQL> declare curs_dyn_l list cursor for
cont>   select res1 where current of curs_dyn_t;
%SQL-F-FLDNOTCRS, Column RESUME was not found in the tables in current scope
SQL>
```

This problem has been corrected in Oracle Rdb Release 7.4.1.0.

6.3 RMU Errors Fixed

6.3.1 Clarification of RMU Unload After_Image Header

Bug 31191945

In prior versions of Oracle Rdb, the LogMiner header returned with each row describes the RDB\$LM_DATA_LEN and RDB\$LM_RECORD_VERSION as SIGNED WORD (or SMALLINT). However, these fields are, in fact, unsigned 16 bit integers.

An Oracle Rdb table can have a total column length that exceeds 32767, and such values were previously displayed in TEXT, DELIMITED_TEXT and DUMP formats as negative numbers.

The following example shows the definition from the record-definition file.

```
DEFINE FIELD RDB$LM_DATA_LEN DATATYPE IS SIGNED WORD.  
DEFINE FIELD RDB$LM_RECORD_VERSION DATATYPE IS SIGNED WORD.
```

The following example shows the generated table definition for SQL. In this case, both the columns `RDB$LM_DATA_LEN` and `RDB$LM_RECORD_VERSION` are described as `SMALLINT` because SQL doesn't have a 16 bit unsigned integer.

```
-- Table definition for LogMiner transaction data 12-MAY-2020 11:47:19.70  
-- From database table "SAMPLE2"  
CREATE TABLE RDB_LM_SAMPLE2 (  
    RDB$LM_ACTION CHAR  
    ,RDB$LM_RELATION_NAME CHAR (31)  
    ,RDB$LM_RECORD_TYPE INTEGER  
    ,RDB$LM_DATA_LEN SMALLINT  
    ,RDB$LM_NBV_LEN SMALLINT  
    ,RDB$LM_DBK BIGINT  
    ,RDB$LM_START_TAD DATE VMS  
    ,RDB$LM_COMMIT_TAD DATE VMS  
    ,RDB$LM_TSN BIGINT  
    ,RDB$LM_RECORD_VERSION SMALLINT  
    ,IDENT INTEGER  
    ,COMMENT CHAR (1000)  
    ,DETAILS CHAR (64250)  
);
```

Applications that provide a callable image to receive the extracted table rows should define the header fields as shown below.

```
#include <ints.h>  
  
#pragma member_alignment __save  
#pragma nomember_alignment  
  
typedef struct { /* LogMiner structure Header */  
    char    rdb$lm_action;  
    char    rdb$lm_relation_name [31];  
    int     rdb$lm_record_type;  
    unsigned short rdb$lm_data_len;  
    short    rdb$lm_nbv_len;  
    dbk      rdb$lm_dbk;  
    int64    rdb$lm_start_tad;  
    int64    rdb$lm_commit_tad;  
    int64    rdb$lm_tsn;  
    unsigned short rdb$lm_record_version;  
} lmrHeader;  
  
#pragma member_alignment __restore
```

In prior versions of Oracle Rdb, the `RMU Unload After_Image` command would incorrectly process these fields as signed and therefore display values as negative values in the `DELIMITED_TEXT`, `DUMP` and `TEXT` format. This has been corrected in Oracle Rdb Release 7.4.1.0.

6.3.2 Unexpected RMU Unload After Bugcheck at AIJEXT\$QSORT_RELEASE_REC

Bug 31191945

In prior releases of Oracle Rdb, it was possible for the `RMU Unload After` command to bugcheck when processing rows for a table with long record sizes (greater than 32767 bytes).

This problem has been corrected in Oracle Rdb Release 7.4.1.0.

Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.4

7.1 Enhancements

7.1.1 OPCOM Messages Now Always Entered in Oracle Rdb Monitor Log File

In prior releases, any events that may have been broadcast by the OpenVMS OPCOM facility would be lost if the database was not configured with NOTIFY IS ENABLED option. With this update, any message that would have been sent to OPCOM is now entered into the Oracle Rdb Monitor log file, regardless of the NOTIFY setting.

In some system environments, this could lead to a much higher growth rate in the size of the monitor log file, especially if features such as database security auditing are enabled. Oracle advises customers to monitor the size of the monitor log file, and if it becomes excessive, implement a procedure to periodically re-open the log and purge older logs. That is, issue the RMU Monitor Reopen command.

7.1.2 RMU Show Statistics Now Adds Lock ID to STALL_INVOKE Subprocess Parameters List

Bug 37516950

The RMU Show Statistics utility provides the ability to call a subprocess if a stall has exceeded the seconds threshold provided on the Alarm qualifier. If a process is waiting for a lock request, it would be helpful if the lock identifier could be passed to the subprocess so that the lock can be examined. For example, the lock could be examined and the user that is blocking other users could then be identified.

With this update, a lock identifier is an additional parameter given to the subprocess. The parameters are now:

- P1 = The database filename
- P2 = The current date and time
- P3 = The process identification (PID) for the stalled user
- P4 = The database stream identifier for the stalled user
- P5 = The value of the alarm threshold in seconds
- P6 = The stall message showing why the user is stalled
- P7 = If the user is waiting on a lock, the lock identifier for the lock, otherwise zero ("00000000")

- P8 = reserved for future use by Oracle

To implement this functionality, you need to do two things:

1. Create a configuration file for RMU Show Statistics. The file will include an entry for STALL_INVOKE that specifies a command to execute when a long stall is detected.
2. Create a command file to be invoked by RMU Show Statistics. The command file will be passed the arguments listed above in P1 - P7.

Please refer to the *Oracle RMU Reference Manual Appendix D: Configuration File Management and User Defined Events*.

A sample configuration file would look like:

```
ALARM = 60;
STALL_INVOKE = "@DISK: [USER] STALL.COM";
```

The command file should include something like the following line of DCL:

```
$ IF P7 .NES. "00000000" THEN RMU/SHOW LOCK/LOCK='P7 /OUTPUT=LOCK.TXT
```

The content of the output file, LOCK.TXT, can then be parsed to determine the PID of the user that is blocking other users.

The RMU Show Statistics command should include the Configure qualifier that specifies the configuration file specification and usually the NoInteractive qualifier so that the command can be run in the background to monitor the stalls.

7.1.3 Enhanced Integrity Checking Before Page Writes

Enhancement 37985413

Additional checks have been added to prevent updated data pages from being written to disk if the area number or page number in the page header have been corrupted. Similarly, the area and page number for the before and after journals (.ruj and .aij) are validated before the journal entries are written. If an error is encountered, then the process will bugcheck and exit.

Such circumstances should be rare, but are added to protect the database. Please report all bugcheck dumps to Oracle Support for further analysis.

7.2 Obsolete Features

None for this release.

Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.3

8.1 Enhancements

8.1.1 The Lock_timeout Qualifier can now be Used With All RMU Reclaim Qualifiers

Previously, the RMU Reclaim Lock_timeout qualifier could only be specified if the RECLAIM Free_pages qualifier was also specified. With this release of Oracle Rdb the Lock_timeout qualifier can be specified for all RMU Reclaim command functions.

If another user is already accessing the current storage area or logical area to be processed the RMU Reclaim command will wait until the other user releases the storage area or logical area, or a lock wait timeout occurs. If no default lock timeout has been defined for the database or no current lock timeout is in effect, or the user wants to specify a lock timeout value to be used just during the execution of the current RMU Reclaim command, the Lock_timeout qualifier can be specified. If a lock wait timeout occurs, a **warning** message will be output that the storage area or logical area could not be processed because of a lock conflict with another user and a **warning** status will be returned by the RECLAIM command. That storage area or logical area will not be processed, but the RECLAIM command will continue processing the next storage area or logical area in the specified list of storage areas or logical areas.

To specify a temporary lock timeout value in seconds that will be in effect only during the execution of the current RMU Reclaim command the syntax is

```
/LOCK_TIMEOUT=seconds
```

The value that must be specified is the maximum time in seconds during which the current RMU Reclaim command will wait to acquire an EXCLUSIVE UPDATE lock on the current storage area or logical area to be processed when accessing an on-line database with other users.

If the Lock_timeout qualifier is not specified the lock wait time value used is the value of the logical name RDM\$BIND_LOCK_TIMEOUT_INTERVAL (if defined), or the LOCK TIMEOUT INTERVAL specified by the SQL CREATE DATABASE or ALTER DATABASE statement. If neither value has been specified and the LOCK_TIMEOUT qualifier is not specified the RMU Reclaim command will wait indefinitely to acquire an EXCLUSIVE UPDATE lock on the current storage area or logical area to be processed.

If `Lock_timeout=0` is specified, the RMU Reclaim command will ignore any lock timeout defaults that may be in effect and wait indefinitely to acquire an EXCLUSIVE UPDATE lock on the current storage area or logical area to be processed.

In the following example, warning messages are output, even if Log is not specified, in the RMU Reclaim commands if there is a lock conflict with another user that cannot be resolved. The optional `Lock_timeout` qualifier is specified to allow the user to specify a temporary lock timeout value to be used for the current RMU Reclaim command.

```
$ RMU/RECLAIM/NOLOG/AREA=AREA1/LOCK_TIMEOUT=600 SAMPLE.RDB
%RMU-W-RCLMARNOTPRC, Area AREA1 could not be processed due to a lock conflict
$
$ RMU/RECLAIM/LOG/LAREA=SAMPLE_TABLE/FREE_PAGES/LOCK_TIMEOUT=1200 -
$ SAMPLE.RDB
%RMU-I-RCLMLAREA, Reclaiming logical area SAMPLE_TABLE in physical
area DISK:[DIRECTORY]AREA1.RDA;1
%RMU-W-RCLMLARNOTPRC, Logical area SAMPLE_TABLE could not be processed
due to a lock conflict
%RMU-I-RCLMLPAGPRC, 0 pages processed for logical area SAMPLE_TABLE in
physical area DISK:[DIRECTORY]AREA1.RDA;1
%RMU-I-RCLMLPAGFREED, 0 clump pages freed for logical area
SAMPLE_TABLE in physical area DISK:[DIRECTORY]AREA1.RDA;1
$
```

8.1.2 RMU Collect Optimizer_Statistics UNIQUE Index Performance Enhancement

Bug 3804717

A performance enhancement has been added to the RMU Collect Optimizer_Statistics and RMU Show Optimizer_Statistics commands. If only cardinality is being collected and the index is unique with a single index segment, that index will not be analyzed to collect the index cardinality because the index cardinality is always derived from the base table. The index will be analyzed in all other cases.

If the Log qualifier is specified for the RMU Collect Optimizer_Statistics or Show Optimizer_Statistics commands a message will be output to indicate that the index did not have to be analyzed to collect index cardinality statistics.

```
$ RMU/COLLECT OPTIMIZER MF_PERSONNEL/STAT=CARDINALITY/LOG
Start loading tables... at 12-JAN-2024 10:44:50.49
Done loading tables.... at 12-JAN-2024 10:44:50.51
Start loading indexes... at 12-JAN-2024 10:44:50.51
%RMU-I-COLTXT_04, Unique Index COLL_COLLEGE_CODE for Table COLLEGES
has same cardinality as table, cardinality not collected
%RMU-I-COLTXT_04, Unique Index DEPARTMENTS_INDEX for Table DEPARTMENTS
has same cardinality as table, cardinality not collected
%RMU-I-COLTXT_04, Unique Index EMPLOYEES_HASH for Table EMPLOYEES has
same cardinality as table, cardinality not collected
```

```
%RMU-I-COLTXT 04, Unique Index EMP EMPLOYEE_ID for Table EMPLOYEES has
  same cardinality as table, cardinality not collected
Done loading indexes... at 12-JAN-2024 10:44:50.56
Start collecting btree index stats... at 12-JAN-2024 10:44:50.76
Done collecting btree index stats... at 12-JAN-2024 10:44:50.77
Start collecting table & hash index stats... at 12-JAN-2024 10:44:50.77
Done collecting table & hash index stats... at 12-JAN-2024 10:44:50.79
Start calculating stats... at 12-JAN-2024 10:44:50.81
Done calculating stats... at 12-JAN-2024 10:44:50.81
Start writing stats... at 12-JAN-2024 10:44:50.84
```

8.1.3 New Partitions Qualifier for RMU Unload Command

Enhancement 34940446

This release of Oracle Rdb supports a new Partitions qualifier as part of RMU Unload. The Partitions qualifier modifies the table being selected to restrict access to just those named partitions.

This change to RMU allows large partitioned tables to be unloaded by physical area, rather than selecting the whole table, or by unloading through a view that identifies selected rows.

The following example shows the names of the partitions, using SQL SHOW STORAGE MAP and several RMU Unload commands to unload each partition into a separate data file.

```
$ sql$
attach 'filename mf_personnel';
show storage map (partition) employees_map;
      EMPLOYEES_MAP
For Table:      EMPLOYEES
Placement Via Index:  EMPLOYEES_HASH
.
.
.
Partition information for storage map:
Comment:      employees partitioned by "00200" "00400"
Compression is:  ENABLED
Partition: (1) SYS_P00079
Storage Area: EMPIDS_LOW
Partition: (2) SYS_P00080
Storage Area: EMPIDS_MID
Partition: (3) SYS_P00081
Storage Area: EMPIDS_OVER
disconnect all;
$
$ rmu/unload mf_personnel employees/partition=(SYS_P00079) SYS_P00079.unl
%RMU-I-DATRECUNL, 37 data records unloaded 17-JAN-2024 09:42:13.45.
$ rmu/unload mf_personnel employees/partition=(SYS_P00080) SYS_P00080.unl
%RMU-I-DATRECUNL, 57 data records unloaded 17-JAN-2024 09:42:13.54.
$ rmu/unload mf_personnel employees/partition=(SYS_P00081) SYS_P00081.unl
%RMU-I-DATRECUNL, 6 data records unloaded 17-JAN-2024 09:42:13.64.
$
```

See also the Transaction_Type option PARTITION_LIST which can be used to reserve partitions for PROTECTED or EXCLUSIVE access to reduce locking during unload. See Section 8.1.4, New PARTITION_LIST Option for the Transaction_type Qualifier for more details.

Usage Notes

- One or more partition names specified as a comma separated list.

```
$ rmu/unload -  
mf_personnel -  
employees/partitions=(SYS_P00079, SYS_P00081) -  
saved_rows.unl
```

If **Partitions=*** is specified then this is synonymous with referencing the base table. If the **Partitions** qualifier lists all the partitions of the table then this is synonymous with **Partitions=***.

- When **Partitions** is used the subset table is treated as **READ ONLY** and therefore **Delete_rows** is not permitted.

```
$ rmu/unload -  
/delete_rows -  
mf_personnel -  
employees/partitions=(*) -  
saved_rows.unl  
%DCL-W-CONFLICT, illegal combination of command elements - check  
documentation  
\DELETE_ROWS\  
$
```

- When **Partitions** is used the indices on the table are not used by the unload query so the table will always be scanned using **SEQUENTIAL ACCESS**.
- The partition names are defined by the storage map for the table. When the storage map is created or altered a partition name may be provided for each storage area as part of the **IN** clause, or Rdb will generate a unique name using the logical area number assigned to each partition.

The Oracle Rdb default name starts with **SYS_P**, such as **SYS_P00079**. Oracle recommends that more descriptive names be assigned to the storage map partitions. The existing names can be altered using the **RENAME PARTITION** clause of **ALTER STORAGE MAP**.

```
SQL> alter storage map JOB_HISTORY_MAP  
cont> rename partition SYS_P00089 to EMPIDS_LOW;  
SQL> alter storage map JOB_HISTORY_MAP  
cont> rename partition SYS_P00090 to EMPIDS_MID;  
SQL> alter storage map JOB_HISTORY_MAP  
cont> rename partition SYS_P00091 to EMPIDS_OVER;  
SQL> commit;
```

Using the **PARTITION** clause when using the **CREATE STORAGE MAP**, and the **ALTER STORAGE MAP** statements allow the database administrator to provide meaningful names to simplify table management and partition references.

8.1.4 New **PARTITION_LIST** Option for the **Transaction_type** Qualifier

This release of Oracle Rdb includes the following new option on the **RMU Unload Transaction_type** qualifier.

- **PARTITION_LIST=(partname1 [...])**

This option accepts a list of partition names for the table being unloaded.

This option is compatible with the EXCLUSIVE, PROTECTED and SHARED options.

This example uses this definition of the EMPLOYEES_MAP

```
SQL> show storage map employees_map
      EMPLOYEES_MAP
For Table:      EMPLOYEES
Placement Via Index:  EMPLOYEES_HASH
Partitioning is:  UPDATABLE
Store clause:      Rename partition SYS_P00079 to EMPIDS_LOW
                  Rename partition SYS_P00080 to EMPIDS_MID
                  Rename partition SYS_P00081 to EMPIDS_OVER

Partition information for storage map:
Compression is:  ENABLED
Partition: (1) EMPIDS_LOW
Storage Area: EMPIDS_LOW
Partition: (2) EMPIDS_MID
Storage Area: EMPIDS_MID
Partition: (3) EMPIDS_OVER
Storage Area: EMPIDS_OVER
SQL>
```

The RMU Unload command in this example unloads a single partition EMPIDS_LOW using the Partitions qualifier and specifies EXCLUSIVE READ for that partition using the Transaction_type qualifier.

```
$ rmu/unload mf_personnel -
  /transaction=(exclusive, partition_list=empids_low) -
  employees/partitions=(empids_low) -
  empids_low.unl
%RMU-I-DATRECUNL, 37 data records unloaded ...
$
```

8.1.5 Oracle Rdb Compilers Now Set the Object Module Severity

This release of Oracle Rdb brings an enhancement to the SQL Module Language compiler and the SQL Precompiler. In line with other OpenVMS compilers these compilers from Oracle Corporation now also establish the severity in the generated object module. This is especially important when the compile results in generated SQL warning messages.

The following example shows that now such warnings also get detected by the LINK command.

```
$      sql$pre /cc /list-
        /object=TEST4.obj -
        TEST4.sc
        finish;      // disconnect default;
        1
%SQL-I-DEPR FEATURE, (1) Deprecated Feature: FINISH is replaced by DISCONNECT
        insert into work_status (status_code) values (user);
        1
%SQL-W-LENMISMAT, (1) Truncating right hand side string for assignment to column STATUS_CODE
$
$      link TEST4,sql$user/library
%LINK-W-WRNNERS, compilation warnings
        in module MM file USERS_DISK:[TESTING]TEST4.OBJ;1
$
```


Oracle recommends that you correct the SQL sources that generate warnings, or disable the warnings using /NOWARNINGS for the SQL Module Language compiler and /SQLOPTIONS=NOWARN for the SQL Precompiler.

You can use the ANALYZE/OBJECT DCL command to examine this setting.

- On an Alpha system this attribute is shown in the END OF MODULE section.

```
26.  END OF MODULE (EOBJ$C_EEOM), 10 bytes
      severity: warnings (1)
      Highest conditional linkage index: 40
```

- On an IA64 system it appears in the Elf Header Information.

Elf Header Information, at file address 0.

Class:	ELF-64
Data:	Little-endian byte order
Elf Header Version:	1.
OS ABI:	OpenVMS
OS ABI Version:	2.
Type:	Object
Machine Architecture:	IA_64
Elf File Version:	1.
VMS Completion Code:	WARNING

8.1.6 New ALTER DATABASE ... REBUILD SYSTEM INDEX Clause

Bugs 5022392, 28854908, 35925682

This release of Oracle Rdb introduces a new REBUILD SYSTEM INDEX clause as part of ALTER DATABASE. This feature allows a database administrator to rebuild selected Rdb system table indices. Reasons to do this might be to improve retrieval due to restructuring or when instructed to do so by Oracle Support.

Syntax

REBUILD SYSTEM INDEX index-name

This example rebuilds the RDB\$INTER_OBJ_SUBOBJ_NDX index that is defined upon the RDB\$INTERRELATIONS table.

```
$ sql$
alter database
  filename MF_PERSONNEL
  rebuild system index RDB$INTER_OBJ_SUBOBJ_NDX
;
```

Usage Notes

- The ALTER DATABASE statement may only specify the REBUILD SYSTEM INDEX clause once. Use multiple ALTER DATABASE statements to rebuild multiple indices.
- The REBUILD SYSTEM INDEX clause requires exclusive access to the system metadata and is therefore an OFFLINE clause. The change is implicitly committed.
- Because ALTER DATABASE can not be rolled back Oracle recommends that a full RMU BACKUP be taken before executing this command.
- The REBUILD SYSTEM INDEX clause is incompatible with ADD CACHE, ADD JOURNAL, ADD STORAGE AREA, ALTER CACHE, ALTER JOURNAL, ALTER STORAGE AREA, DROP CACHE, DROP JOURNAL, DROP STORAGE AREA, and RESERVE clauses. Specify these clauses in a separate ALTER DATABASE statement.
- The index name provided to the REBUILD SYSTEM INDEX clause must be one of the defined system table indices. Use the SQL SHOW SYSTEM TABLES statement to see the name of the system tables.

```
SQL> show system tables
System tables in database with filename SQL$DATABASE
RDB$COLLATIONS
RDB$CONSTRAINTS
RDB$CONSTRAINT_RELATIONS
RDB$DATABASE
RDB$FIELDS
RDB$FIELD_VERSIONS
RDB$GRANTED_PROFILES
RDB$INDEX_SEGMENTS
RDB$INDICES
RDB$INTERRELATIONS
.
.
.
SQL>
```

Use the SQL SHOW TABLE (INDEX) statement to display the names of the system indices for a table.

```
SQL> show table (index) RDB$INTERRELATIONS
Information for table RDB$INTERRELATIONS

Compression is disabled.

Indexes on table RDB$INTERRELATIONS:
RDB$INTER_CON_NAM_NDX          with column RDB$CONSTRAINT_NAME
  Duplicates are allowed
  Type is Sorted Ranked
    Duplicates are Compressed Bitmaps
  Key suffix compression is DISABLED

RDB$INTER_ENTITY_NDX          with column RDB$ENTITY_NAME1
                              and column RDB$ENTITY_NAME2
  Duplicates are allowed
  Type is Sorted Ranked
    Duplicates are Compressed Bitmaps
  Key suffix compression is DISABLED
```

```

RDB$INTER_OBJ_SUBOBJ_NDX          with column RDB$OBJECT_NAME
                                   and column RDB$SUBOBJECT_NAME

  Duplicates are allowed
  Type is Sorted Ranked
  Duplicates are Compressed Bitmaps
  Key suffix compression is DISABLED

```

- Oracle recommends using the RDMS\$SET_FLAGS logical name defined to the value "INDEX_STATS" prior to running interactive SQL. This will enable tracing during execution of the rebuild.

```

$ define/user RDMS$SET_FLAGS "INDEX_STATS"
$ sql$
alter database
  filename MF_PERSONNEL
  rebuild system index RDB$INTER_OBJ_SUBOBJ_NDX
;
.
.
.
$

```

- A small subset of the Rdb system indices are required during the rebuild action and may not be rebuilt using the REBUILD SYSTEM INDEX clause.

8.1.7 New Implementation of the SQL IN Predicate

In prior releases of Oracle Rdb SQL would expand the IN predicate into a series of equivalent OR and equal (=) conditions. In this release of Oracle Rdb a new compact format is used to encode and transmit queries from the client to the server. This format significantly reduces the size of the queries, triggers, constraints and procedure definitions that use the IN predicate, with related reductions in processing times for the query compiler. This new format also allows several opportunities for new and future optimizations. See *Section 12.2.1, New Optimizations for the SQL IN Predicate* for more details and examples of the optimizations available in this release of Oracle Rdb.

Usage Notes

- The generated query has a different structure from that generated in prior versions. Any query outlines that are activated by matching the query id value will no longer match. Customers are advised to recreate the query outlines to match the new query id. If applications use the OPTIMIZE USING clause then that name will still be used to locate the appropriate outline and no changes should be required.
- The current releases of Oracle CDD Repository, Replication Option for Rdb, and Oracle TRACE do not support this new format. This support is forthcoming for these products (V7.4.1). Until such updated versions are installed Oracle recommends that customers define this logical name system wide.

```

$ define /system SQL$DISABLE_IN_CLAUSE "Y"

```

Use RMU/Show Logical_name SQL\$DISABLE_IN_CLAUSE /Description to read more information.

- Applications that query via the remote interface to older Oracle Rdb releases (7.4.1.2 and earlier) will adjust to the capabilities of that version when using Interactive SQL, or Dynamic SQL.

However, if the SQL Module Language compiler or the SQL Precompiler are used to compile applications using the current version of Oracle Rdb but run against a different run-time database on an older Rdb version then it is possible that the IN predicate will fail to execute. Define this logical name prior to compiling any source modules.

```
$! TEST A: go remote to see that the support fails
$! >>> expect: %RDB-E-INVALID_BLR, request BLR is incorrect at offset ...
$ SQL$PRE/CC-
    /OBJECT=SCC_NEW_IN_CLAUSE -
    SCC_NEW_IN_CLAUSE
$ LINK SCC_NEW_IN_CLAUSE,SQL$USER/LIBRARY
$
$ define/user TARGET_DATABASE SCC_NEW_IN_CLAUSE_DB_REM
$ RUN SCC_NEW_IN_CLAUSE
%RDB-E-INVALID_BLR, request BLR is incorrect at offset 226
$
$! TEST B: go remote to see that it now work after compiling with disabled clause
$ define/user SQL$DISABLE_IN_CLAUSE "Y"
$ SQL$PRE/CC-
    /OBJECT=SCC_NEW_IN_CLAUSE -
    SCC_NEW_IN_CLAUSE
$ LINK SCC_NEW_IN_CLAUSE,SQL$USER/LIBRARY
$
$ define/user TARGET_DATABASE SCC_NEW_IN_CLAUSE_DB_REM
$ RUN SCC_NEW_IN_CLAUSE
last-name -> Canonica
last-name -> Dietrich
last-name -> Gehr
last-name -> Johnson
last-name -> O'Sullivan
last-name -> Tarbassian
last-name -> Wood
last-name -> Ziemke
8 rows fetched.
$
```

Applications compiled with prior versions of Oracle Rdb (7.4.1.2 or earlier) will still use the older query format. Such applications only need to be recompiled to take advantage of the new optimizations.

8.1.8 Updated Sample Code for RDB_CYPHER

This release of Oracle Rdb includes a new sample program RDB_CYPHER. It is written in the C language and uses a different approach to using Rdb modules to provide context for many calls to the OpenVMS Encryption services.

Note

This program supersedes the original sample program RDB_CYPHER.B32. That older example source code is still provided as reference.

Introduction

This sample program demonstrates calling encryption and description routines provided by the OpenVMS ENCRYPT library, but could also be adapted to locally written, or open source routines.

Note

Please refer to the *VSI OpenVMS Utility Routines Manual, Chapter 11 Encryption (ENCRYPT) Routines* for more details on the routines referenced in the sample SQL external function definitions.

The Encryption Routines call interface is also described by DCL HELP ENCRYPT_Routine.

The sample program is described in these sections:

1. Choosing the interface.
2. Implementing the external module and routines.
3. Compiling, linking and deploying this module.
4. Important notes for the implementation.

Oracle recommends a study of the encryption services be made before reading further.

Choosing the Interface

OpenVMS provides two routines to encrypt and decrypt a single field; ENCRYPT\$ENCRYPT_ONE_RECORD and ENCRYPT\$DECRYPT_ONE_RECORD. These routines are called to encrypt each table column values, and subsequently decrypt those column values for display or other computation.

Two external routines are defined to accept input and call the appropriate encryption service. Note that these services require the use of a defined key which in this sample code is created by the routine ENCRYPT\$DEFINE_KEY but could also be defined outside the application using DCL commands; ENCRYPT /CREATE_KEY.

For the purposes of this example the routines ENCRYPT\$DEFINE_KEY and ENCRYPT\$DELETE_KEY are called to manage the key.

Implementing the External Module and Routines

Applications that include encryption support are expected to call these routines frequently. Therefore, the calls to manage the key are not part of the per column encryption and decryption processing since they need not be repeated. Instead they are implemented in the external module's NOTIFY routine.

Each external routine definition can reference a routine (that can not be written in SQL) which is called at special intervals during external routine usage. From our example when the database connect (attach) occurs the first activation of routines in the module will execute the CONNECT action (RDB_K_RTX_

NOTIFY_CONN_START¹). This is specified using the NOTIFY clause as part of the EXTERNAL routine definition.

```
notify ENCD_NOTIFY on connect
```

Although Rdb supports standalone external routines for the purposes of this example we need a common context. Therefore, all the external routines are defined within a single module (see the CREATE MODULE Statement in the *Oracle Rdb Server SQL Reference Manual*).

The support code for the external routines will be written in C, and the module will include three routines: ENCD_NOTIFY; the notify routine, ENCRYPT_CCN; routine to encrypt the credit card number, and DECRYPT_CCN; routine to decrypt the credit card number.

Any database intending to call ENCRYPT_CCN and DECRYPT_CCN must have their interfaces defined in the database so that queries can reference them by name and so the parameter can be passed correctly. These SQL CREATE statements are included in the SQLSSAMPLE:RDB_CYPHER.C source code. Extract them and create a SQL command procedure which can be used to define the interface on any databases that need these services.

```
create module ENCRYPTION_SERVICES
  procedure ENCRYPT_CCN (
    in :plain_text_dsc      char (20) character set unspecified by descriptor,
    out :cypher_text_dsc    char (24) character set unspecified by descriptor);
  external
  name ENCRYPT_CCN
  location 'SYS$SHARE:RDB_CYPHER.EXE' with all logical_name translation
  language C
  parameter style GENERAL
  notify ENCD_NOTIFY on connect
  comment is 'Encrypt the CCN using a predefined key'
;
  procedure DECRYPT_CCN (
    in :cypher_text_dsc      char (24) character set unspecified by descriptor,
    out :plain_text_dsc      char (20) character set unspecified by descriptor);
  external
  name DECRYPT_CCN
  location 'SYS$SHARE:RDB_CYPHER.EXE' with all logical_name translation
  language C
  parameter style GENERAL
  notify ENCD_NOTIFY on connect
  comment is 'Decrypt the CCN using a predefined key'
;
end module;
```

Once created the module should be granted (using the SQL GRANT statement) appropriate protections on the module as required by the security policies of the organization. Note that individual routines in a module do not have separate protections.

¹ See SYS\$LIBRARY:DSRI.H

Compiling, Linking and Deploying this Module

The sample module must be compiled and linked. The resulting executable image should then be copied to SYS\$SHARE. The image should also be installed to ensure correct and valid usage.

```
$ CC RDB_CYPHER.C
$ LINK /SHARE RDB_CYPHER.OBJ+SYS$INPUT:/OPTION
symbol_vector=( -
    ENCRYPT_CCN=procedure, -
    DECRYPT_CCN=procedure, -
    ENCD_NOTIFY=procedure)
$ COPY RDB_CYPHER.EXE SYS$COMMON:[SYSLIB] /PROT=W:RE
$ INSTALL REPLACE SYS$SHARE:RDB_CYPHER.EXE /OPEN/SHARE/HEAD
```

Now when SQL queries, triggers or stored procedures execute the ENCRYPT_CCN and DECRYPT_CCN routines the shareable image will be activated from a known secure location.

Important Notes for the Implementation

This C module is supplied solely as a template example of calling the OpenVMS encryption routines from an Rdb stored procedure. The expected use of this module would be to encrypt and decrypt a 20 byte text field containing a credit card number.

You are free to use this example code and extend it; adding more routines for different sized fields, changing the key definition module, and so on.

Further considerations:

- How is NULL handled? External routines do not accept NULL values so the interface might need a jacket SQL routine that passes the nullness of the data as a second parameter.
- If columns include duplicate data for the same table it might be more secure to add salting (randomized values) to the plain text prior to encryption to obscure any patterns. Upon decrypting that randomized data is not returned to the caller.
- This example only uses a single key for inception but different columns could have different encryption keys. That may allow, for instance, different users to decrypt a subset of the data based on the application requirements.
- Protect the KEY. If the key is lost then the encrypted data is no longer accessible, so don't lose it. Also protect the security of the encryption key. Coding it as plain text in this source file is for example purposes only.

8.2 Obsolete Features

None for this release.

Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.2

9.1 Enhancements

9.1.1 Substring Function Changes

This release of Oracle Rdb provides native versions of the functions SUBSTR and SUBSTRB which were previously provided as external functions in the SQL_FUNCTIONSnn.SQL script.

Note

Those external functions will remain defined in the SQL_FUNCTIONS74.SQL function script to support applications that currently use them, or wish to continue to use those functions.

These new native functions have the following advantages over the external functions:

- These functions are now part of the Oracle Rdb SQL language and no external functions for SUBSTR or SUBSTRB need be defined in the database.
- These functions are polymorphic; meaning they adapt to the passed character string - both length and character set. Therefore, the result will be derived from the length of the input string length and the length field (if specified).
The SQL external functions defined in SQL_FUNCTIONS74.SQL were typically limited to CHAR and VARCHAR inputs of 2000 characters, and returned a VARCHAR (2000) typed result even when operating on shorter string inputs.
- These builtin functions now accept CHAR and VARCHAR strings up to 65535 octets in length (the actual length in CHARACTERS will depend on the character set being used).
- These changes will typically result in much lower virtual memory usage and better run-time performance for queries. In particular applications using ORACLE dialects (such as those using OCI Services for Rdb) should expect a considerable reduction in query size.

Note

Applications created using the SQL Module Language or SQL Pre-compiler in prior releases will continue to use the SQL external functions until those application sources are recompiled.

The substring family of functions operate similarly with these differences:

- SUBSTRING (char_string FROM start_position [FOR substring_length])
This is the ANSI and ISO SQL Database Language implementation. Optional USING CHARACTERS or USING OCTETS clauses adjusts the interpretation of the start_position and substring_length.
- SUBSTR (char_string , start_position [, substring_length])
For SUBSTR start_position and substring_length represent units as characters. SUBSTR is similar in operation to SUBSTRING with the USING CHARACTERS clause with these extra rules.
 - * If start_position is 0, then it is treated as 1.
 - * If start_position is positive, then Oracle Rdb counts from the beginning of char_string to find the first character of the result.
 - * If start_position is negative, then Oracle Rdb counts backward from the end of char_string. For example, to return the last 3 characters of a string you can pass start_position as -3 and the substring_length as 3. If the computed beginning of char_string is prior to the start of the string then a zero length string is the result.
- SUBSTRB (char_string , start_position [, substring_length])
For SUBSTRB start_position and substring_length represent units as octets. SUBSTRB is similar in operation to SUBSTRING with the USING OCTETS clause with the same rules as described for SUBSTR.

If the dialect is set to an Oracle dialect (ORACLE LEVEL1, ORACLE LEVEL2, or ORACLE LEVEL3) these functions will return NULL if the resulting string was zero length.

Changes in Query Signature related to Query Outlines

The database administrator should be aware that queries using SUBSTR or SUBSTRB generated by SQL will be differently structured from queries using the SQL external functions even if the results remain the same.

Similarly, the SUBSTRING builtin function now directly supports the Oracle dialect semantics which consider a zero length string result as being equivalent to NULL. So when using SUBSTRING under the Oracle dialects (ORACLE LEVEL1, ORACLE LEVEL2, and ORACLE LEVEL3) the actual query will be changed.

Therefore, any query outline previously created for such queries will need to be redefined. The generated query id (signature) changes even if the strategy does not.

Consider this example which shows that under Oracle Rdb V7.3 SUBSTRING queries for Oracle Dialects were rewritten with CASE expressions to ensure the correct result semantics.

```

SQL> select employee_id, last_name, first_name
cont> from employees
cont> where substring (last_name from 2 for 1) = '''
cont> ;
~S: Outline "QO_0_73_MODE_3" used
Tables:
  0 = EMPLOYEES
Leaf#01 FFirst 0:EMPLOYEES Card=100
  Bool: CASE (WHEN (CHAR_LENGTH (SUBSTRING (0.LAST_NAME FROM 2 FOR 1)) = 0)
    THEN NULL ELSE SUBSTRING (0.LAST_NAME FROM 2 FOR 1)) = '''
  BgrNdx1 EMP_LAST_NAME [0:0] Fan=12
  Bool: CASE (WHEN (CHAR_LENGTH (SUBSTRING (0.LAST NAME FROM 2 FOR 1)) = 0)
    THEN NULL ELSE SUBSTRING (0.LAST_NAME FROM 2 FOR 1)) = '''
EMPLOYEE_ID  LAST_NAME      FIRST_NAME
00171        D'Amico         Aruwa
00190        O'Sullivan      Rick
2 rows selected
SQL>

```

In this case a query outline QO_0_73_MODE_3 was defined (and was used) which matched the query signature. However, when run under this release of Oracle Rdb with a simplified query the outline is no longer matched.

```

SQL> select employee_id, last_name, first_name
cont> from employees
cont> where substring (last_name from 2 for 1) = '''
cont> ;
Tables:
  0 = EMPLOYEES
Leaf#01 FFirst 0:EMPLOYEES Card=101
  Bool: SUBSTRING (0.LAST NAME FROM 2 FOR 1) = '''
  BgrNdx1 EMP_LAST_NAME [0:0] Fan=12
  Bool: SUBSTRING (0.LAST_NAME FROM 2 FOR 1) = '''
EMPLOYEE_ID  LAST_NAME      FIRST_NAME
00171        D'Amico         Aruwa
00190        O'Sullivan      Rick
2 rows selected
SQL>

```

Either the query could be modified to explicitly use the OPTIMIZE USING clause or a new query outline created to match the new query signature.

9.1.2 Support for EMPTY_BLOB and EMPTY_CLOB in SQL

This release of Oracle Rdb adds the new SQL functions EMPTY_BLOB () and EMPTY_CLOB ().

Purpose

EMPTY_BLOB and EMPTY_CLOB return an empty handle that can be used to initialize a LIST OF BYTE VARYING variable or as a value expression in an INSERT, UPDATE, or REPLACE statement to initialize a LIST OF BYTE VARYING column. An EMPTY indication means that the LIST is initialized, but not populated with data.

Examples

The following example initializes the BADGE_PHOTO column of the sample EMPLOYEES table to EMPTY using the UPDATE statement.

```
SQL> alter table EMPLOYEES
cont>      add column BADGE_PHOTO list of byte varying
cont> ;
SQL>
SQL> update EMPLOYEES
cont>      set BADGE_PHOTO = empty_blob ();
101 rows updated
SQL>
```

The following example initializes the BADGE_PHOTO column of the sample EMPLOYEES table to EMPTY using the INSERT statement.

```
SQL> insert into EMPLOYEES
cont>      values ('99001', 'McAlister', 'James', 'A',
cont>              '33 Running Deer Road', '', 'Hollis', 'NH', '03049',
cont>              'M', date vms'3-May-1989', '2',          /* part-time */
cont>              empty_blob()                          /* no photo yet */
cont>      );
1 row inserted
SQL>
```

9.1.3 New ANY_VALUE Aggregate Function

This release of Oracle Rdb adds the SQL function ANY_VALUE.

Format

```
ANY_VALUE ( [ ALL | DISTINCT ] value_expr ) [ FILTER ( WHERE boolean_expr ) ]
```

Arguments

- **ALL, and DISTINCT:** These keywords are supported by ANY_VALUE although they have no effect on the result of the query and are effectively ignored.
- **value_expr:** The expression can be a column, constant, variable, function or a complex expression.
- **boolean_expr:** the FILTER clause can be used to further filter the rows that participate in the aggregation.

Description

ANY_VALUE is an aggregate function that returns a single non-deterministic value of value_expr.

Use ANY_VALUE to optimize a query that has a GROUP BY clause. ANY_VALUE returns a value of an expression in a group. It is optimized to return the first value.

It ensures that there are no comparisons for any incoming row and also eliminates the necessity to specify every column as part of the GROUP BY clause. Because it does not compare values, ANY_VALUE returns a value more quickly than MIN or MAX in a GROUP BY query.

NULL values in the expression are ignored.

Returns any value within each group based on the ORDER BY specification.
Returns NULL if all rows in the group have NULL expression values.

Example

The following example shows the use of ANY_VALUE. When computing the aggregation of the rows values of SALARY_END which are NULL are excluded. In cases where all rows are excluded the ANY_VALUE function returns NULL.

```
SQL> select
cont>     employee_id, any_value (salary_end), count(*)
cont>     from salary_history
cont>     where employee_id in ('00241', '00227', '00225', '00164')
cont>     group by employee_id
cont> ;
EMPLOYEE_ID          2-Mar-1981          4
00164
00225          NULL          1
00227          NULL          1
00241          NULL          1
4 rows selected
SQL>
```

9.1.4 Updated Diagnostics for the RMU REPAIR Command

With this release of Oracle Rdb the RMU Repair command has been updated to output improved messages when changes are made to the database.

- %RMU-W-FULLAIJBKUP, partially-journaled changes made; database may not be recoverable

This message warns the user that, even though After Image Journaling is enabled, not all structural changes made to the database by the Repair operation are journaled. Therefore, the user should consider backing up the database once the repair operation completes.

This message is only output if After Image Journaling is enabled for the database.

```
$ RMU/REPAIR/ABM MF_PERSONNEL
%RMU-W-FULLAIJBKUP, partially-journaled changes made; database may not
be recoverable
%RMU-W-DOFULLBCK, full database backup should be done to ensure future recovery
```

- %RMU-W-DOFULLBCK, full database backup should be done to ensure future recovery

This message warns the user to do a full database backup after the repair operation completes to ensure that the structural changes made by the Repair operation can be restored by the RMU Restore command and subsequent RMU Recover from database AIJ files.

- %RMU-I-REPAIR_DEF, operation not specified; Reconstructing the SPAM pages for specified areas

This message informs the user that with no directives included on the RMU Repair command it defaulted to the /SPAM operation.

```
$ RMU/REPAIR MF_PERSONNEL
%RMU-I-REPAIR_DEF, operation not specified; Reconstructing the SPAM
  pages for specified areas
%RMU-W-DOFULLBCK, full database backup should be done to ensure future
  recovery
```

9.1.5 Updated Diagnostics for the RMU MOVE_AREA Command

With this release of Oracle Rdb the RMU Move_Area command has been updated to output improved messages when changes are made to the database.

- %RMU-W-FULLAIJBKUP, partially-journalled changes made; database may not be recoverable

This message warns the user that the actual move of database files to another device or directory is not journaled. Therefore, the user should consider performing a full backup of the database once the **move** operation completes.

The RMU Restore of the full database backup will recover the device and directory changes that were made and a subsequent RMU Recover will update the moved areas as needed.

The following example shows the %RMU-W-FULLAIJBKUP message being output for a move operation which moves the database root file of a database with After Image Journaling enabled.

```
$ RMU/MOVE AREA /ROOT=disk:[directory] /NOLOG MF_PERSONNEL
%RMU-W-FULLAIJBKUP, partially-journalled changes made; database may not
  be recoverable
%RMU-I-AIJRSTAVL, 1 after-image journal available for use
%RMU-I-AIJRSTMOD, 1 after-image journal marked as "modified"
%RMU-I-AIJISON, after-image journaling has been enabled
%RMU-W-DOFULLBCK, full database backup should be done to ensure future
  recovery
```

9.1.6 Hash Join Enhancements and Support

Status: BETA

Introduction

Table joins are typically solved using a mix of index and sorting strategies. In some cases there may not be any suitable index to use resulting in the solution selected by the optimizer requiring an extra sorting step.

The use of Hash Join eliminates the need to sort data by using an in-memory hash table to filter the join data.

This feature can be enabled by using any one of the following methods:

- Using the OPTIMIZE FOR HASH JOIN clause
The optimizer will attempt to apply Hash Join on joins within the query. See SQL HELP SELECT for more information.
- Using SET FLAGS 'HASHING(JOINS)' statement

The optimizer will attempt to apply Hash Join on joins within all subsequent queries within this session. Use SET FLAGS 'NOHASHING' to disable this action within the session.

- Define the logical name RDMS\$SET_FLAGS "HASHING(JOINS)"
Similar to using the SET FLAGS statement, defining RDMS\$SET_FLAGS to the value "HASHING(JOINS)" will enable this feature and will be applied to any applications that perform joins. Deassign this logical name (as well as RDMS\$ENABLE_HASH_JOIN, see below) to disable this action within the process.
- Define the logical name RDMS\$ENABLE_HASH_JOIN to true
Defining this logical name to true ("T", "t", "Y", "y" or "1") instructs the Rdb optimizer to attempt to use in-memory Hash Join to solve queries. Deassign this logical name (as well as RDMS\$SET_FLAGS, see above) to disable this action within the process.
Use the RMU /SHOW LOGICAL RDMS\$ENABLE_HASH_JOIN /DESCRIPTION command to see details.
- Create a query outline
The outline should use the JOIN BY HASH clause to adjust specific parts of the query solution. Refer to SQL HELP CREATE OUTLINE for further information on query outlines.
The created query outline will be used when referenced by name in the OPTIMIZE USING clause, or by any query matching the query signature (query outline id).

Oracle Corporation appreciates any and all feedback on this feature.

Changes for This Release

This release of Oracle Rdb includes an improved design of the Hash Join feature with these enhancements:

- Reduced VM usage for Hash Join solution
Copies of the join keys is reduced for duplicate values.
- Reduced CPU usage during Hash Join operation
Previously CPU usage increased as more duplicates were encountered in the hash table entries. The current version reduces the CPU time by no longer re-scanning the duplicate chains.
- Relaxed restrictions when Hash Join is using sequential retrieval
Sequential retrieval is now more widely applied when no indices are defined or available for the join columns.
- Improved FULL OUTER JOIN support
Wrong results related to Hash Join strategy, that occurred with Full Outer Join in prior releases, have been corrected.
- Improved diagnostics reported when Hash Join is rejected
The SET FLAGS statement or the equivalent logical name RDMS\$SET_FLAGS definition can be used to enable the display of these diagnostics. See the Examples.

Current Restrictions

Currently, the use of Hash Join may be rejected by the optimizer for a number of reasons as specified below. The reported reasons for the rejection can be displayed by running the query after setting the flags STRATEGY and DETAIL(2).

1. CHAR or VARCHAR characteristics of Index Columns differ

When CHAR or VARCHAR columns are used during hashing the length, character set and collating sequence of the compared columns must be the same.

Example:

```
~Sh: Hash Join #!2S inner rejected-TAB1_NDX and (generated-index)
      column 0 lengths differ
~Sh: Hash Join #!11S inner rejected-TAB1_NDX and (generated-index)
      column 1 collating sequences differ
~Sh: Hash Join #!3S inner rejected-TAB1_NDX and (generated-index)
      column 2 character sets differ
```

2. Data types in index columns differ

The columns used for hashing must have the same datatypes.

Example:

```
~Sh: Hash Join #5S inner rejected-TAB_NDX and (generated-index)
      column 0 datatypes differ
```

3. Date/time types in index columns differ

If the join is between TIMESTAMP and DATE VMS columns then Oracle Rdb will consider this as non-matching data type.

Example:

```
~Sh: Hash Join #!ULS inner rejected-TAB3_NDX and (generated-index)
      column 2 date/time types differ
```

4. The number of index columns differ

The number of columns to compare in both streams must be the same.

Example:

```
~Sh: Hash Join #6S inner rejected-TAB1_NDX and (generated-index)
      number of columns differ
```

5. Sort orders in index columns differ

The sort ordering for columns (ASC and DESC clauses) must match.

Example:

```
~Sh: Hash Join #!ULS inner rejected-TAB2_NDX2 and (generated-index)
      column 1 sort orders differ
```

6. Integer precision in index columns differ

Hash Join cannot be used if it is required to compare integer columns with different precision.

Example:

```
~Sh: Hash Join #6S inner rejected-TAB1_NDX and TAB2_NDX column 0 scales differ
```

7. OR Retrieval on outer or inner table

Hash Join cannot be used if either stream relies on the results of an OR-retrieval optimization.

Example:

```
~Sh: Hash Join #9S inner rejected-OR retrieval on inner table
~Sh: Hash Join #19S outer rejected-OR retrieval on outer table
```

8. Outer or inner column not equivalent with join columns

Hash Join can only be done using columns that are considered to be in the same equivalence group and thus may be compared.

Example:

```
~Sh: Hash Join #12S outer rejected-X1 column 0 not equivalent with join columns
~Sh: Hash Join #3S inner rejected-(generated-index)
      column 0 not equivalent with join columns
```

9. Index columns not provided as join columns

Hash join using the named index will be rejected if not all index columns are provided values (no partial index key matches are allowed).

Example:

```
~Sh: Hash Join #1S outer rejected-values of all outer
      TAB1_NDX columns not provided
~Sh: Hash Join #16S inner rejected-values of all inner
      TAB2_NDX columns not provided
```

10. Outer or inner index column not a join predicate

Hash Join can only be done when the matching columns for the inner and outer streams take part in the join predicate.

Example:

```
~Sh: Hash Join #15S inner rejected-TAB1_NDX column 1 not a join predicate
```

11. Outer or inner stream not a table

Hash Join requires that streams access a table directly or via an index.

Example:

```
~Sh: Hash Join #11S inner rejected-inner stream not a table
~Sh: Hash Join #14S outer rejected-outer stream not a table
```

12. Outer or inner stream not comply with outlines

Hash Join cannot be done if a stream cannot be optimized in accordance with the supplied outline.

Example:

```
~Sh: Hash Join #4S inner rejected-does not comply with outline
~Sh: Hash Join #12S outer rejected-does not comply with outline
```

13. Query contains aggregate outer join boolean

Hash Join is not available if the outer join relies on a check on the results of an aggregation.

Example:

```
~Sh: Hash Join #3S inner rejected-Q2 contains Aggregate OJ boolean
```

14. Reverse scan on outer or inner table

If either stream uses Reverse Scan then Hash Join is not possible.

Example:

```
~Sh: Hash Join #7S outer rejected-reverse scan on outer table
```

15. Sequential retrieval on outer or inner table not available

If the value to be matched is not derived from an actual or derived column Hash Join cannot utilize Sequential Retrieval

Example:

```
~Sh: Hash Join #4S outer rejected-sequential retrieval on outer table
~Sh: Hash Join #12S inner rejected-sequential retrieval on inner table
```


Examples

This section contains examples of attempts to use Hash Join when restrictions are encountered.

The following query reverts to cross strategy when full compliance with the outline was not possible:

First, create a query outline. This example uses an Rdb generated outline for the query produced using SET FLAGS 'OUTLINE' with the name LOJ_NEST_HSH.

Example 9–1 Example 1: Using OPTIMIZE USING

```
drop outline LOJ_NEST_HSH;
create outline LOJ_NEST_HSH
id 'FE43684748E33D965A065EF6545EB1BB'
mode 0
as (
  query (
    -- For loop
    subquery (
      subquery (
        P 2      access path sequential
      )
      join by cross to
      subquery (
        S 0      access path sequential
        join by hash to
        P 1      access path sequential
      )
    )
  )
)
compliance optional      ;
```

Now execute the query with this query outline.

```
set flags 'strategy,detail(2)';

select snum, s.city, status, pnum, p.city, weight
  from s left join p
    on s.city = p.city and
      (exists (select pnum from p where pnum = 'P0'))
OPTIMIZE USING LOJ_NEST_HSH
;
~S: Outline "LOJ_NEST_HSH" used
~Sh: Hash Join #3S inner rejected-Q2 contains Aggregate OJ boolean
~S: Full compliance with the outline was not possible
Tables:
  0 = S
  1 = P
  2 = P
```

(continued on next page)

Example 9–1 (Cont.) Example 1: Using OPTIMIZE USING

```
Cross block of 2 entries Q1
Cross block entry 1
  Aggregate-F1: 0:COUNT-ANY (<subselect>) Q3
  Conjunct: 2.PNUM = 'P0'
  Get      Retrieval sequentially of relation 2:P  Card=6  JCard=1
Cross block entry 2
  Match      (Left Outer Join)  Inner_TTBL Q2
  Outer loop
  Match Key:0.CITY
  Sort: 0.CITY(a)
  Get Retrieval sequentially of relation 0:S  Card=5  JCard=5
  Inner loop
  Match Key:1.CITY
  Temporary relation
  Sort: 1.CITY(a)
  Get Retrieval sequentially of relation 1:P  Card=6  JCard=10
```

S.SNUM	S.CITY	S.STATUS	P.PNUM	P.CITY	P.WEIGHT
S5	Athens	30	NULL	NULL	NULL
S1	London	20	NULL	NULL	NULL
S4	London	20	NULL	NULL	NULL
S2	Paris	10	NULL	NULL	NULL
S3	Paris	30	NULL	NULL	NULL

5 rows selected

If the query outline definition is changed to use COMPLIANCE MANDATORY then the query using it will fail to run if JOIN BY HASH is not possible.

```
create outline LOJ_NEST_HSH
id 'FE43684748E33D965A065EF6545EB1BB'
mode 0
as (
  query (
    -- For loop
    subquery (
      subquery (
        P 2      access path sequential
      )
      join by cross to
      subquery (
        S 0      access path sequential
        join by hash to
        P 1      access path sequential
      )
    )
  )
)
compliance mandatory ;
```

(continued on next page)

Example 9–1 (Cont.) Example 1: Using OPTIMIZE USING

```
set flags 'strategy,detail(2)';
select snum, s.city, status, pnum, p.city, weight
  from s left join p
    on s.city = p.city and
      (exists (select pnum from p where pnum = 'P0'))
OPTIMIZE USING LOJ_NEST_HSH
;
~S: Outline "LOJ_NEST_HSH" used
~Sh: Hash Join #3S inner rejected-Q2 contains Aggregate OJ boolean
~S: Full compliance with the outline was not possible
%RDB-E-QRYCOMP_FAILED, query or routine failed to compile
-RDMS-F-OUTLINE_FAILED, could not comply with mandatory query outline directives
```

The following query reverts to cross strategy when full compliance with the query outline is not possible.

Example 9–2 Example 2: Using the OPTIMIZE FOR HASH JOIN clause

```
select snum, s.city, status, pnum, p.city, weight
  from s left join p
    on s.city = p.city and
      (exists (select pnum from p where pnum = 'P0'))
OPTIMIZE FOR HASH JOIN
;
~S: Outline "QO_FE43684748E33D96_00000000" used
~Sh: Hash Join abandoned because Q2 contains Aggregate OJ boolean
~S: Full compliance with the outline was not possible
Tables:
  0 = S
  1 = P
  2 = P
Cross block of 2 entries  Q1
Cross block entry 1
  Aggregate-F1: 0:COUNT-ANY (<subselect>) Q3
  Conjunct: 2.PNUM = 'P0'
  Get      Retrieval sequentially of relation 2:P
Cross block entry 2
  Match      (Left Outer Join)  Inner_TTBL Q2
  Outer loop
    Match_Key:0.CITY
    Sort: 0.CITY(a)
    Get Retrieval sequentially of relation 0:S
  Inner loop
    Match_Key:1.CITY
    Temporary relation
    Sort: 1.CITY(a)
    Get Retrieval sequentially of relation 1:P
```

(continued on next page)

Example 9–2 (Cont.) Example 2: Using the OPTIMIZE FOR HASH JOIN clause

S.SNUM	S.CITY	S.STATUS	P.PNUM	P.CITY	P.WEIGHT
S5	Athens	30	NULL	NULL	NULL
S1	London	20	NULL	NULL	NULL
S4	London	20	NULL	NULL	NULL
S2	Paris	10	NULL	NULL	NULL
S3	Paris	30	NULL	NULL	NULL

5 rows selected

9.2 Obsolete Features

None for this release.

Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.1

10.1 Enhancements

10.1.1 Optional Builtin Function RDB\$\$IS_ROW_FRAGMENTED

Oracle Rdb supports an optional builtin function that can determine if a row is fragmented. The function, `RDB$$IS_ROW_FRAGMENTED` must be declared as a function using the attributes and properties as shown below.

```
declare function RDB$$IS_ROW_FRAGMENTED
  (in :dbk char(8) character set unspecified)
  returns integer;
```

The following example shows the usage on the `WORK_STATUS` table in the `PERSONNEL` database.

```
SQL> declare function RDB$$IS_ROW_FRAGMENTED
cont>   (in :dbk char(8) character set unspecified)
cont>   returns integer;
SQL>
SQL> select dbkey, RDB$$IS_ROW_FRAGMENTED (dbkey) from work_status;
          DBKEY
          99:10:12          0
          99:10:13          0
          99:10:14          0
3 rows selected
```

Usage Notes

- This routine may only be used from Interactive and Dynamic SQL.
- Only valid DBKEY values should be passed to the function.
- If the DBKEY passed is not the current row, then additional I/O may be required to fetch the target row.
- If the DBKEY is for a vertically partitioned table, then only the fragmented state of the primary segment is reported. There is currently no programmatic method to determine fragmented secondary segments.
- Temporary table and information table rows are never fragmented as they reside in virtual memory only.
- Fragmentation occurs when either the row is too large to fit entirely on a page or an existing row was updated to a larger size and no space existed at that time for the expanded row. The first case requires that the page size be changed for the area. However, for the second case, a DELETE and INSERT of the row might remove the fragmentation. In that case, this function allows the DBA to identify candidate fragmented rows. Fragmentation may occur when compression is enabled and the compressed row size changes due to changed data, NULL values replaced with non-NULL values, or ALTER


TABLE or ALTER DOMAIN statements that have increased the size of columns.

10.1.2 WITH READ ONLY Clause for CREATE and ALTER VIEW Statements

This release of Oracle Rdb adds support for READ ONLY view definitions, using CREATE VIEW, DECLARE LOCAL TEMPORARY VIEW or ALTER VIEW statements.

Syntax

check-option-clause =



Under normal circumstances, views are considered to be READ ONLY and Rdb will prevent INSERT, UPDATE and DELETE through those views when the select expression uses one of these clauses:

- * Includes the DISTINCT operator to eliminate duplicate rows from the result table
- * Names more than one table or view in the FROM clause
- * Uses a derived table as the row source for a FROM clause
- * Includes an aggregate function in the select list
- * Contains a UNION, EXCEPT, MINUS, INTERSECT, GROUP BY, or HAVING clause

With this release the database administrator can also force a view to be READ ONLY by applying the WITH READ ONLY clause, even if the factors listed above are not true; i.e. a view that would normally be updatable is considered read-only.

This clause (WITH READ ONLY) and the WITH CHECK OPTION clause are mutually exclusive. Any CHECK OPTION constraint previously defined for the view will be deleted when WITH READ ONLY is used. Conversely if the view is altered to successfully add a CHECK OPTION then the READ ONLY attribute is removed.

The following example shows creating a view on a base table and forcing the view to be read-only.

```
SQL> create view SHOW_CURRENT_SALARY
cont> as
cont> select employee_id, salary_amount
cont> from salary_history
cont> where salary_end is null
cont> with read only
cont> ;
```

10.1.3 Enhanced Conversion of Date/time String Literals

In prior releases of Oracle Rdb, string literals assigned to DATE VMS columns, or which were CAST to DATE VMS data types were translated by the SQL interface prior to being passed to the Oracle Rdb Server for execution. This small optimization could reduce or avoid the actual CAST operation.

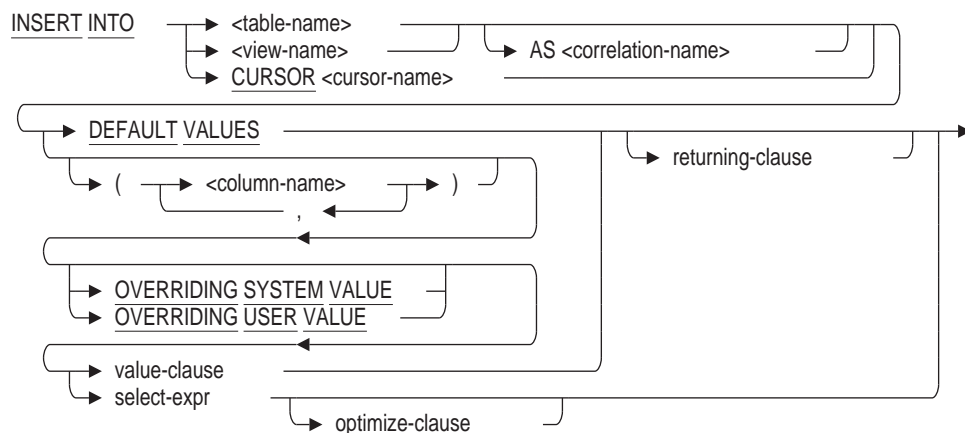
This release of Oracle Rdb expands this support to string literals assigned to DATE ANSI, TIME, TIMESTAMP and INTERVAL data types. If the format of the string is invalid then an error is immediately reported. This is especially a benefit when applications were compiled using the SQL Precompiler or SQL Module Language which now reports the improper format during compile instead of being deferred to runtime.

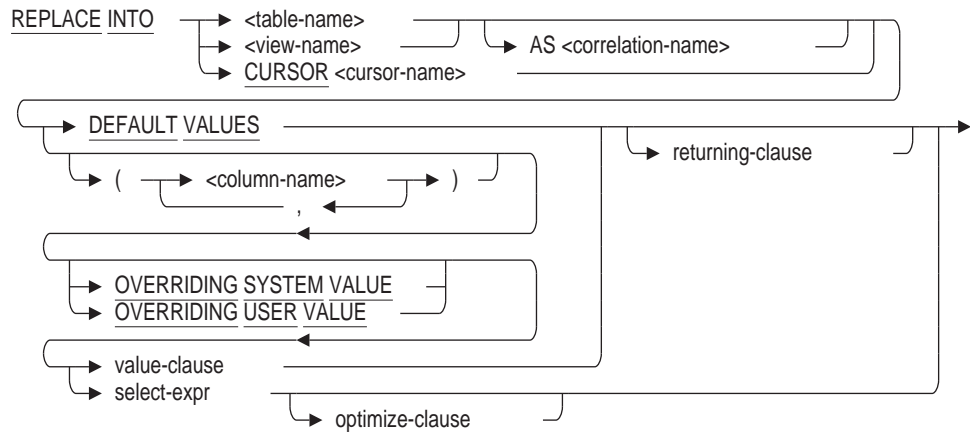
```
SQL> select cast ('1-Jan-2021' as date ansi) from rdb$database;  
%SQL-F-DATCONERR, Data conversion error for string '1-Jan-2021'  
-COSI-F-IVTIME, invalid date or time  
SQL> select cast ('2021-1-1' as date ansi) from rdb$database;  
  
2021-01-01  
1 row selected  
SQL>
```

10.1.4 Support for OVERRIDING Clause in INSERT and REPLACE Statements

This release of Oracle Rdb supports the ANSI and ISO SQL Database Language Standard OVERRIDING clause for the INSERT statement. Oracle Rdb also extends this support to the REPLACE statement. The OVERRIDING USER VALUE and OVERRIDING SYSTEM VALUE clauses affect the handling of inserts to any generated columns during INSERT or REPLACE statements. The OVERRIDING clause appears before the VALUES clause or before the SELECT clause as part of the INSERT and REPLACE statements.

Syntax





- * The OVERRIDING SYSTEM VALUE clause instructs Rdb that the GENERATED, IDENTITY or AUTOMATIC AS columns will be updated with user supplied values and therefore no generated values will be created. Such a clause would be used if a table was being reloaded after maintenance and the database administrator wanted to retain the saved generated values.

Note

The INSERT or REPLACE statements can use the DEFAULT keyword in place of a column value. When the column being updated is a generated or automatic column then the OVERRIDING clause has no effect on that column as it will be the same in either case.

```

replace into SALES_EMPLOYEES (employee_id, last_name, first_name)
  overriding system value
  values (default, 'Myotte', 'Daniel');
  
```

This clause is similar to the SET FLAGS 'AUTO_OVERRIDE' feature.

The following example shows the use of OVERRIDING SYSTEM VALUE in the case of propagating a daily sales table to the yearly accumulated sales table. In this case we don't want the new generated values for these columns as that was already done by the INSERT into the daily sales table.

```

SQL> --> Now we want to perform end-of-day processing.
SQL> set transaction
cont>   read write
cont>   reserving DAILY_SALES for exclusive write,
cont>   YEARLY_SALES for exclusive write;
SQL>
SQL> --> move daily sales
SQL> insert into YEARLY_SALES
cont>   overriding system value
cont>   select * from DAILY_SALES
cont> ;
7 rows inserted
SQL>
SQL> truncate table DAILY_SALES;
SQL>
SQL> commit;
SQL>
  
```


- * The OVERRIDING USER VALUE clause instructs Rdb that the GENERATED, IDENTITY or AUTOMATIC AS columns will be generated by the database system and that any user supplied values will be ignored. Such a clause would be used when column names were wild carded by a SELECT clause and therefore avoids enumerating all non-generated column names.

The following example shows copying all daily sales to a log table which will generate new sales_id values.

```
SQL> --> move daily sales and generate new column values for automatic columns
SQL> insert into SALES_LOG
cont>      overriding user value
cont>      select * from DAILY_SALES
cont> ;
```

10.1.5 RMU Recover Progress_report Qualifier and Ctrl-T Display

In this release a Progress_report qualifier has been added to the RMU Recover command. The Progress_report=n qualifier, where n is the time interval in seconds, displays the performance and progress of the database RMU RECOVER operation at timed intervals to SYS\$OUTPUT.

The same display can also be output whenever Ctrl-T is typed during the RMU Recover operation. Ctrl-T must have been previously enabled at the DCL level using SET CONTROL=T. SET CONTROL=T requires that SET TERMINAL/BROADCAST is enabled for the display terminal.

The RMU Recover performance and progress display has the following format.

- The first line is the file specification of the after image journal (AIJ) currently being recovered.
- The second line displays the number of megabytes that have been read and processed from the AIJ file during the current interval, the percentage of the AIJ file that has been processed, the number of megabytes that are currently being read per second, and the estimated completion time for the processing of this AIJ file.

The Progress_report interval that has been specified in this case is 1 second.

```
DEVICE: [DIRECTORY]TEST_JOURNAL.AIJ;1
Read    21 MB (29%) at    21 MB/s, estimated completion time 11:50:29.80
```

The RMU Recover performance and progress display will be different if the Format=NEW_TAPE qualifier has been specified with the RMU Recover command. This is due to limitations caused by the way AIJ data in this format is processed.

- The first line is the file specification of the current temporary AIJ work file being recovered.
- The second line displays the number of megabytes that have been read and processed from the AIJ work file during this interval, and the number of megabytes that are currently being read per second.

The Progress_report interval that has been specified in this case is 1 second.

```
DEVICE: [DIRECTORY]AIJ_WORKG8V0RS6M99D1GJKG4I80.AIJ;
Read    36 MB at    36 MB/s
```

Syntax

The command line syntax for the RMU Recover command `Progress_report` qualifier is

```
/PROGRESS_REPORT [=seconds]
```

This qualifier cannot be negated and is not the default. The default value for the `Progress_report` display interval is 60 seconds. The minimal value for the `Progress_report` display interval is 1 second.

Examples

The following example shows the recovery of one backed up AIJ file with a progress report interval of 1 second. A zero value indicates no data was read from the AIJ file during that interval. The completion time estimate is an approximation not guaranteed to be exact.

```
$ SHOW TIME
11-AUG-2020 16:04:39
$ rmu/recover/LOG/root=DEVICE:[DIRECTORY]glory.rdb/PROGRESS_REPORT=1 -
  DEVICE:[DIRECTORY]backup_after.baij
%RMU-I-LOGRECDB, recovering database file DEVICE:[DIRECTORY]GLORY.RDB;1
%RMU-I-LOGOPNAIJ, opened journal file DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  at 11-AUG-2020 16:04:39.19
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 10 MB ( 7%) at 10 MB/s, estimated completion time 16:04:52.75
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 19 MB (13%) at 9 MB/s, estimated completion time 16:04:55.39
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 52 MB (35%) at 32 MB/s, estimated completion time 16:04:45.16
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 74 MB (49%) at 21 MB/s, estimated completion time 16:04:46.57
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 74 MB (49%) at 0 KB/s, estimated completion time 16:04:44.19
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 74 MB (49%) at 0 KB/s, estimated completion time 16:04:45.19
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 74 MB (49%) at 0 KB/s, estimated completion time 16:04:46.19
%RMU-I-LOGRECSTAT, transaction with TSN 225 committed
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 78 MB (52%) at 3 MB/s, estimated completion time 16:05:04.80
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 100 MB (68%) at 22 MB/s, estimated completion time 16:04:50.26
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 118 MB (79%) at 17 MB/s, estimated completion time 16:04:50.88
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 141 MB (95%) at 23 MB/s, estimated completion time 16:04:50.45
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 148 MB (99%) at 6 MB/s, estimated completion time 16:04:51.20
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 148 MB (99%) at 0 KB/s, estimated completion time 16:04:52.19
DEVICE:[DIRECTORY]BACKUP_AFTER.BAIJ;1
  Read 148 MB (99%) at 0 KB/s, estimated completion time 16:04:53.19
%RMU-I-LOGRECSTAT, transaction with TSN 226 committed
%RMU-I-AIJONEDONE, AIJ file sequence 0 roll-forward operations completed
%RMU-I-LOGRECOVR, 2 transactions committed
%RMU-I-LOGRECOVR, 0 transactions rolled back
%RMU-I-LOGRECOVR, 0 transactions ignored
%RMU-I-AIJNOACTIVE, there are no active transactions
%RMU-I-AIJSUCCESS, database recovery completed successfully
%RMU-I-AIJNXTSEQ, to continue this AIJ file recovery, the sequence number
  needed will be 1
%RMU-I-AIJALLDONE, after-image journal roll-forward operations completed
```

```

%RMU-I-LOGSUMMARY, total 2 transactions committed
%RMU-I-LOGSUMMARY, total 0 transactions rolled back
%RMU-I-LOGSUMMARY, total 0 transactions ignored
%RMU-I-AIJSUCCE, database recovery completed successfully
%RMU-I-AIJFNLSEQ, to start another AIJ file recovery, the sequence number
needed will be 1
%RMU-I-AIJNOENABLED, after-image journaling has not yet been enabled
$ SHOW TIME
11-AUG-2020 16:04:53

```

The following example shows the recovery of two AIJ files with a progress report interval of 1 second. A zero value indicates no data was read from the AIJ file during that interval. The completion time estimate is an approximation not guaranteed to be exact.

```

$ SHOW TIME
11-AUG-2020 16:08:20
$ rmu/recover/root=DEVICE:[DIRECTORY]glory.rdb/PROGRESS_REPORT=1 -
DEVICE:[DIRECTORY]backup_after1.aij, -
DEVICE:[DIRECTORY]backup_after2.aij
%RMU-I-LOGRECDB, recovering database file DEVICE:[DIRECTORY]GLORY.RDB;1
%RMU-I-LOGOPNAIJ, opened journal file DEVICE:[DIRECTORY]BACKUP_AFTER1.AIJ;1
at 11-AUG-2020 16:08:20.90
DEVICE:[DIRECTORY]BACKUP_AFTER1.AIJ;1
Read 33 MB (44%) at 33 MB/s, estimated completion time 16:08:23.12
DEVICE:[DIRECTORY]BACKUP_AFTER1.AIJ;1
Read 64 MB (86%) at 30 MB/s, estimated completion time 16:08:23.23
DEVICE:[DIRECTORY]BACKUP_AFTER1.AIJ;1
Read 74 MB (99%) at 9 MB/s, estimated completion time 16:08:23.92
DEVICE:[DIRECTORY]BACKUP_AFTER1.AIJ;1
Read 74 MB (99%) at 0 KB/s, estimated completion time 16:08:24.90
DEVICE:[DIRECTORY]BACKUP_AFTER1.AIJ;1
Read 74 MB (99%) at 0 KB/s, estimated completion time 16:08:25.90
%RMU-I-LOGRECSTAT, transaction with TSN 225 committed
%RMU-I-LOGRECSTAT, transaction with TSN 227 committed
%RMU-I-AIJONEDONE, AIJ file sequence 0 roll-forward operations completed
%RMU-I-LOGRECOVR, 2 transactions committed
%RMU-I-LOGRECOVR, 0 transactions rolled back
%RMU-I-LOGRECOVR, 0 transactions ignored
%RMU-I-AIJACTIVE, 1 active transaction not yet committed or aborted
%RMU-I-LOGRECSTAT, transaction with TSN 226 is active
%RMU-I-AIJSUCCE, database recovery completed successfully
%RMU-I-AIJNXTSEQ, to continue this AIJ file recovery, the sequence number
needed will be 1
%RMU-I-LOGOPNAIJ, opened journal file DEVICE:[DIRECTORY]BACKUP_AFTER2.AIJ;1
at 11-AUG-2020 16:08:26.63
DEVICE:[DIRECTORY]BACKUP_AFTER2.AIJ;1
Read 32 MB (44%) at 32 MB/s, estimated completion time 16:08:28.90
DEVICE:[DIRECTORY]BACKUP_AFTER2.AIJ;1
Read 67 MB (90%) at 34 MB/s, estimated completion time 16:08:28.82
DEVICE:[DIRECTORY]BACKUP_AFTER2.AIJ;1
Read 73 MB (99%) at 6 MB/s, estimated completion time 16:08:29.68
DEVICE:[DIRECTORY]BACKUP_AFTER2.AIJ;1
Read 73 MB (99%) at 0 KB/s, estimated completion time 16:08:30.63
DEVICE:[DIRECTORY]BACKUP_AFTER2.AIJ;1
Read 73 MB (99%) at 0 KB/s, estimated completion time 16:08:31.63
%RMU-I-LOGRECSTAT, transaction with TSN 226 committed
%RMU-I-AIJONEDONE, AIJ file sequence 1 roll-forward operations completed
%RMU-I-LOGRECOVR, 1 transaction committed
%RMU-I-LOGRECOVR, 0 transactions rolled back
%RMU-I-LOGRECOVR, 0 transactions ignored
%RMU-I-AIJNOACTIVE, there are no active transactions
%RMU-I-AIJSUCCE, database recovery completed successfully
%RMU-I-AIJNXTSEQ, to continue this AIJ file recovery, the sequence number
needed will be 2

```

```
%RMU-I-AIJALLDONE, after-image journal roll-forward operations completed
%RMU-I-LOGSUMMARY, total 3 transactions committed
%RMU-I-LOGSUMMARY, total 0 transactions rolled back
%RMU-I-LOGSUMMARY, total 0 transactions ignored
%RMU-I-AIJSUCCE, database recovery completed successfully
%RMU-I-AIJFNLSEQ, to start another AIJ file recovery, the sequence number
needed will be 2
%RMU-I-AIJNOENABLED, after-image journaling has not yet been enabled
$ SHOW TIME
11-AUG-2020 16:08:31
```

10.1.6 Additional Options for the RMU Set Database Command

This release of Oracle Rdb adds the following qualifiers to the RMU Set Database statement. This command is an OFFLINE command and requires exclusive access to the target database.

- **/NODES_MAX**

Sets the number of nodes that are permitted to attach to the database. This command is equivalent to the SQL ALTER DATABASE ... NUMBER OF CLUSTER NODES statement.

This qualifier can only be applied to multi-file databases. RMU will report an error for single file databases. Use SQL EXPORT DATABASE and IMPORT DATABASE to change this value for a single file database.

```
$ rmu/set data/node=1 personnel
%RMU-F-MFDBONLY, operation is not allowed on single-file databases
%RMU-F-FTL_RMU, Fatal error for RMU operation at 29-OCT-2020 13:55:55.33
```

- **/RESERVE**

This clause alters the reserve limit of the database. One or more of the keywords AREAS, CACHES, JOURNALS, or SEQUENCES with new values can be specified.

This qualifier can only be applied to multi-file databases. RMU will report an error for single file databases. Use SQL EXPORT DATABASE and IMPORT DATABASE to change this value for a single file database.

```
$ rmu/set data/res=area=10 personnel
%RMU-F-MFDBONLY, operation is not allowed on single-file databases
%RMU-F-FTL_RMU, Fatal error for RMU operation at 29-OCT-2020 13:59:19.94
```

- **AREAS**

Reserves extra storage area entries to allow subsequent ALTER DATABASE ... ADD STORAGE AREA statements.

- **CACHES**

Reserves extra row cache entries to allow subsequent ALTER DATABASE ... ADD CACHE statements.

- **JOURNALS**

Reserves extra after image journal entries to allow subsequent ALTER DATABASE ... ADD JOURNAL statements, or RMU Set After_Journal command.

- **SEQUENCES**

Reserves sequence entries. This action should be taken when a CREATE TABLE with IDENTITY or a CREATE SEQUENCE statement fails due to insufficient sequence table entries; RDMS-F-SEQTBLFUL, sequence table is full.

```
SQL> create sequence NEW_PRODUCT_CODES;
%RDB-E-NO META UPDATE, metadata update failed
-RDMS-F-SEQTBLFUL, sequence table is full
SQL>
```

- **/USERS_MAX**

Sets the number of users that are permitted to attach to the database. This command is equivalent to the SQL ALTER DATABASE ... NUMBER OF USERS statement.

This qualifier can only be applied to multi-file databases. RMU will report an error for single file databases. Use SQL EXPORT DATABASE and IMPORT DATABASE to change this value for a single file database.

```
$ rmu/set data/users=11 personnel
%RMU-F-MFDBONLY, operation is not allowed on single-file databases
%RMU-F-FTL_RMU, Fatal error for RMU operation at 29-OCT-2020 13:58:29.03
```

10.1.7 New Summary_only Qualifier to RMU Dump Audit Command

This release of Oracle Rdb adds the SUMMARY_ONLY qualifier to RMU Dump Audit. This allows the database administrator to see a list of databases that have entries recorded in the named AUDIT\$JOURNAL.

Neither the /FORMAT nor the /TYPE qualifiers are permitted with /SUMMARY_ONLY. The database parameter is ignored.

The following example generates a file containing the database names used by that version of the SECURITY.AUDIT\$JOURNAL.

```
$ define/nolog RMU_AJ SYS$COMMON:[SYSMGR]SECURITY.AUDIT$JOURNAL;8398
$ rmu/dump/audit -
    "" -
    RMU_AJ -
    /since=TODAY -
    /log -
    /summary_only -
    /output=audit_dump.txt
$
```

10.1.8 New RMU Verify Minimize_conflicts Qualifier

In this release of Oracle Rdb, a /MINIMIZE_CONFLICTS qualifier has been added to the RMU Verify command. The /MINIMIZE_CONFLICTS qualifier attempts to reduce database page lock conflicts at timed intervals during the verify operation.

Syntax

- /MINIMIZE_CONFLICTS[=n]

This qualifier requests that RMU Verify periodically attempt to reduce locks on buffers during operation. The value of (n) is the time interval in seconds. The minimal value that can be specified is 1 second. The default time interval is 30 seconds.

- /NOMINIMIZE_CONFLICTS

This qualifier requests that RMU Verify request that RMU not release locking periodically. This returns RMU to prior default behavior.

- If this qualifier is omitted then the default is assumed to be MINIMIZE_CONFLICTS=30

Examples

The following example shows examples for using this new MINIMIZE_CONFLICTS feature.

```
$!  
$! Default - minimize database page lock conflicts at 30 second  
$! intervals  
$!  
$ RMU/VERIFY/ALL/NOLOG TEST.RDB  
$!  
$! Default time interval - minimize database page lock conflicts  
$! at 30 second intervals  
$!  
$ RMU/VERIFY/ALL/NOLOG/MINIMIZE_CONFLICTS TEST.RDB  
$!  
$! Minimize database page lock conflicts at 1 second intervals  
$!  
$ RMU/VERIFY/ALL/NOLOG/MINIMIZE_CONFLICTS=1 TEST.RDB  
$!  
$! Do not minimize page lock conflicts  
$!  
$ RMU/VERIFY/ALL/NOLOG/NOMINIMIZE_CONFLICTS TEST.RDB  
$
```

10.1.9 New OPTION=GENERATED Added to RMU Extract Command

This release of Oracle Rdb includes a new GENERATED option for RMU Extract. In prior releases RMU Extract ITEM=UNLOAD and ITEM=LOAD would generate load commands that assumed all the columns were updatable.

The option FULL can be used to generate syntax that loads every field by name and includes virtual columns (AUTOMATIC AS, GENERATED, IDENTITY and COMPUTED BY) as commented out names. Therefore, editing was required to uncomment GENERATED column names so they could be reloaded. In addition the /VIRTUAL=AUTOMATIC qualifier needed to be added to the RMU Load and RMU Unload commands.

Now using OPTION=(GENERATED) will instruct RMU Extract to generate more appropriate DCL commands for unloading and re-loading data in tables that contain GENERATED columns.

The following example shows a portion of a generated DCL procedure when only **OPTION=FULL** is used.

```
$ RMU/EXTRACT/ITEM=UNLOAD/OPTION=FULL SAMPLE_DB
.
.
.
$ CREATE SAMPLE0.COLUMNS
! Columns list for table SAMPLE0
! in ...
! Created by RMU Extract for Oracle Rdb ... on 29-JAN-2021 13:20:28.40
! Virtual: IDENT_COL
DETAILS
! Virtual: LAST_UPDATE
$ RMU/UNLOAD -
  USER1:[TESTING.DATABASES]MF_PERSONNEL_SQL.RDB -
  /FIELDS="@SAMPLE0.COLUMNS" -
  SAMPLE0 -
  SAMPLE0.UNL
$
```

The following example shows a portion of a generated DCL procedure when **OPTION=(GENERATED)** is used.

```
$ RMU/EXTRACT/ITEM=UNLOAD/OPTION=GENERATED SAMPLE_DB
.
.
.
$ CREATE SAMPLE0.COLUMNS
! Columns list for table SAMPLE0
! in ...
! Created by RMU Extract for Oracle Rdb ... on 29-JAN-2021 13:23:27.76
IDENT COL
DETAILS
LAST_UPDATE
$ RMU/UNLOAD -
  USER1:[TESTING.DATABASES]MF_PERSONNEL_SQL.RDB -
  /FIELDS="@SAMPLE0.COLUMNS" -
  /VIRTUAL=AUTOMATIC -
  SAMPLE0 -
  SAMPLE0.UNL
$
```

10.1.10 Changed Behavior for the Noedit_filename Qualifier in RMU Backup After_Journal Command

In prior releases of Oracle Rdb the **/NOEDIT_FILENAME** qualifier on the RMU Backup After_Journal was ignored. With this release it takes on a new meaning as described below:

– **/EDIT_FILENAME**

As with previous versions, this qualifier defines the editing to be performed for the output backup file name. This editing is performed on the provided backup filename, or if "" is specified the default backup filename defined in the database.

This qualifier replaces any **EDIT_FILENAME** defined for the database.

– **/NOEDIT_FILENAME**

This qualifier negates any prior usage on the command of the /EDIT_FILENAME qualifier and also instructs RMU to ignore the EDIT_FILENAME defined by the SQL ALTER DATABASE statement, or RMU Set After_Journal command. This is a change of behavior from prior versions and supports the enhancements made to the RMU Set After_Journal command which allows the defaults to be defined for the MANUAL backup processing.

No editing is performed on the provided backup filename, or if "" is specified the default backup filename defined in the database is used without changes.

- Neither /EDIT_FILENAME nor /NOEDIT_FILENAME was used.

In this case RMU Backup After_Journal will use the default if defined in the database by SQL ALTER DATABASE statement, or RMU Set After_Journal command.

10.2 Obsolete Features

10.2.1 Comma Statement Separator in Trigger Body No Longer Supported

The syntax for trigger actions in the CREATE TRIGGER statement has, in the past, supported the comma (,) as well as the semicolon (;) as statement separators. The use of the comma separator has been problematic in Oracle Rdb SQL because it conflicts in various places with the comma used as an element separator within some statements. For example, the TRACE statement allows a comma separated list of values and the INSERT INTO ... SELECT ... FROM statement allows a comma separated list of table names in the FROM clause. In these cases, a comma cannot be used as a statement separator because the current statement appears to be continued.

Future versions of Oracle Rdb are expected to include enhancements to the TRIGGER action syntax which will allow other statements to include comma as an element separator. Therefore, the comma statement separator is now no longer supported.

Any scripts or applications that include the CREATE TRIGGER statement must now be modified to use only the semicolon (;) as a separator.

This change does not affect existing database triggers, only new triggers defined using the CREATE TRIGGER statement. The RMU Extract Item=TRIGGER command already generates semicolon separators in extracted CREATE TRIGGER statements.

Enhancements And Changes Provided in Oracle Rdb Release 7.4.1.0

11.1 Enhancements

11.1.1 PCSI Support for Rdb Kit Installation and Deinstallation

Whenever Oracle Rdb is installed or deinstalled, Oracle Rdb will be registered in the PCSI software product database. This will allow users to use the PCSI `PRODUCT SHOW HISTORY` and `PRODUCT SHOW PRODUCT` commands to display information about releases of Oracle Rdb that have been installed or deinstalled. This information will also be helpful as input whenever a Service Request (SR) is submitted to Oracle Support.

The following lines will now be displayed during the installation of Oracle Rdb, showing that the installation has been registered in the PCSI database.

```
The following product has been selected:
ORCL I64VMS RDB74 V7.4-100          Transition (registration)

The following product will be registered:
ORCL I64VMS RDB74 V7.4-100          DISK$NODE84_2:[VMS$COMMON.]

File lookup pass starting ...

Portion done: 0%
...100%

File lookup pass completed search for all files listed in the product's PDF
Total files searched: 0   Files present: 0   Files absent: 0

The following product has been registered:
ORCL I64VMS RDB74 V7.4-100          Transition (registration)
%VMSINSTAL-I-MOVEFILES, Files will now be moved to their target directories...
```

Registration in the PCSI software product database allows a user to use commands such as the following to track what Oracle Rdb releases are currently installed and the history of any past product installations and deinstallations.

```
$ PRODUCT SHOW HISTORY/SINCE
-----
PRODUCT                KIT TYPE   OPERATION  VAL DATE
-----
ORCL I64VMS RDB74 V7.4-100  Transition  Reg Product (U) 10-JUN-2020
-----

1 item found

$ PRODUCT SHOW HISTORY RDB7*
-----
PRODUCT                KIT TYPE   OPERATION  VAL DATE
-----
ORCL I64VMS RDB74 V7.4-100  Transition  Reg Product (U) 10-JUN-2020
-----

1 item found
```

```
$ PRODUCT SHOW PRODUCT RDB7*
```

PRODUCT	KIT TYPE	STATE
ORCL I64VMS RDB74 V7.4-100	Transition	Installed

```
1 item found
```

The following lines will now be displayed during the deinstallation of Oracle Rdb, showing that the removal of the release has been registered in the PCSI database. Deinstallation is performed by executing the DCL procedure SYSSMANAGER:RDB\$DEINSTALL_DELETE.COM. Please refer to section "Deleting Versions of Oracle Rdb" in the Oracle Rdb Installation Guide for further details.

```
The following product has been selected:
```

ORCL I64VMS RDB74 V7.4-100	Transition (registration)
----------------------------	---------------------------

```
The following product will be removed from destination:
```

ORCL I64VMS RDB74 V7.4-100	DISK\$NODE84_2:[VMS\$COMMON.]
----------------------------	-------------------------------

```
Portion done: 0%...100%
```

```
The following product has been removed:
```

ORCL I64VMS RDB74 V7.4-100	Transition (registration)
----------------------------	---------------------------

The example below shows the additional information that will be displayed by the PCSI PRODUCT commands as a result of the deinstallation of a release of Oracle Rdb.

```
$ PRODUCT SHOW HISTORY/SINCE
```

PRODUCT	KIT TYPE	OPERATION	VAL	DATE
ORCL I64VMS RDB74 V7.4-100	Transition	Remove	-	10-JUN-2020
ORCL I64VMS RDB74 V7.4-100	Transition	Reg Product	(U)	10-JUN-2020

```
2 items found
```

```
$ PRODUCT SHOW HISTORY RDB7*
```

PRODUCT	KIT TYPE	OPERATION	VAL	DATE
ORCL I64VMS RDB74 V7.4-100	Transition	Remove	-	10-JUN-2020
ORCL I64VMS RDB74 V7.4-100	Transition	Reg Product	(U)	10-JUN-2020

```
2 items found
```

```
$ PRODUCT SHOW PRODUCT RDB7*
```

PRODUCT	KIT TYPE	STATE
---------	----------	-------

```
0 items found
```

11.1.2 Some Aggregate Functions Inherit Source Column EDIT STRING

Oracle Rdb supports EDIT STRING inheritance for these functions when using Interactive SQL.

- MAX, MEDIAN, MIN, FIRST_VALUE, LAST_VALUE

When the input type matches the output type, then the EDIT STRING from the source column is inherited to improve the readability of the aggregate.

- CAST

When the datatype of the CAST includes a domain with the EDIT STRING.

The following example shows the EDIT STRING being used.

```
SQL> create domain DOM_TST integer(2) edit string '$(9)9.99';
SQL>
SQL> create table TST
cont>      (a integer(2) edit string '$(9)9.99'
cont>      ,c char(10)
cont>      );
SQL>
SQL> insert into TST
cont>      values (100, 100, 'A');
1 row inserted
SQL> insert into TST
cont>      values (233, 233, 'B');
1 row inserted
SQL>
SQL> --> column with explicit edit string
SQL> select min (a), max (a), cast (a as DOM_TST)
cont> from TST
cont> group by a
cont> ;
cont> ;

          $100.00          $100.00          $100.00
          $233.00          $233.00          $233.00
2 rows selected
SQL>
SQL> select first_value (a) within group (order by b desc),
cont>      last_value (a) within group (order by b desc),
cont>      median (a)
cont> from TST
cont> ;

          $233.00          $100.00          $166.50
1 row selected
SQL>
```

Use the SET DISPLAY NO EDIT STRING statement to disable this behavior.

11.1.3 Enhanced LIKE Table Support in CREATE TABLE Statement

This release of Oracle Rdb introduces support for the ANSI and ISO SQL Language Standard syntax for the LIKE table clause. It also adds new EXCLUDING and INCLUDING clauses to the LIKE clause within the CREATE TABLE statement.

In prior releases of Oracle Rdb, a table can be created using syntax similar to the following:

```
SQL> create table RETIRED_EMPLOYEES
cont>      like EMPLOYEES
cont>      ;
SQL>
```

This statement copies the definitions of each column as well as DEFAULT values defined for those source columns. SQL also allows additional columns and constraints to be defined for the new table.

```
SQL> create table RETIRED_EMPLOYEES
cont>      like EMPLOYEES
cont>      (retirement_date DATE
cont>      ,check (retirement_date > birthday) not deferrable
cont>      );
SQL>
```

This syntax is retained for backward compatibility with prior releases of Oracle Rdb.


The syntax for a similar feature in the ANSI/ISO SQL Database Language moves the LIKE clause into the section that defines the columns and constraint. This adds the ability to copy column definitions from more than one table, control how GENERATED, AUTOMATIC, IDENTITY and COMPUTED columns are inherited, as well as define the column ordering; this is determined by the order of the listed columns and tables.

```
SQL> create table RETIRED_EMPLOYEES
cont>      (retirement_date DATE
cont>      ,like EMPLOYEES
cont>          including COMPUTED
cont>          excluding DEFAULTS
cont>      ,check (retirement_date > birthday) not deferrable
cont>      ,unique (employee_id)
cont>      ,hr_authorizations LIST OF BYTE VARYING
cont>      );
SQL>
```

By default, GENERATED, AUTOMATIC, IDENTITY and COMPUTED columns are not copied but columns representing the same data types are created instead.

Syntax


column-constraint-list =



ansi-like-table-clause =



like-attributes =



Usage Notes

- When using the LIKE clause to copy a table definition, the creator of the new table must have REFERENCES or SELECT privilege granted for the referenced table.
- By default, Rdb includes the column protections (access control lists) and comments for any copied column. These new clauses allow the database administrator to suppress the copying of that metadata.

- The LIKE clause can be used multiple times within a CREATE TABLE statement. However, if the copied tables include any duplicate column names, then an error will be reported. Only one IDENTITY column can be defined or inherited. Use the INCLUDING IDENTITY clause, if necessary, to inherit the attributes from the referenced table.

The default behavior is EXCLUDING COMPUTED, GENERATED, IDENTITY column details. In this case, non-generated columns will be created which contain the same data type attributes. Default values defined for the source tables are not automatically inherited. Use the INCLUDING DEFAULTS clause to control this behavior.

Note: For backward compatibility with previous versions of Oracle Rdb, the LIKE clause used outside the column-constraint-list defaults to INCLUDING GENERATED, INCLUDING IDENTITY, INCLUDING COMPUTED and INCLUDING DEFAULTS. The like-attributes may not be specified in this location and therefore these defaults may not be changed.

- The clauses EXCLUDING GENERATED or INCLUDING GENERATED apply to columns defined using the GENERATED ... AS (expr) and AUTOMATIC ... AS (expr) syntax. When EXCLUDING is used or implied, the generated (or automatic) column is converted to a simple base column with the same data types.
- The clauses EXCLUDING IDENTITY or INCLUDING IDENTITY apply to columns defined using the GENERATED ... AS IDENTITY and IDENTITY (...) syntax. When EXCLUDING is used or implied, the identity column is converted to a simple base column with the same data types.
- The clauses EXCLUDING COMPUTED or INCLUDING COMPUTED apply to columns defined using the COMPUTED BY expr syntax. When EXCLUDING is used or implied, the computed by column is converted to a simple base column with the same data types. Note that the column will require space in the defined table, which isn't true for COMPUTED BY columns.
- When the LIKE clause is used within the column-constraint-list, then EXCLUDING DEFAULTS is assumed. Use the INCLUDING DEFAULTS if you wish the inherited columns to have DEFAULTS inherited from the source table.
- The LIKE clause is only used to inherit the column definitions from the referenced table. Once the table is created with LIKE clauses, subsequent changes to the source table are not propagated to the created tables.

Examples

The following example shows the use of the LIKE clause to inherit columns from various template tables.

```

SQL> create table NAMES_REC
cont>      (LAST_NAME          LAST_NAME_DOM
cont>        ,FIRST_NAME        FIRST_NAME_DOM
cont>        ,MIDDLE_INITIAL     MIDDLE_INITIAL_DOM
cont>        );
SQL>
SQL> create table ADDRESS_REC
cont>      (ADDRESS_DATA_1      ADDRESS_DATA_1_DOM
cont>        ,ADDRESS_DATA_2      ADDRESS_DATA_2_DOM
cont>        ,CITY                 CITY_DOM
cont>        ,STATE                STATE_DOM
cont>        ,POSTAL_CODE          POSTAL_CODE_DOM
cont>        );
SQL>
SQL> create table employees
cont>      (EMPLOYEE_ID          ID_DOM not null
cont>        ,like NAMES_REC      including DEFAULTS
cont>        ,like ADDRESS_REC     including DEFAULTS
cont>        ,SEX                  SEX_DOM
cont>        ,BIRTHDAY             DATE_DOM
cont>        ,STATUS_CODE          STATUS_CODE_DOM
cont>        );
SQL>

```

The resulting CREATE TABLE for the EMPLOYEES table is easier to read and allows for consistency among similar definitions.

```

SQL> show table (column) EMPLOYEES;
Information for table EMPLOYEES

```

Columns for table EMPLOYEES:

Column Name	Data Type	Domain
-----	-----	-----
EMPLOYEE_ID	CHAR(5)	ID_DOM
Not Null constraint EMPLOYEES_EMPLOYEE_ID_NOT_NULL		
LAST_NAME	CHAR(14)	LAST_NAME_DOM
FIRST_NAME	CHAR(10)	FIRST_NAME_DOM
MIDDLE_INITIAL	CHAR(1)	MIDDLE_INITIAL_DOM
ADDRESS_DATA_1	CHAR(25)	ADDRESS_DATA_1_DOM
ADDRESS_DATA_2	CHAR(20)	ADDRESS_DATA_2_DOM
CITY	CHAR(20)	CITY_DOM
STATE	CHAR(2)	STATE_DOM
POSTAL_CODE	CHAR(5)	POSTAL_CODE_DOM
SEX	CHAR(1)	SEX_DOM
BIRTHDAY	DATE VMS	DATE_DOM
STATUS_CODE	CHAR(1)	STATUS_CODE_DOM

```

SQL>

```

11.1.4 RMU Reclaim Free_pages Qualifier Frees Unused Data Page Clumps

There is an additional Free_pages qualifier for the RMU Reclaim command. This qualifier is used to free unused data page clumps that are allocated in uniform storage areas. It will free all unused page clumps in an entire uniform storage area or all unused page clumps in one or more specified table or index logical areas in uniform storage areas. Any deleted dbkeys and locked space on pages will also be freed.

Command Qualifiers

/[NO]FREE_PAGES

Nofree_pages is the default.

Other qualifiers may be used in conjunction with the Free_pages qualifier.

/AREA[=storage-area-name-list]

Area is used to specify a list of uniform storage area names to process. The wild card syntax AREA=* can be specified for processing all uniform storage areas in the database.

The default for the Area qualifier is all uniform storage areas in the database.

/LAREA=logical-area-name-list

Larea is used to specify a list of individual table or index logical area names to process.

There is no default for the Larea qualifier. A list of logical area names must be specified. The logical area name will be used to determine the storage area where the logical area is located. If the logical area is partitioned among multiple storage areas, each logical area partition will be processed.

This qualifier can only be specified if the Free_pages qualifier is specified.

/LOCK_TIMEOUT=seconds

Lock_timeout is used to specify a lock timeout value that will be in effect during the execution of the RMU Reclaim Free_pages command.

Lock_timeout can only be specified if Free_pages is also specified. The value specified with this qualifier is the maximum time in seconds during which the current RMU Reclaim Free_pages command will wait to acquire an exclusive update lock on the current storage area or logical area to be processed when accessing an on-line database with other users.

If Lock_timeout is not specified, one of the following values will be used, in the specified order of precedence.

1. The value of the logical name RDM\$BIND_LOCK_TIMEOUT_INTERVAL, if it has been specified.
2. The "LOCK TIMEOUT INTERVAL" specified by the SQL CREATE or ALTER DATABASE command is used.
3. The RMU Reclaim Free_pages command will wait indefinitely to acquire an exclusive update lock on the current storage area or logical area to be processed.

If /LOCK_TIMEOUT=0 is specified, the RMU Reclaim Free_pages command will ignore any lock timeout defaults that may be in effect and wait indefinitely to acquire an exclusive update lock on the current storage area or logical area to be processed.

Usage Notes

- The Free_pages command can be used when the database is active. Please note that RMU will lock affected areas during processing, which may reduce concurrency.
- Free_pages is not a default qualifier for the RMU Reclaim command. If the Free_pages qualifier is not specified, the RMU Reclaim command will implement the default functionality of freeing deleted dbkeys and locked space in mixed and uniform database storage areas.

- RMU Reclaim Free_Pages can be interrupted at any time; any work in progress will be rolled back. Actions that are completed will have each been committed: if processing a list of logical area names (/LAREA), a commit is performed after each logical area and if processing a list of storage areas (/AREA), a commit is performed after each storage area. Note that tables and indices which are partitioned have multiple logical areas that share the same name as the table or index.
- If Free_pages is specified without either the Larea or Area qualifier, all the uniform storage areas in the database will be processed.
- The Area and Larea qualifiers cannot both be specified in the same RMU Reclaim command.
- If a mixed storage area name is specified with the Area qualifier or the name of a logical area in a mixed storage area is specified with the Larea qualifier, a warning message will be output and a warning status will be returned by the Reclaim command. That storage or logical area will not be processed but the Reclaim command will continue processing the next storage or logical area in the specified list of storage areas or logical areas.
- If a lock wait timeout occurs, a warning message will be output and a warning status will be returned by the Reclaim command. That storage or logical area will not be processed but the Reclaim command will continue processing the next storage area or logical area in the specified list of storage or logical areas.
- The RMU Reclaim Free_pages functionality replaces that provided by RMU REPAIR /INITIALIZE=FREE_PAGES. The main advantage of Reclaim is that it can be run on an active database.

Examples

Examples using /AREA

The following examples show the Free_pages qualifier with the Area qualifier to free unused page clumps for one or more named storage areas.

```
$ RMU/RECLAIM/LOG/AREA=ABM_AREA1/FREE_PAGES ABM_SAMPLE.RDB
%RMU-I-RCLMAREA, Reclaiming area ABM_AREA1
%RMU-I-RCLMPAGPRC, 2138 pages processed for area ABM_AREA1
%RMU-I-RCLMPAGFREED, 1992 clump pages freed for area ABM_AREA1
$
$ RMU/RECLAIM/FREE_PAGES/AREA=(MFDBA2,MFDBA1)/LOG MFDB
%RMU-I-RCLMAREA, Reclaiming area MFDBA2
%RMU-I-RCLMPAGPRC, 13 pages processed for area MFDBA2
%RMU-I-RCLMPAGFREED, 4 clump pages freed for area MFDBA2
%RMU-I-RCLMAREA, Reclaiming area MFDBA1
%RMU-I-RCLMPAGPRC, 13 pages processed for area MFDBA1
%RMU-I-RCLMPAGFREED, 4 clump pages freed for area MFDBA1
$
$ RMU/RECLAIM/FREE_PAGES/AREA=*/LOG MFDB
%RMU-I-RCLMAREA, Reclaiming area DISK:[DIRECTORY]MFDB.RDA;1
%RMU-I-RCLMPAGPRC, 701 pages processed for area
DISK:[DIRECTORY]MFDB.RDA;1
%RMU-I-RCLMPAGFREED, 220 clump pages freed for area
DISK:[DIRECTORY]MFDB.RDA;1
%RMU-I-RCLMAREA, Reclaiming area DISK:[DIRECTORY]MFDBA1.RDA;1
%RMU-I-RCLMPAGPRC, 13 pages processed for area
DISK:[DIRECTORY]MFDBA1.RDA;1
%RMU-I-RCLMPAGFREED, 0 clump pages freed for area
DISK:[DIRECTORY]MFDBA1.RDA;1
%RMU-I-RCLMAREA, Reclaiming area DISK:[DIRECTORY]MFDBA2.RDA;1
```



```

%RMU-I-RCLMPAGPRC, 13 pages processed for area
DISK: [DIRECTORY]MFDBA2.RDA;1
%RMU-I-RCLMPAGFREED, 0 clump pages freed for area
DISK: [DIRECTORY]MFDBA2.RDA;1
%RMU-I-RCLMAREA, Reclaiming area DISK: [DIRECTORY]MFDBA3.RDA;1
%RMU-I-RCLMPAGPRC, 13 pages processed for area
DISK: [DIRECTORY]MFDBA3.RDA;1
%RMU-I-RCLMPAGFREED, 0 clump pages freed for area
DISK: [DIRECTORY]MFDBA3.RDA;1
$
$ RMU/RECLAIM/FREE_PAGES/AREA/LOG MFDB
%RMU-I-RCLMAREA, Reclaiming area DISK: [DIRECTORY]MFDB.RDA;1
%RMU-I-RCLMPAGPRC, 701 pages processed for area
DISK: [DIRECTORY]MFDB.RDA;1
%RMU-I-RCLMPAGFREED, 220 clump pages freed for area
DISK: [DIRECTORY]MFDB.RDA;1
%RMU-I-RCLMAREA, Reclaiming area DISK: [DIRECTORY]MFDBA1.RDA;1
%RMU-I-RCLMPAGPRC, 13 pages processed for area
DISK: [DIRECTORY]MFDBA1.RDA;1
%RMU-I-RCLMPAGFREED, 0 clump pages freed for area
DISK: [DIRECTORY]MFDBA1.RDA;1
%RMU-I-RCLMAREA, Reclaiming area DISK: [DIRECTORY]MFDBA2.RDA;1
%RMU-I-RCLMPAGPRC, 13 pages processed for area
DISK: [DIRECTORY]MFDBA2.RDA;1
%RMU-I-RCLMPAGFREED, 0 clump pages freed for area
DISK: [DIRECTORY]MFDBA2.RDA;1
%RMU-I-RCLMAREA, Reclaiming area DISK: [DIRECTORY]MFDBA3.RDA;1
%RMU-I-RCLMPAGPRC, 13 pages processed for area
DISK: [DIRECTORY]MFDBA3.RDA;1
%RMU-I-RCLMPAGFREED, 0 clump pages freed for area
DISK: [DIRECTORY]MFDBA3.RDA;1
$

```

Examples using /LAREA

The following examples show the Free_pages qualifier with the Larea qualifier to free unused page clumps for one or more named table and index logical areas.

```

$ RMU/RECLAIM/LOG/LAREA=SAMPLE_TABLE/FREE_PAGES ABM_SAMPLE.RDB
%RMU-I-RCLMLAREA, Reclaiming logical area SAMPLE_TABLE in physical
area DISK: [DIRECTORY]ABM_AREA1.RDA;1
%RMU-I-RCLMLPAGPRC, 2008 pages processed for logical area SAMPLE_TABLE
in physical area DISK: [DIRECTORY]ABM_AREA1.RDA;1
%RMU-I-RCLMLPAGFREED, 1992 clump pages freed for logical area
SAMPLE_TABLE in physical area DISK: [DIRECTORY]ABM_AREA1.RDA;1
$
$ RMU/RECLAIM/LOG/LAREA=(SAMPLE_TABLE,SAMPLE_TABLE2)/FREE_PAGES
ABM_SAMPLE.RDB
%RMU-I-RCLMLAREA, Reclaiming logical area SAMPLE_TABLE in physical
area DISK: [DIRECTORY]ABM_AREA1.RDA;1
%RMU-I-RCLMLPAGPRC, 2008 pages processed for logical area SAMPLE_TABLE
in physical area DISK: [DIRECTORY]ABM_AREA1.RDA;1
%RMU-I-RCLMLPAGFREED, 1992 clump pages freed for logical area
SAMPLE_TABLE in physical area DISK: [DIRECTORY]ABM_AREA1.RDA;1
%RMU-I-RCLMLAREA, Reclaiming logical area SAMPLE_TABLE2 in physical
area DISK: [DIRECTORY]ABM_AREA2.RDA;1
%RMU-I-RCLMLPAGPRC, 2008 pages processed for logical area
SAMPLE_TABLE2 in physical area DISK: [DIRECTORY]ABM_AREA2.RDA;1
%RMU-I-RCLMLPAGFREED, 1992 clump pages freed for logical area
SAMPLE_TABLE2 in physical area DISK: [DIRECTORY]ABM_AREA2.RDA;1
$
$ RMU/RECLAIM/LOG/LAREA=NDX_NAME/FREE_PAGES TEST_DATABASE.RDB
%RMU-I-RCLMLAREA, Reclaiming logical area NDX_NAME in
physical area DISK: [DIRECTORY]DB_DEFAULT.RDA;1
%RMU-I-RCLMLPAGPRC, 12 pages processed for logical area
NDX_NAME in physical area DISK: [DIRECTORY]DB_DEFAULT.RDA;1

```

```
%RMU-I-RCLMLPAGFREED, 3 clump pages freed for logical area
NDX_NAME in physical area DISK:[DIRECTORY]DB_DEFAULT.RDA;1
$
```

Examples using /LAREA and partitioned index logical area

The following example shows the Free_pages qualifier specified in an RMU Reclaim command with the Larea qualifier to free unused page clumps in an index logical area partitioned among different storage areas. Each index logical area partition is processed separately.

```
$ RMU/RECLAIM/FREE_PAGES/LOG/LAREA=INDEXA TEST_DATABASE.RDB
%RMU-I-RCLMLAREA, Reclaiming logical area INDEXA in physical area
DISK:[DIRECTORY]INDEXA_1.RDA;1
%RMU-I-RCLMLPAGPRC, 8 pages processed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_1.RDA;1
%RMU-I-RCLMLPAGFREED, 3 clump pages freed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_1.RDA;1
%RMU-I-RCLMLAREA, Reclaiming logical area INDEXA in physical area
DISK:[DIRECTORY]INDEXA_2.RDA;1
%RMU-I-RCLMLPAGPRC, 8 pages processed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_2.RDA;1
%RMU-I-RCLMLPAGFREED, 3 clump pages freed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_2.RDA;1
%RMU-I-RCLMLAREA, Reclaiming logical area INDEXA in physical area
DISK:[DIRECTORY]INDEXA_3.RDA;1
%RMU-I-RCLMLPAGPRC, 8 pages processed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_3.RDA;1
%RMU-I-RCLMLPAGFREED, 3 clump pages freed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_3.RDA;1
%RMU-I-RCLMLAREA, Reclaiming logical area INDEXA in physical area
DISK:[DIRECTORY]INDEXA_4.RDA;1
%RMU-I-RCLMLPAGPRC, 8 pages processed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_4.RDA;1
%RMU-I-RCLMLPAGFREED, 3 clump pages freed for logical area INDEXA in
physical area DISK:[DIRECTORY]INDEXA_4.RDA;1
$
```


Examples showing warnings

In the following examples, warning messages are output even if /LOG is not specified in the RMU Reclaim Free_pages commands if a mixed storage area is specified or a logical area could not be processed because of a lock conflict with another user. The optional Lock_timeout qualifier is specified.

```
$ RMU/RECLAIM/LOG/AREA=DEPARTMENTS/FREE_PAGES MF_PERSONNEL.RDB
%RMU-W-RCLMMIXIGN, Mixed area DEPARTMENTS not processed if
RMU/RECLAIM/FREE_PAGES
$
$ RMU/RECLAIM/NOLOG/AREA=ABM_AREA1/FREE_PAGES/LOCK_TIMEOUT=600 -
$ _ABM_SAMPLE.RDB
%RMU-W-RCLMARNOTPRC, Area ABM_AREA1 could not be processed due to a
lock conflict
$
$ RMU/RECLAIM/LOG/LAREA=SAMPLE_TABLE/FREE_PAGES/LOCK_TIMEOUT=1200 -
$ _ABM_SAMPLE.RDB
%RMU-I-RCLMLAREA, Reclaiming logical area SAMPLE_TABLE in physical
area DISK:[DIRECTORY]ABM_AREA1.RDA;1
%RMU-W-RCLMARNOTPRC, Logical area SAMPLE_TABLE could not be processed
due to a lock conflict
%RMU-I-RCLMLPAGPRC, 0 pages processed for logical area SAMPLE_TABLE in
physical area DISK:[DIRECTORY]ABM_AREA1.RDA;1
%RMU-I-RCLMLPAGFREED, 0 clump pages freed for logical area
SAMPLE_TABLE in physical area DISK:[DIRECTORY]ABM_AREA1.RDA;1
$
```

11.1.5 CREATE DEFAULT AUDIT Command Supports OR REPLACE Clause

This release of Oracle Rdb enhances the CREATE DEFAULT AUDIT statement by allowing the OR REPLACE clause.



Arguments

- **OR REPLACE**

If the OR REPLACE clause is used and the referenced object-type exists, then it will be modified using the specified audit flags and comment. Any attributes that are not specified will assume their default values.

Note

Any protections granted to the object by the GRANT statement are not replaced. They would need to be removed using the REVOKE statement.

If the referenced object-type does not exist, then it will be created as if the OR REPLACE clause was not used.

Example

This example shows the CREATE DEFAULT AUDIT statement adding a new table object (which is always named RDB\$DEFAULT_AUDIT_TABLE).

```
SQL> create default audit
cont>     for table
cont>     all privileges
cont>     comment is 'Add a default audit table so we can inherit an ACL'
cont> ;
SQL>
SQL> grant select, delete, update, insert, show on rdb$default_audit_table to
testuser2;
SQL> grant select, dbctrl on rdb$default_audit_table to testuser3;
SQL> grant show on rdb$default_audit_table to public;
SQL>
SQL> --> display the attributes for the default table
SQL> --> note: that TESTUSER1 was the executor of the CREATE statement
SQL> show protection on table rdb$default_audit_table;
Protection on Table RDB$DEFAULT_AUDIT_TABLE
  (IDENTIFIER=[TEST,TESTUSER3],ACCESS=SELECT+DBCTRL)
  (IDENTIFIER=[TEST,TESTUSER2],ACCESS=SELECT+INSERT+UPDATE+DELETE+SHOW)
  (IDENTIFIER=[TEST,TESTUSER1],ACCESS=SELECT+INSERT+UPDATE+DELETE+SHOW+
CREATE+ALTER+DROP+DBCTRL+REFERENCES)
  (IDENTIFIER=[*,*],ACCESS=SHOW)
SQL> show audit on table rdb$default_audit_table;
Audit information for Table RDB$DEFAULT_AUDIT_TABLE
Audit Privileges:
  ALL
Alarm Privileges:
  ALL
SQL> show table (comment) rdb$default_audit_table;
Information for table RDB$DEFAULT_AUDIT_TABLE
```

Comment on table RDB\$DEFAULT_AUDIT_TABLE:
Add a default audit table so we can inherit an ACL

A global temporary table.
On commit Delete rows

SQL>

At some later time, the table template object can be created or replaced using the CREATE OR REPLACE DEFAULT AUDIT statement.

```
SQL> create or replace default audit
cont>     for table
cont>     type is (audit)
cont>     privileges (success, failure)
cont>     comment is 'Only audit SUCCESS and FAILURE'
cont> ;
SQL>
SQL> --> show that the protections are retained by OR REPLACE
SQL> show protection on table rdb$default_audit_table;
Protection on Table RDB$DEFAULT_AUDIT_TABLE
  (IDENTIFIER=[TEST,TESTUSER3],ACCESS=SELECT+DBCTRL)
  (IDENTIFIER=[TEST,TESTUSER2],ACCESS=SELECT+INSERT+UPDATE+DELETE+SHOW)
  (IDENTIFIER=[TEST,TESTUSER1],ACCESS=SELECT+INSERT+UPDATE+DELETE+SHOW
+CREATE+ALTER+DROP+DBCTRL+REFERENCES)
  (IDENTIFIER=[*,*],ACCESS=SHOW)
SQL> show audit on table rdb$default_audit_table;
Audit information for Table RDB$DEFAULT_AUDIT_TABLE
  Audit Privileges:
    SUCCESS,FAILURE

SQL> show table (comment) rdb$default_audit_table;
Information for table RDB$DEFAULT_AUDIT_TABLE

Comment on table RDB$DEFAULT_AUDIT_TABLE:
Only audit SUCCESS and FAILURE

A global temporary table.
On commit Delete rows

SQL>
SQL> commit;
SQL>
```

11.1.6 System Privileges Feature

This release of Oracle Rdb introduces Database System Privileges - an enhancement for database security.

Introduction

A typical Oracle Rdb database will have a security policy defined by granting privileges to users of the database and, when objects are created (tables, sequences, and so on), granting access to those objects. In addition, roles (also known as rights identifiers) may be granted specific access and any user assigned that role may inherit its access rights.

These granted privileges are stored in an access control list (ACL) with an entry for a user or role known as an access control entry (ACE).

Please refer to the *Oracle Rdb Server SQL Reference Manual* GRANT statement and REVOKE statement sections for more detailed descriptions and examples.

Security Definitions

All objects are within the DATABASE security domain.

The primary security objects are: ASSERTION (an assertion is a form of constraint defined independently of a table definition. This feature is currently not available in SQL but maps to the standalone RDO constraints), CATALOG, COLLATING SEQUENCE, DOMAIN, OUTLINE, PROCEDURE, PROFILE, TABLE, ROLE, SCHEMA, SEQUENCE, SYNONYM, and USER.

For the purposes of system privileges, Oracle Rdb treats modules, functions and procedures as a single class - namely PROCEDURE.

Some primary security objects have an associated ACL that protects that object and controls access to sub-objects. TABLE includes the following sub-objects: VIEW, STORAGE MAP, and TRIGGER.

When TABLE, SEQUENCE, and PROCEDURE objects are created, they are implicitly given an ACL that grants ALL PRIVILEGES to the creator and NO PRIVILEGES to the PUBLIC (also known as [*,*]). The database administrator can override this default ACL for new primary objects:

1. For tables and views, the PUBLIC access control entry will be inherited from the database access control entry for the DEFAULT user. Note: not all OpenVMS systems have a DEFAULT user defined in the system user authorization file (UAF) so that would be created by the system manager if required.

```
SQL> grant show
cont>      on database alias paysys
cont>      to default
cont> ;
SQL>
.
.
.
SQL> create table paysys.CONTROL_TABLE
cont>      (identifier_value integer generated by default as identity
cont>      );
SQL>
SQL> show protection on table paysys.CONTROL_TABLE;
Protection on Table PAYSYS.CONTROL_TABLE
  (IDENTIFIER=[ADMIN,DATABASE],ACCESS=SELECT+INSERT+UPDATE+DELETE+
  SHOW+CREATE+ALTER+DROP+DBCTRL+REFERENCES)
  (IDENTIFIER=[*,*],ACCESS=SHOW)
SQL>
```

2. Alternately, the CREATE DEFAULT AUDIT statement can be used to define template security objects in the database. These special objects are used for audit and protection inheritance. When an object (table, sequence, etc) is created, then the access control list from the security template is inherited. Note that choice (1) above will also be applied to the TABLE and VIEW default audit templates.

This example shows the **default audit** object creation and then being granted a default access for PUBLIC and specific access for users and roles.

```


SQL> create default audit
cont>     alias PAYSYS
cont>     for table
cont> ;
SQL>
SQL> grant select
cont>     on table PAYSYS.rdb$default_audit_table
cont>     to m_smith, b_lee, s_jain
cont> ;
SQL>
SQL> grant select, insert, delete, update
cont>     on table PAYSYS.rdb$default_audit_table
cont>     to paysys_admin
cont> ;
SQL>
SQL> grant show
cont>     on table PAYSYS.rdb$default_audit_table
cont>     to PUBLIC
cont> ;
SQL>
.
.
.
SQL> create table paysys.CONTROL_TABLE
cont>     (identifier_value integer generated by default as identity
cont>     );
SQL>
SQL> show protection on table paysys.CONTROL_TABLE;
Protection on Table PAYSYS.CONTROL_TABLE
  (IDENTIFIER=PAYSYS_ADMIN,ACCESS=SELECT+INSERT+UPDATE+DELETE)
  (IDENTIFIER=[RDB,S_JAIN],ACCESS=SELECT)
  (IDENTIFIER=[DEV,B_LEE],ACCESS=SELECT)
  (IDENTIFIER=[AUDITOR,M_SMITH],ACCESS=SELECT)
  (IDENTIFIER=[ADMIN,DATABASE],ACCESS=SELECT+INSERT+UPDATE+DELETE+
    SHOW+CREATE+ALTER+DROP+DBCTRL+REFERENCES)
  (IDENTIFIER=[*,*],ACCESS=SHOW)
SQL>


```

Database Vault Feature

The security policy implemented through ACLs can be overridden at runtime by a suitably privileged OpenVMS user. That is, an OpenVMS power user might be able to attach and select data from a table even if they do not have database or table access granted by an access control entry. This override ability can be limited on a per database level by enabling the DATABASE VAULT attribute using ALTER DATABASE ... DATABASE VAULT IS ENABLED. When DATABASE VAULT is enabled, only the access control lists and database system privileges are used to determine access to database objects. This can prevent accidental override of the security policy. See the *Oracle Rdb Server SQL Reference Manual* DATABASE VAULT Appendix for further details.

Syntax





Arguments

- **ALL PRIVILEGES**

All privileges will be granted (or revoked) from the listed users and roles.

- **system-privilege**

Refer to Table 11–1, System Privileges for a list of supported system privileges.

- **TO username**

TO role-name

TO PUBLIC

Specifies the user name, role name, or the PUBLIC user to which you want to grant the system privilege. The PUBLIC user is the user name associated with all anonymous users who access the database.

Note

Oracle recommends that you only grant system privileges to trusted users. If system privileges are granted to roles then only assign those roles to trusted users.

If the database is defined as SECURITY CHECK IS INTERNAL and the user or role name exists as an operating system user or rights identifier, Oracle Rdb will automatically create the user or role name when you issue the GRANT statement.

- **FROM username**

FROM role-name

FROM PUBLIC

Specifies the user, role, or the PUBLIC user from which the specified role is to be revoked.

Database System Privileges

System privileges are associated with specific users (CREATE USER) and roles (CREATE ROLE) within the database. They are assigned by a user with SECURITY privilege on the database by the GRANT statement and removed by the REVOKE statement.

For example,

```

SQL> create user J_JONES identified externally;
SQL> create role DB_PROGRAMMER identified externally;
SQL>
SQL> grant create any sequence, create any procedure,
cont> create any temporary table to DB_PROGRAMMER;
SQL> grant create any trigger, create any index to J_JONES;
  
```


In this example, the database uses the default SECURITY CHECKING IS EXTERNAL therefore the granting of the rights identifier DB_PROGRAMMER must be performed by OpenVMS. Use the AUTHORIZE utility to grant the rights identifier to specific users.

```
$ run sys$system:authorize
UAF> GRANT/ID DB_PROGRAMMER J_JONES
UAF>
```

When J_JONES attaches to the database, they will receive the benefits of their own granted system privileges (if any) as well as those granted to the assigned role (DB_PROGRAMMER). This is true even for DATABASE VAULT protected databases.

If a database is defined as SECURITY CHECKING IS INTERNAL, then the GRANT statement is used to associate roles created by CREATE ROLE with specific users.

```
SQL> grant DB_PROGRAMMER to J_JONES;
SQL>
```

Refer to the *Oracle Rdb Server SQL Reference Manual* GRANT Statement: Roles section for more details and examples.

Implied System Privileges

In prior releases of Oracle Rdb, there were implied system privileges based on the database level CREATE, ALTER, DROP and SECURITY privileges that may have been granted to the user via the database ACL.

CREATE implies the permission to CREATE ANY object (except PROFILE, ROLE, and USER), and ALTER implies the permission to ALTER any object (except PROFILE, ROLE, and USER), DROP implies the permission to DROP ANY object (except PROFILE, ROLE, and USER), and SECURITY implies the permission to CREATE, ALTER, and DROP ANY PROFILE, ROLE or USER.

These implied system privileges can be displayed when attaching to a database and performing a SHOW PRIVILEGE ON DATABASE command. The implied system privileges (if any) are displayed. This is similar to the access privileges shown by this command. They reflect the ACL on the database as well as inherited access based on special privileges such as DBADM and OpenVMS privileges.

```
SQL> show privileges on database rdb$dbhandle;
Privileges on Alias RDB$DBHANDLE
  (IDENTIFIER=[RDB,RDBUSER2],ACCESS=SELECT+CREATE+ALTER+DROP)

Current system privileges:
  Granted Create Any
    COLLATING SEQUENCE, DOMAIN, ASSERTION, SESSION, OUTLINE, PROCEDURE,
    TABLE, SEQUENCE, VIEW
  Granted Alter Any
    COLLATING SEQUENCE, DOMAIN, ASSERTION, DATABASE, OUTLINE
  Granted Drop Any
    COLLATING SEQUENCE, DOMAIN, ASSERTION, DATABASE, OUTLINE
SQL>
```

Note that ALTER and DROP for primary database objects are not implicitly inherited from the database ALTER and DROP database privileges. This is because these operations are controlled by the object's own access control list.

Fine Tuning

Although system privileges are now available in Oracle Rdb, the database administrator is not required to use them. The database will operate as it did in previous releases.

Making use of this enhanced privilege system requires that the database administrator perform these tasks.

- Use CREATE USER for each database user to which system privileges need to be granted.
- Use CREATE ROLE for any rights identifiers used to fine tune access.

Not all users and roles that access the database need to be created in the database as Rdb will still use them as in prior releases to select matching access control entries. However, users and roles must be created to allow the database administrator to grant system privileges as these objects are used to store the current privilege set.

- Use the GRANT statement to manage the CREATE ANY, ALTER ANY, DROP ANY, and TRUNCATE ANY privileges and assign them to users and roles.
- Use REVOKE on the DATABASE ALIAS to remove the CREATE, ALTER, DROP or SECURITY privileges that were previously granted to those uses. This step is required so that those database level privileges do not interfere with the fine control of system privileges.

This change will limit those users according to the new security policy. The SHOW USER and SHOW ROLE statements will display all the granted system privileges. The SHOW PRIVILEGES ON DATABASE will show the augmented system privileges based on:

- Database ACL entries,
- Granted system privileges for the attaching user,
- Granted system privileges for all rights (roles) assigned to the user,
- OpenVMS privileges (when DATABASE VAULT is enabled there will be none used)

All access control is established at ATTACH time so changes made to the USER, ROLE or DATABASE access will not have an effect until the next database ATTACH.

Note

Oracle recommends that you only grant system privileges to trusted users. If system privileges are granted to roles then only assign those roles to trusted users.

Table 11–1 System Privileges

Operation Type	Object Type	Description
ALL PRIVILEGES		Can be used to GRANT or REVOKE all system privileges to a USER or ROLE.
ALTER ANY ...	ASSERTION	Permits the holder to execute the RDO CHANGE CONSTRAINT command and the SQL ALTER CONSTRAINT and COMMENT ON CONSTRAINT statements.
	CATALOG	Permits the holder to execute the ALTER CATALOG statement. The database must have multiscHEMA enabled; ALTER DATABASE ... MULTISCHEMA IS ON;
	COLLATING SEQUENCE	Permits the holder to execute the ALTER COLLATING SEQUENCE statement.
	DOMAIN	Permits the holder to execute the ALTER DOMAIN statement.
	DATABASE	Permits the holder to execute the ALTER DATABASE and COMMENT ON DATABASE statements.
	INDEX	Permits the holder to execute the ALTER INDEX statement.
	OUTLINE	Permits the holder to execute the ALTER OUTLINE statement.
	PROCEDURE	Permits the holder to execute the ALTER FUNCTION, ALTER MODULE and ALTER PROCEDURE statements.
	PROFILE	Permits the holder to execute the ALTER PROFILE and ALTER DEFAULT PROFILE statements. When the profile exists, the holder can also execute the CREATE OR REPLACE PROFILE and CREATE OR REPLACE DEFAULT PROFILE statements.
	ROLE	Permits the holder to execute the ALTER ROLE statement.
	SEQUENCE	Permits the holder to execute the ALTER SEQUENCE statement. When the sequence exists, the holder can also execute the CREATE OR REPLACE SEQUENCE statement.
	SCHEMA	Permits the holder to execute the ALTER SCHEMA statement. The database must have multiscHEMA enabled: ALTER DATABASE ... MULTISCHEMA IS ON;
	STORAGE MAP	Permits the holder to execute the ALTER STORAGE MAP statement.

(continued on next page)

Table 11–1 (Cont.) System Privileges

Operation Type	Object Type	Description
	SYNONYM	Permits the holder to execute the ALTER SYNONYM statement. When the synonym exists, the holder can also execute the CREATE OR REPLACE SYNONYM statements. The database must have synonyms enabled: ALTER DATABASE ... SYNONYMS ARE ENABLED;
	TABLE	Permits the holder to execute the ALTER TABLE statement.
	TEMPORARY TABLE	Permits the holder to execute the ALTER TEMPORARY TABLE statement or CREATE INFORMATION TABLE statement when CREATE ANY TABLE privilege is not granted.
	TRIGGER	Permits the holder to execute the ALTER TRIGGER statement.
	USER	Permits the holder to execute the ALTER USER statement.
	VIEW	Permits the holder to execute the ALTER VIEW statement. When the view exists, the holder can also execute the CREATE OR REPLACE VIEW statement.
CREATE ...	SESSION	Permits the holder to execute ATTACH, CONNECT, DECLARE ALIAS, SET SESSION AUTHORIZATION, and other session starting statements.
CREATE ANY ...	ASSERTION	Permits the holder to execute the RDO DEFINE CONSTRAINT command. The SQL equivalent to DEFINE CONSTRAINT would be a table level constraint. Such definitions are managed by TABLE privileges, therefore this privilege does not apply to SQL.
	CATALOG	Permits the holder to execute the CREATE CATALOG statement. The database must have multiscHEMA enabled: ALTER DATABASE ... MULTISCHEMA IS ON;
	COLLATING SEQUENCE	Permits the holder to execute the CREATE COLLATING SEQUENCE statement.
	DOMAIN	Permits the holder to execute the CREATE DOMAIN statement.
	INDEX	Permits the holder to execute the CREATE INDEX statement.
	OUTLINE	Permits the holder to execute the CREATE OUTLINE statement.

(continued on next page)

Table 11–1 (Cont.) System Privileges

Operation Type	Object Type	Description
	PROCEDURE	Permits the holder to execute the CREATE FUNCTION, CREATE MODULE and CREATE PROCEDURE statements.
	PROFILE	Permits the holder to execute the CREATE PROFILE and CREATE DEFAULT PROFILE statements.
	ROLE	Permits the holder to execute the CREATE ROLE statement.
	SEQUENCE	Permits the holder to execute the CREATE SEQUENCE statement.
	SCHEMA	Permits the holder to execute the CREATE SCHEMA statement. The database must have multiscHEMA enabled: ALTER DATABASE ... MULTISCHEMA IS ON;
	STORAGE MAP	Permits the holder to execute the CREATE STORAGE MAP statement.
	SYNONYM	Permits the holder to execute the CREATE SYNONYM statement. The database must have synonyms enabled: ALTER DATABASE ... SYNONYMS ARE ENABLED;
	TABLE	Permits the holder to execute the CREATE TABLE statement.
	TEMPORARY TABLE	Permits the holder to execute the CREATE TEMPORARY TABLE and CREATE INFORMATION TABLE statements when CREATE ANY TABLE privilege is not granted.
	TRIGGER	Permits the holder to execute the CREATE TRIGGER statement.
	USER	Permits the holder to execute the CREATE USER statement.
	VIEW	Permits the holder to execute the CREATE VIEW statement.
DROP ANY ...	ASSERTION	Permits the holder to execute the RDO DELETE CONSTRAINT command or the SQL DROP CONSTRAINT statement.
	CATALOG	Permits the holder to execute the DROP CATALOG statement. The database must have multiscHEMA enabled: ALTER DATABASE ... MULTISCHEMA IS ON;
	COLLATING SEQUENCE	Permits the holder to execute the DROP COLLATING SEQUENCE statement.
	DOMAIN	Permits the holder to execute the DROP DOMAIN statement.
	DATABASE	Permits the holder to execute the DROP DATABASE statement.

(continued on next page)

Table 11–1 (Cont.) System Privileges

Operation Type	Object Type	Description
	INDEX	Permits the holder to execute the DROP INDEX statement.
	OUTLINE	Permits the holder to execute the DROP OUTLINE statement.
	PROCEDURE	Permits the holder to execute the DROP FUNCTION, DROP MODULE and DROP PROCEDURE statements.
	PROFILE	Permits the holder to execute the DROP PROFILE and DROP DEFAULT PROFILE statements.
	ROLE	Permits the holder to execute the DROP ROLE statement.
	SEQUENCE	Permits the holder to execute the DROP SEQUENCE statement.
	SCHEMA	Permits the holder to execute the DROP SCHEMA statement. The database must have multischema enabled: ALTER DATABASE ... MULTISCHHEMA IS ON;
	STORAGE MAP	Permits the holder to execute the DROP STORAGE MAP statement.
	SYNONYM	Permits the holder to execute the DROP SYNONYM statement. The database must have synonyms enabled: ALTER DATABASE ... SYNONYMS ARE ENABLED;
	TABLE	Permits the holder to execute the DROP TABLE statement.
	TEMPORARY TABLE	Permits the holder to execute the DROP TEMPORARY TABLE statement or DROP INFORMATION TABLE statement when DROP ANY TABLE privilege is not granted.
	TRIGGER	Permits the holder to execute the DROP TRIGGER statement.
	USER	Permits the holder to execute the DROP USER statement.
	VIEW	Permits the holder to execute the DROP VIEW statement.
TRUNCATE ...	ANY TABLE	Permits the holder to execute the TRUNCATE TABLE statement. This privilege effectively allows the user to temporarily disable triggers during the TRUNCATE operation and assume DELETE access to the table.
	TABLE	This privilege is similar to the TRUNCATE ANY TABLE privilege but requires that the user also be granted DELETE access to the table. Permits the user to execute the TRUNCATE TABLE statement without further privilege checking for BEFORE and AFTER DELETE triggers.

Usage Notes

- You must have the SECURITY privilege on the database to grant a system privilege to a user or a role.
- You must have the SECURITY privilege on the database to revoke a system privilege from a user or a role.
- The TEMPORARY TABLE class of privileges is considered a subset of TABLE and allows the database administrator to grant privileges to a user but only for logical tables not physical (base) tables.
- If the user has CREATE ANY TABLE, then the CREATE ANY TEMPORARY TABLE privilege is not required. Similarly, ALTER ANY TEMPORARY TABLE and DROP ANY TEMPORARY TABLE are not used if ALTER ANY TABLE or DROP ANY TABLE is granted.
- The SHOW PRIVILEGES ON DATABASE statement displays the current active system privileges. This is based on the current user system privileges, system privileges inherited from granted roles, inherited privileges from the database access control list and OpenVMS process privileges.

It does not display any privileges for features which are not enabled for the database. CATALOG and SCHEMA privileges will not be displayed if MULTISCHEMA is not enabled. SYNONYM privileges will not be displayed if SYNONYMS are not enabled. However, the SHOW USER and SHOW ROLE statements will display all granted privileges even if that privilege has no application in the current database configuration.

11.1.7 Database Vault Feature

This release of Oracle Rdb introduces the DATABASE VAULT functionality.

The goal of DATABASE VAULT is to avoid accidental database access by an OpenVMS privileged user when the database security policy (ACL) should prevent such access.

This feature allows the database administrator to enforce an access policy for all attached database users by disabling the use of OpenVMS privileges as overrides to the database access control list.

DATABASE VAULT can be enabled by any of these commands.

- The SQL CREATE DATABASE ... DATABASE VAULT IS ENABLED statement.
- The SQL ALTER DATABASE ... DATABASE VAULT IS ENABLED statement.
- The SQL IMPORT DATABASE ... DATABASE VAULT IS ENABLED statement.

Note

If a database with DATABASE VAULT enabled is exported, then an IMPORT DATABASE will implicitly execute the DATABASE VAULT IS ENABLED action without that clause being required on the statement.

- RMU/SET DATABASE/DATABASE_VAULT=ENABLED command

DATABASE VAULT can be disabled by any of these commands.

- The SQL ALTER DATABASE ... DATABASE VAULT IS DISABLED statement.
- The SQL IMPORT DATABASE ... DATABASE VAULT IS DISABLED statement.
- RMU/SET DATABASE/DATABASE_VAULT=DISABLED command

Before executing any of these commands, the user must be granted (at least temporarily) the rights identifier RDBVMS\$DATABASE_VAULT_MANAGER that is added to the system during installation.

For example, the SET RIGHTS_LIST DCL command can be used to temporarily enable it.

```
$ SET RIGHTS_LIST /ENABLE RDBVMS$DATABASE_VAULT_MANAGER
$
$ rmu/set database mf_personnel/database_vault=enable
%RMU-I-MODIFIED, Database state modified
%RMU-W-DOFULLBCK, full database backup should be done to ensure future
recovery
$
$ SET RIGHTS_LIST /DISABLE RDBVMS$DATABASE_VAULT_MANAGER
```

Please refer to the *Oracle Rdb Server SQL Reference Manual* Appendix J for more details. This includes a description of the new DBVAULT audit class that can be used to audit changes to the DATABASE VAULT settings of a database.

11.1.8 SET FLAGS Keyword for Hash Join Feature - HASHING

This release of Oracle Rdb has added a flag to control the HASH JOIN feature of the Rdb optimizer. The HASHING flag can be used with the SET FLAGS statement or the RDMSS\$SET_FLAGS logical name to enable this feature. When enabled, the optimizer will attempt to engage the HASH JOIN feature during query solution. The default is NOHASHING(JOINS).

The following example shows the HASHING flag in use.

```
SQL> set flags 'strategy,detail(2)';
SQL> set flags 'hashing(joins)';
SQL>
SQL> select e.employee_id, e.birthday, jh.job_start
cont> from employees e, job_history jh
cont> where e.employee_id = jh.employee_id
cont> and jh.job_end is null
cont> ;
Tables:
  0 = EMPLOYEES
  1 = JOB_HISTORY
Conjunct: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID
Hash Q1
Outer Build
Match Key:0.EMPLOYEE_ID
Get      Retrieval by index of relation 0:EMPLOYEES
Index name EMP_EMPLOYEE_ID [0:0]
Inner Probe
Match Key:1.EMPLOYEE_ID
Conjunct: MISSING (1.JOB_END)
Get      Retrieval by index of relation 1:JOB_HISTORY
Index name JH_EMPLOYEE_ID [0:0]
Table=0:EMPLOYEES #Buckets=131 #Hits=61 #Collisions=39 #Dups=0 #Dups_Chain=0
Load_Factor= 4.656488549618321E-001
E.EMPLOYEE_ID E.BIRTHDAY JH.JOB_START
```

```

00164          28-Mar-1947    21-Sep-1981
00165          15-May-1954     8-Mar-1981
.
.
.
100 rows selected
SQL>

```

To disable the flag, use 'NOHASHING(JOINS)'. The setting is displayed by the SHOW FLAGS statement.

See also the new logical name RDM\$ENABLE_HASH_JOIN. Defining this logical name to true ("T", "t", "Y", "y" or "1") instructs the Rdb optimizer to try to use in-memory HASH JOIN to solve queries.

11.1.9 JOIN BY HASH Clause in CREATE OUTLINE Statement

This release of Oracle Rdb adds a new JOIN BY HASH clause to the CREATE OUTLINE statement and the OPTIMIZE OUTLINE clause of the select statement.

The following example shows the new syntax and the resulting query strategy.

```

SQL> create outline QO_1
cont> id '352E2736F133A6A322A3C935DB2CBE12'
cont> mode 0
cont> as (
cont>   query (
cont>     subquery (
cont>       EMPLOYEES 0 access path index EMP_EMPLOYEE_ID
cont>       join by hash to
cont>       SALARY_HISTORY 1 access path index SH_EMPLOYEE_ID
cont>     )
cont>   )
cont> )
cont> compliance optional;
SQL>
SQL> set flags 'strategy,detail(2)';
SQL>
SQL> select
cont>   e.employee_id, sh.salary_start, sh.salary_amount
cont> from
cont>   employees e
cont>   inner join
cont>   salary_history sh on (e.employee_id = sh.employee_id
cont>                        and sh.salary_end is null)
cont> optimize using QO_1
cont> ;
~S: Outline "QO_1" used
Tables:
  0 = EMPLOYEES
  1 = SALARY_HISTORY
Conjunct: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID
Hash Q1
Outer Build
Match Key:0.EMPLOYEE_ID
Index only retrieval of relation 0:EMPLOYEES
Index name EMP_EMPLOYEE_ID [0:0]
Inner Probe
Match Key:1.EMPLOYEE_ID
Conjunct: MISSING (1.SALARY_END)
Get Retrieval by index of relation 1:SALARY_HISTORY
Index name SH_EMPLOYEE_ID [0:0]
Table=0:EMPLOYEES #Buckets=131 #Hits=61 #Collisions=39 #Dups=0 #Dups_Chain=0
Load_Factor= 4.656488549618321E-001

```


E.EMPLOYEE_ID	SH.SALARY START	SH.SALARY_AMOUNT
00164	14-Jan-1983	\$51,712.00
00165	1-Jul-1982	\$11,676.00
.		
.		
.		

Note

This syntax, JOIN BY HASH, cannot be applied remotely to an older version of Oracle Rdb. An error such as this will be returned.

```
SQL> create outline QO_1
cont> id '352E2736F133A6A322A3C935DB2CBE12'
cont> mode 0
cont> as (
cont>   query (
cont>     subquery (
cont>       EMPLOYEES 0 access path index EMP_EMPLOYEE_ID
cont>       join by hash to
cont>       SALARY_HISTORY 1 access path index SH_EMPLOYEE_ID
cont>     )
cont>   )
cont> )
cont> compliance optional;
%SQL-F-UNSUPVER, Operation is unsupported for version of database
-SQL-F-UNSUPFEATURE, feature JOIN BY HASH is not supported
```

11.1.10 Hash Join Feature

Status: BETA

This release of Oracle Rdb includes a new optimization method known as Hash Join. A Hash Join is performed by hashing (mapping) one set of data into virtual memory based on the join columns and reading the other table to probe into the hash table to locate matching rows.

Typically, a Hash Join has a lower cost compared to the alternate of sorting when the hash table can be held entirely in memory, with the total cost amounting to very little other than the cost of reading the data sets. The cost rises if the hash table has to be spilled over to a temporary file.

Hash Join is only used for equi-joins. In general, Hash Join is a better solution for joining large numbers of rows in an equi-join.

Applications that in the past used Match Join might (unknowingly) rely on the implicit use of SORT during the query solution. However, as no implicit SORT is performed, the data might appear in a different order with Hash Join. Adding an ORDER BY will sort the result data but no longer sort the inputs to the join.

Note

Not all queries will be solved using HASH JOIN if this optional feature is enabled. Use the SET FLAGS 'STRATEGY,DETAIL(3)' to see a report of the current restrictions which cause the HASH JOIN method to be rejected.

Enabling HASH JOIN

The optimizer, by default, does not try the HASH JOIN method in solutions. This action must be enabled by the programmer in one of the following ways. The optimizer will then include HASH JOIN as part of its solution matching, but it may be rejected for various reasons.

- Define the logical name `RDMS$ENABLE_HASH_JOIN`.
Defining this logical name to true ("T", "t", "Y", "y" or "1") instructs the Rdb optimizer to try to use in-memory HASH JOIN to solve queries.
- Defining the logical name `RDMS$SET_FLAGS` or using the `SET FLAGS` statement with the string 'HASHING(JOINS)'. See Section 11.1.8 for more details.
- The new `JOIN BY HASH` clause in `CREATE OUTLINE` statement and the `OPTIMIZE OUTLINE` clause of the select statement. See Section 11.1.9 for more details.
- Specifying the `OPTIMIZE FOR HASH JOIN` on the select statement.

```
SQL> select e.employee_id, (sh.salary_end - sh.salary_start) month (3)
cont> from employees e, salary_history sh
cont> where e.employee_id = sh.employee_id
cont>         and e.employee_id <= '00164'
cont>         and sh.salary_end is not null
cont> optimize for hash join
cont> ;
.
.
.
```

Example

This example shows the strategy used by the optimizer when HASH JOIN is enabled.

```
SQL> set flags 'STRATEGY,DETAIL(2)';
SQL> set flags 'HASHING(JOINS)';
SQL>
SQL> select e.employee_id, (sh.salary_end - sh.salary_start) month (3)
cont> from employees e, salary_history sh
cont> where e.employee_id = sh.employee_id
cont>         and e.employee_id <= '00164'
cont>         and sh.salary_end is not null
cont> ;
Tables:
  0 = EMPLOYEES
  1 = SALARY_HISTORY
Conjunct: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID
Hash Q1
Outer Build
Match Key:0.EMPLOYEE_ID
Index only retrieval of relation 0:EMPLOYEES
Index name EMP_EMPLOYEE_ID [0:1]
Keys: 0.EMPLOYEE_ID <= '00164'
Inner Probe
Match Key:1.EMPLOYEE_ID
Conjunct: NOT MISSING (1.SALARY_END)
Conjunct: 1.EMPLOYEE_ID <= '00164'
Get Retrieval by index of relation 1:SALARY_HISTORY
Index name SH_EMPLOYEE_ID [0:1]
Keys: 1.EMPLOYEE_ID <= '00164'
Table=0:EMPLOYEES #Buckets=47 #Hits=1 #Collisions=0 #Dups=0 #Dups_Chain=0
```

```

Load_Factor= 2.127659574468085E-002
E.EMPLOYEE_ID
00164          008
00164          006
00164          016
3 rows selected
SQL>

```

Now try the same query with NOHASHING.

```

SQL> set flags 'NOHASHING(JOINS)';
SQL>
SQL> select e.employee_id, (sh.salary_end - sh.salary_start) month (3)
cont> from employees e, salary_history sh
cont> where e.employee_id = sh.employee_id
cont>        and e.employee_id <= '00164'
cont>        and sh.salary_end is not null
cont> ;
Tables:
  0 = EMPLOYEES
  1 = SALARY_HISTORY
Conjunct: 0.EMPLOYEE_ID = 1.EMPLOYEE_ID
Match Q1
Outer loop      (zig-zag)
Match_Key:0.EMPLOYEE_ID
Index_Key:EMPLOYEE_ID
Index only retrieval of relation 0:EMPLOYEES
Index name EMP_EMPLOYEE_ID [0:1]
Keys: 0.EMPLOYEE_ID <= '00164'
Inner loop      (zig-zag)
Match_Key:1.EMPLOYEE_ID
Index_Key:EMPLOYEE_ID
Conjunct: NOT MISSING (1.SALARY_END)
Conjunct: 1.EMPLOYEE_ID <= '00164'
Get Retrieval by index of relation 1:SALARY_HISTORY
Index name SH_EMPLOYEE_ID [0:1]
Keys: 1.EMPLOYEE_ID <= '00164'
E.EMPLOYEE_ID
00164          008
00164          006
00164          016
3 rows selected
SQL>

```

11.1.11 New ALTER DATABASE ... LOAD ACL IDENTIFIERS Clause

In this release of Oracle Rdb, the database administrator can automatically and simply create users and roles in the database. This clause, **LOAD ACL IDENTIFIERS**, is part of the **ALTER DATABASE** statement and can be run as often as necessary to add new users and roles derived from the existing access control lists (ACLs) granted to the database and database objects.

The following example shows this clause on a sample database:

```

SQL> attach 'filename WAREHOUSE';
SQL>
SQL> show roles;
Roles in database with filename WAREHOUSE
  No roles found
SQL> show users;
Users in database with filename WAREHOUSE
  No users found
SQL>
SQL> disconnect all;
SQL>
SQL> alter database
cont>      filename WAREHOUSE
cont>
cont>      load acl identifiers
cont> ;
SQL>
SQL> attach 'filename WAREHOUSE';
SQL>
SQL> show roles;
Roles in database with filename WAREHOUSE
  CDD$EXTENDER
  CDD$SYSTEM
  STORES_CUST_READABLE
  STORES_EXTRACT_TEXT
  STORES_MAIL_TEXT
  STORES_PRINT_TEXT
  STORES_USER
SQL> show users;
Users in database with filename WAREHOUSE
  FLEE
  ISMITH
  JJONES
  KSTJOHN
  WH_QUERY_1
  WH_QUERY_2
  WH_QUERY_3
  WH_QUERY_4
  WH_QUERY_5
SQL>
SQL> disconnect all;
SQL>

```

When the database administrator uses the GRANT statement to give access to users and OpenVMS rights identifiers (aka roles), they are recorded in the access control lists for each object; database, table, view, column, sequence, module, and routine. This clause of the ALTER DATABASE statement reads every ACL in the database and creates USER and ROLE definitions if necessary.

Note

Some access control entries (ACEs) may use OpenVMS group identifiers (PUBLIC, [*,*], [ADMIN,*], [*,], etc), or special modifier rights identifiers (BATCH, DIALUP, INTERACTIVE, LOCAL, NETWORK, REMOTE) which are not valid users and roles - these will be ignored by the LOAD ACL IDENTIFIERS clause.

In addition to the DBADM privilege required to use ALTER DATABASE, this clause also requires SECURITY on the database. Alternately, the user must be granted the ALTER ANY DATABASE, CREATE ANY USER and CREATE ANY ROLE database system privilege.

11.1.12 ALTER TABLE Statement Allows Change to READ ONLY Table

This release of Oracle Rdb adds the ability to change a table to READ ONLY access. Once committed, no other application or interactive SQL session may modify rows in the table. You can issue database definition statements (DDL) as long as they do not modify any table data. Operations on indices associated with the table are allowed when the table is in READ ONLY mode. To revert to a read-write table, the clause READ WRITE can be applied.

While the table is READ ONLY the following restrictions apply:

- The data manipulation statements INSERT, UPDATE, DELETE may not modify rows in the table.
- Table updates via LIST cursor may fail during the OPEN or CLOSE statement depending on the cursor declaration.
- The SELECT statement using the FOR UPDATE clause may fail because it tries to apply UPDATE semantics.

```
SQL> select * from SAMPLE_TABLE for update;
%RDB-E-READ_ONLY_REL, relation SAMPLE_TABLE was reserved for read access;
updates not allowed
```

- The TRUNCATE TABLE statement is not permitted to truncate rows from the table.
- ALTER TABLE ... ADD COLUMN is permitted unless a DEFAULT is added for the new column (either explicitly or implicitly from a domain reference). In this case, Rdb would normally execute an UPDATE statement to include the default into each pre-existing row.
- Most other database definition statements are permitted. For instance, CREATE INDEX, ALTER INDEX ... REBUILD ALL PARTITIONS, DROP INDEX, can all be performed while the table is in this state.

The following example shows the diagnostic reported by Oracle Rdb.

```
SQL> alter table SAMPLE_TABLE
cont>   read only
cont> ;
SQL>
SQL> show table (comment) SAMPLE_TABLE;
Information for table SAMPLE_TABLE

Comment on table SAMPLE_TABLE:
Samples table

Table is set READ ONLY

SQL> truncate table SAMPLE_TABLE;
%RDB-E-NO META UPDATE, metadata update failed
-RDB-E-READ_ONLY_REL, relation SAMPLE_TABLE was reserved for read access;
updates not allowed
SQL> commit;
SQL>
```

11.1.13 NULLS FIRST and NULLS LAST Options for ORDER BY Clause

This release of Oracle Rdb adds the NULLS FIRST and NULLS LAST options to the ORDER BY clause. These options control the ordering of the NULL values relative to the ordering of the key data. This is demonstrated by the four examples shown below.

```
SQL> select employee_id, salary_start, salary_end, salary_amount
cont> from salary_history sh
cont> where employee_id = '00164'
cont> order by salary_end asc nulls first
cont> ;
EMPLOYEE_ID  SALARY_START  SALARY_END  SALARY_AMOUNT
00164         14-Jan-1983  NULL        $51,712.00
00164         5-Jul-1980   2-Mar-1981  $26,291.00
00164         2-Mar-1981   21-Sep-1981 $26,291.00
00164         21-Sep-1981  14-Jan-1983 $50,000.00
4 rows selected
SQL>
SQL> select employee_id, salary_start, salary_end, salary_amount
cont> from salary_history sh
cont> where employee_id = '00164'
cont> order by salary_end asc nulls last
cont> ;
EMPLOYEE_ID  SALARY_START  SALARY_END  SALARY_AMOUNT
00164         5-Jul-1980   2-Mar-1981  $26,291.00
00164         2-Mar-1981   21-Sep-1981 $26,291.00
00164         21-Sep-1981  14-Jan-1983 $50,000.00
00164         14-Jan-1983  NULL        $51,712.00
4 rows selected
SQL>
SQL> select employee_id, salary_start, salary_end, salary_amount
cont> from salary_history sh
cont> where employee_id = '00164'
cont> order by salary_end desc nulls first
cont> ;
EMPLOYEE_ID  SALARY_START  SALARY_END  SALARY_AMOUNT
00164         14-Jan-1983  NULL        $51,712.00
00164         21-Sep-1981  14-Jan-1983 $50,000.00
00164         2-Mar-1981   21-Sep-1981 $26,291.00
00164         5-Jul-1980   2-Mar-1981  $26,291.00
4 rows selected
SQL>
SQL> select employee_id, salary_start, salary_end, salary_amount
cont> from salary_history sh
cont> where employee_id = '00164'
cont> order by salary_end desc nulls last
cont> ;
EMPLOYEE_ID  SALARY_START  SALARY_END  SALARY_AMOUNT
00164         21-Sep-1981  14-Jan-1983 $50,000.00
00164         2-Mar-1981   21-Sep-1981 $26,291.00
00164         5-Jul-1980   2-Mar-1981  $26,291.00
00164         14-Jan-1983  NULL        $51,712.00
4 rows selected
SQL>
```

11.1.14 Enhancements to RMU Unload After_Image (LogMiner) Interface

This release of Oracle Rdb adds support for the following new options for the RMU Unload After_Image (aka LogMiner) command.

- The RMU Unload After_Image command can now output an XML script containing the updates to the selected table or tables.
- When the FORMAT qualifier specifies one of DELIMITED_TEXT, DUMP or XML, the TRIM option can also be used to trim leading and/or trailing spaces.
- The SYMBOLS qualifier now accepts the keyword LOCAL (default) or GLOBAL. The default behavior is to generate local scope DCL symbols. When Symbols=GLOBAL is used, these symbols have global scope.

11.1.14.1 XML Option to FORMAT Qualifier

When using FORMAT=XML, the following options can also be specified:

- CHARACTER_ENCODING_XML

When using RMU Unload After_Image Format=XML, the XML header record will, by default, use the character encoding "ISO-8859-1". For example, this will appear in the header of the XML file.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
```

This encoding (ISO-8859-1) is Latin 1 and covers encoding of many European character sets. However, this encoding is not adequate if you use other character encoding for Asian languages, or languages not covered by this ISO Standard.

This release of Oracle Rdb adds an option, CHARACTER_ENCODING_XML, that allows the command procedure to specify an alternate character encoding. For example, if you wish to have the XML header describe UTF8, then specify the qualifier /FORMAT=(XML,CHAR="utf-8").

```
<?xml version="1.0" encoding="utf-8"?>
```

- DATA_XML_NULL

This option accepts one of the following keywords which control the output of NULL column values: DROP, NIL_ATTRIBUTE, or EMPTY. If this option is not specified, the default is EMPTY as shown in the following example.

This is a fragment of the XML data generated by RMU Unload After_Image with the qualifier /FORMAT=(XML,TRIM=TRAILING) defaulting to DATA_XML_NULL=EMPTY.

```
<ROW>
  <RDB_LM_ACTION>M</RDB_LM_ACTION>
  <RDB_LM_RELATION_NAME>SAMPLE2</RDB_LM_RELATION_NAME>
  <RDB_LM_RECORD_TYPE>32</RDB_LM_RECORD_TYPE>
  <RDB_LM_DATA_LEN>65254</RDB_LM_DATA_LEN>
  <RDB_LM_NBV_LEN>3</RDB_LM_NBV_LEN>
  <RDB_LM_DBK>60:14417:0</RDB_LM_DBK>
  <RDB_LM_START_TAD>2020-05-11T10:20:50.88</RDB_LM_START_TAD>
  <RDB_LM_COMMIT_TAD>2020-05-11T10:20:50.89</RDB_LM_COMMIT_TAD>
  <RDB_LM_TSN>712</RDB_LM_TSN>
  <RDB_LM_RECORD_VERSION>1</RDB_LM_RECORD_VERSION>
  <IDENT>101</IDENT>
  <COMMENT/>
  <DETAILS/>
</ROW>
```

If DATA_XML_NULL is specified as DROP, then that column's value is omitted from the XML record. If DATA_XML_NULL is specified as NIL_ATTRIBUTE, then the XML tag attribute for the column is output as xsi:nil="true".

- **TRIM**

This option allows the data values for the columns to have trailing and leading spaces and horizontal tab characters removed from the columns.

11.1.14.2 TRIM Option

When the FORMAT selected is one of DELIMITED_TEXT, DUMP or XML, then RMU can be instructed to trim trailing and/or leading spaces and horizontal tab characters from the columns.

- The default when FORMAT is XML or DELIMITED_TEXT is no trimming. The default when FORMAT is DUMP is TRIM_TRAILING.
- The TRIM option is not compatible with FORMAT=BINARY and FORMAT=TEXT.
- One of the following keywords can be specified for TRIM: TRAILING, LEADING and BOTH. If TRIM is specified without qualification, then TRAILING is assumed.

The following example shows the use of format XML with the TRIM=BOTH option.

```
$      RMU/UNLOAD-
        /AFTER_IMAGE -
        /INCLUDE=ACTION: (COMMIT,DELETE,NOMODIFY) -
        /LOG-
        /TABLE=(name=SAMPLE2, output=SAMPLE3.DAT) -
        /FORMAT=(XML,TRIM=BOTH) -
        /ORDER_AIJ_FILES-
        USER1:[TESTER.LOGMINER] LOGMINER_DB -
        USER1:[TESTER.LOGMINER] AIJ_BU_*.BAIJ
$
```

Usage Notes

- When unloading rows with columns that have many trailing spaces, then using FORMAT=(XML,TRIM) or FORMAT=(DELIMITED_TEXT,TRIM) can, in some cases, reduce the size of the output file without loss of significant data.
- When dumping rows with long columns that have many trailing spaces, then using /FORMAT=(DUMP,TRIM) can significantly reduce the size of the output file. Therefore, RMU Unload After_Image implicitly enables TRIM=TRAILING.

11.1.14.3 SYMBOLS Qualifier

The SYMBOLS qualifier now accepts the keyword LOCAL (default) or GLOBAL. The default behavior is to generate local scope DCL symbols. When Symbols=GLOBAL is used, these symbols have global scope.

Arguments

Symbols

Symbols=LOCAL

Symbols=GLOBAL

NoSymbols

Specifies whether DCL symbols are to be created, indicating information about records extracted for each table.

The default is Symbols, which causes local symbols to be created. Use Symbols=GLOBAL to have RMU define global symbols. Use NoSymbols to prevent creation of any DCL symbols.

If a large number of tables are being unloaded, too many associated symbols may be created and the CLI symbol table space can become exhausted. The error message "LIB-F-INSCLIMEM, insufficient CLI memory" is returned in this case. Specify the Nosymbols qualifier to prevent creation of the symbols.

11.1.15 Named Partition Support for RESERVING Clause

This release of Oracle Rdb adds the ability to use named partitions in the RESERVING clause of the SET TRANSACTION or DECLARE TRANSACTION statements. In prior versions, only partition numbers were allowed.

The partition names might be system generated (as shown below for the EMPLOYEES_MAP from the MF_PERSONNEL database) or they can be defined as part of the CREATE STORAGE MAP statement.

The following example shows the partition numbers as well as the system generated partition names under the *Partition information for storage map* output.

```
SQL> show storage map employees_map
      EMPLOYEES_MAP
For Table:      EMPLOYEES
Placement Via Index:  EMPLOYEES_HASH
Partitioning is:  UPDATABLE

Partition information for storage map:
Compression is:  ENABLED
Partition: (1) SYS_P00079
Storage Area: EMPIDS_LOW
Partition: (2) SYS_P00080
Storage Area: EMPIDS_MID
Partition: (3) SYS_P00081
Storage Area: EMPIDS_OVER

SQL>
```

Usage Notes

- The PARTITION clause accepts a list of partition names or a list of partition ordinal values. You may not mix numeric and named notations.
- Duplicate partition names in the RESERVING clause will cause an exception. Review the RESERVING clause and correct the partition names.

```

set transaction
  read write
  evaluating
    job_history_foreign1 at verb time
    ,salary_history_foreign1 at verb time
  reserving
    employees partition (SYS_P00080, SYS_P00080) for exclusive write
;
%RDB-E-BAD_TPB_CONTENT, invalid transaction parameters in the transaction
parameter block (TPB)
-RDMS-E-DUPPARTNAME, partition SYS_P00080 for table EMPLOYEES already used

```

- **Unknown partition names in the RESERVING clause (which might occur due to a change in the storage map definition) will cause an exception. Use the SHOW STORAGE MAP statement to review the partition names.**

```

create module mod_testing1a
  language sql
  procedure proc_xa ();
begin not atomic
set transaction
  read write
  evaluating salary_history_foreign1 at verb time
  reserving employees partition (SYS_P00080, "UNKNOWN", SYS_P00081)
    ,departments for protected write;

  commit;
end;
end module;
%RDB-E-NO_META_UPDATE, metadata update failed
-RDB-E-BAD_TPB_CONTENT, invalid transaction parameters in the transaction
parameter block (TPB)
-RDMS-F-PARTNEXTS, partition "UNKNOWN" does not exist in this map or index
"EMPLOYEES_MAP"

```

11.2 Obsolete Features

11.2.1 RMU Backup No Longer Supports HUFFMAN or LZSS Compression, Use ZLIB Instead

This release of Oracle Rdb removes the compression options HUFFMAN and LZSS from the RMU Backup and RMU Backup After_Journal commands.

These older compression algorithms are much slower than the default ZLIB compression. Orders of magnitude more CPU is required in some cases. If you receive a diagnostic as shown in the following example, then Oracle recommends accepting the default, changing the DCL command procedure, or RMU PLAN file to explicitly state ZLIB. Additionally, ZLIB compression allows the database administrator to determine levels of compression efficiency, from less time to more effective compression. Refer to *Oracle RMU Reference Manual* for more details.

The following example shows the new diagnostic reported by RMU.

```

$ rmu/backup-
  /disk_file-
  /list_plan=plan_1.plan-
  /execute-
  /compress=lzss-
  /parallel=Executor_Count=3 -
  sql$database -
  sav_11.rbf,sav_12.rbf,sav_13.rbf
%RMU-E-NOTSUPFORVER, The function COMPRESSION.LZSS is not supported for Oracle
Rdb V7.4-100
-RMU-I-COMPUSEZLIB, use the default, or specify ZLIB compression
%RMU-F-FTL_BCK, Fatal error for BACKUP operation at 18-JUN-2020 15:45:23.27
$

```

Also note that neither HUFFMAN nor LZSS are accepted by the RMU Set After_Journal Backups qualifier. That command has always required the preferred ZLIB algorithm for compression.

Documentation Corrections, Additions and Changes

This chapter describes problems, restrictions and additional documentation relating to Oracle Rdb version 7.4.1.4.

Note

Documentation for all Oracle Rdb logical names is available using the `RMU/SHOW LOGICAL_NAMES /DESCRIPTION` command. The command accepts wildcard logical-names (OpenVMS style * and % characters). Omit the logical-name parameter to see a complete list of relevant logical names, and their current values. For example,

```
$ RMU/SHOW LOGICAL_NAMES /UNDEFINED
"RDMS$AUTO_READY" = Undefined
"RDM$BIND_ABS_GLOBAL_STATISTICS" = Undefined
"RDM$BIND_ABS_LOG_FILE" = Undefined
.
.
.
"SORTWORKX" = Undefined
"SORTWORKY" = Undefined
"SORTWORKZ" = Undefined
$
```

The `/DESCRIPTION` qualifier can then be used to view more details for the selected logical names.

```
$ RMU/SHOW LOGICAL_NAMES *workfiles /DESCRIPTION
"RDMS$BIND_SORT_WORKFILES" = Undefined
```

This logical name specifies how many work files the OpenVMS Sort utility uses if work files are required. The default is 2 and the maximum number is 36. The work files can be individually located by defining the `SORTWORKn` logical names (where n is from 0 through 9, and A through Z).

You can increase the efficiency of Oracle Rdb sort operations, which use the OpenVMS SORT utility, by assigning the location of the temporary sort work files to different devices. These assignments are made by using up to logical names, `SORTWORK0` through `SORTWORK9`, and `SORTWORKA` through `SORTWORKZ`.

Note: Interfaces such as RMU Load, RMU Verify, RMU Optimize After Journal, SQL IMPORT DATABASE and RDO IMPORT also use this logical name but do not support values greater than 10.

```
$
```

12.1 Documentation Corrections

12.1.1 Oracle Rdb Release 7.4 New Features Manual

A document has been provided which contains all of the New Features Chapters from all Rdb 7.4 Release Notes. This document will be included in saveset A of the Rdb kit and installed to SYSSHELP. It is called RDB_NEWFEATURES_74xx and will be available in postscript, text and PDF format.

12.1.2 Oracle Rdb Timer Queue Limit Increased

Bug 28451930

In prior releases of Oracle Rdb the limit for timer queue entries was set at 100. This may cause applications and server processes to fail if the number of attaches exceeds this limit. Often multiple timer queue entries are required for some attaches.

These timer queue entries are not associated with the process level TQELM set for the process. This internal Oracle Rdb limit controls a fixed sized vector used in the Oracle Rdb Kernel.

In Oracle Rdb Release 7.4.1.1 this internal limit has been increased to 128.

12.1.3 Undocumented /TRANSACTION=EXCLUSIVE Option for RMU Populate_Cache

The EXCLUSIVE option was not documented in prior releases of Oracle Rdb. This is a revised section for the RMU Populate_Cache Command and will appear in a future update to the *Oracle RMU Reference Manual*.

- Transaction_Type=option

Allows you to specify the transaction mode for the transactions used to perform the populate cache operation. Valid options are:

- Automatic
- Exclusive
- Read_Only
- Noread_Only

You must specify an option if you use this qualifier.

If you do not specify this qualifier, the default is Transaction_Type=Automatic. This qualifier specifies that Oracle RMU is to determine the transaction mode. For example, if READ ONLY transactions can not be performed, an equivalent READ WRITE transaction will be started.

The Transaction_Type=Read_Only qualifier specifies the transactions used to perform the analyze operation be set to read-only mode.

The Transaction_Type=Noread_Only qualifier specifies that the transactions used for the analyze operation be set to read/write mode. You might select this option if you want to avoid the growth of snapshot files that occurs during a long running read-only transaction and are willing to incur the cost of increased locking that occurs during a read/write transaction.

The Transaction_Type=Exclusive readies all accessed storage areas in EXCLUSIVE mode. Therefore, it avoids the growth of snapshot files (no updates will be made to any snapshot file) and will use limited locking.

Note

Using Transaction_Type=Exclusive might interfere with other actions occurring on the database; especially note that ONLINE backup and some actions of the Row Cache server are incompatible with this Transaction_Type.

12.1.4 Action of DECLARE TRANSACTION When SQL\$DATABASE is Defined

Bug 5705204

The following clarification is missing from the DECLARE TRANSACTION Statement documentation in the *Oracle Rdb Server SQL Reference Manual*.

- In Dynamic SQL and Interactive SQL, the DECLARE TRANSACTION statement is applied to all attached databases. When those databases are disconnected, the declared default transaction is discarded. If there are no attached databases, then the DECLARE TRANSACTION is used as a session wide default.

When the first executable statement in a Dynamic SQL or an Interactive SQL session is a DECLARE TRANSACTION statement and the logical name SQL\$DATABASE is defined, then the DECLARE statement will implicitly attach to that database. The declared default transaction is applied to that database attach. A subsequent DISCONNECT statement will disconnect from the database and discard the default transaction.

12.1.5 RDB\$USAGE Field Values

The following information is missing from the *Oracle Rdb Server SQL Reference Manual*.

The table Rdb\$INTERRELATIONS records much of the dependency information when one object references another in the database. Such information is used by DROP ... RESTRICT statements to prevent an object being deleted when it is required by some other object. For instance, a function may use one or more columns from a table in a query. That table and its columns will be recorded with a value in RDB\$USAGE of 'Storage Map'.

Many reported errors include text from the RDB\$USAGE field to explain the type of dependency preventing the DROP from succeeding. These text strings are described in Table 12–1, Rdb\$USAGE Field Values.

Table 12–1 Rdb\$USAGE Field Values

Field Value	Description
Computed Field	A Computed by or Automatic column references this table, view, column or function
Constraint	Constraint definition references table, view, column, sequence or function
Storage Map	Storage map references table and column
View	View definition requires table, view, column, sequence or function
View Field	View column requires table, view, column, sequence or function
Trigger	Trigger definition requires table, view, column, sequence or function
RelConstraint	A table (relation) constraint references a table
Domain Constraint (VALID IF)	A domain constraint (or VALID IF) references this routine or sequence
Requires	This table, temporary table (with module name), or index is used by a query outline
Procedure	Procedure definition requires table, view, column, sequence or function
Function	Function definition requires table, view, column, sequence or function
Default Txn Reserving	A stored module DECLARE TRANSACTION references a table or view in the RESERVING clause
Default Txn Evaluating	A stored module DECLARE TRANSACTION references a constraint in the EVALUATING clause
Lang Semantics	A stored function, procedure or trigger uses wildcard for column list. This includes SELECT *, or INSERT with an omitted column list
Cast As Domain	A CAST function referenced a domain name
Temp Table Using Domain	A DECLARE LOCAL TEMPORARY TABLE used a domain for a columns data type
Computed Column in Temp Table	A computed by or automatic column defined by a DECLARE LOCAL TEMPORARY TABLE or DECLARE LOCAL TEMPORARY VIEW references this object
Module Variable Default Value	A module global DECLARE statement used a DEFAULT clause. Table, view, function, domain and sequence dependencies are recorded
Referenced by Synonym	When a synonym is created, a dependency is stored
Default Value	A table column uses a DEFAULT clause. Table, view, function, domain and sequence dependencies are recorded
Constraint Index	When SET FLAGS 'AUTO_INDEX' is active, any constraint definition will define an index matching the columns of the constraint
Module Variable	This module variable uses this domain
Routine Parameter	Not currently used. Reserved for future use

(continued on next page)

Table 12–1 (Cont.) Rdb\$USAGE Field Values

Field Value	Description
Temp Table Reference	A DECLARE LOCAL TEMPORARY TABLE references this table in the LIKE clause
Storage Map Function	When CREATE STORAGE MAP is executed, a system routine is created to reflect the mapping. Those column dependencies are recorded

12.1.6 Updated Documentation for RDMS\$BIND_CODE_OPTIMIZATION Logical

The logical name RDMS\$BIND_CODE_OPTIMIZATION has different uses on Alpha and Integrity systems.

- **Alpha Systems**

In prior releases of Oracle Rdb, queries and procedures were rendered as small Alpha instruction sequences which were used at runtime to implement all or part of the query strategy. This meant that executable code and data could reside on any Alpha page (8192 byte) sections.

Recent performance testing shows that OpenVMS for Alpha running on an Alpha emulator (such as Stromasys, Inc's CHARON-AXP) can greatly benefit from isolating the generated code instructions from updatable data areas. This change in memory management can be enabled by defining this logical to the value 3. Please note that such systems will typically require more virtual memory for the alternate memory management.

This logical name can be defined to the value (0) to disable the alternate memory management or defined as (3) to enable the alternate memory management. The default is 0 on Alpha systems and other values are undefined.

For more information on Stromasys, Inc, please check their web site at: <https://www.stromasys.com/solutions/charon-axp/>

- **Integrity systems**

When Oracle Rdb first became available on the Integrity platform, it used a portable instruction interpreter. Now, by default, Oracle Rdb generates native IA64 instructions to implement part of the query and procedure execution code. This native code execution in general delivers improved performance but at the expense of more virtual memory usage.

This logical name can be defined to the value (0) to disable the native instruction compiler and use the interpreted instructions, or defined as (2) to enable the native IA64 instruction generator. The default is 2 on Integrity systems and other values are undefined.

The RMU Show Logical_Names command can be used to display the latest documentation for this and other logical names used by Oracle Rdb.

```
$ rmu /show Logical_Names RDMS$BIND_CODE_OPTIMIZATION/Description
"RDMS$BIND_CODE_OPTIMIZATION" = Undefined
...etc...
```


12.1.7 New RDM\$BIND_AIJ_BUFFER_POOL_COUNT Logical

The description of this new logical was inadvertently left out of the Rdb Release 7.3.1.3 Release Notes.

A new RDM\$BIND_AIJ_BUFFER_POOL_COUNT logical has been added for use by various components of Oracle Rdb. This logical name can be used to change the number of buffers in the internal buffer pool used by the RMU/UNLOAD AFTER_JOURNAL command to select and process journal records from database transactions. If the logical name is not defined, a default of 10 buffers will be used.

The RDM\$BIND_AIJ_BUFFER_POOL_COUNT logical allows a whole integer number of buffers between 4 and 16 to be specified for the RMU/UNLOAD AFTER_JOURNAL buffer pool. The current value of the RDM\$BIND_AIJ_BUFFER_POOL_COUNT logical can be displayed using the RMU/SHOW LOGICAL_NAMES command if it is defined as a system logical.

The following examples show that the RMU/SHOW LOGICAL_NAMES command can be used to show if RDM\$BIND_AIJ_BUFFER_POOL_COUNT is currently defined as a system logical. They also show that RDM\$BIND_AIJ_BUFFER_POOL_COUNT can be defined as a process logical, and that if a value is defined for RDM\$BIND_AIJ_BUFFER_POOL_COUNT that is not a whole integer between 4 and 16, an error will be returned by the RMU/UNLOAD AFTER_JOURNAL command and the RMU/UNLOAD AFTER_JOURNAL command will be aborted without processing any After Image Journal records.

```
$ RMU/SHOW LOGICAL_NAMES RDM$BIND_AIJ_BUFFER_POOL_COUNT
$
$!
$ RMU/SHOW LOGICAL_NAMES/UNDEFINED RDM$BIND_AIJ_BUFFER_POOL_COUNT
"RDM$BIND_AIJ_BUFFER_POOL_COUNT" = Undefined
$!
$ DEFINE/SYSTEM RDM$BIND_AIJ_BUFFER_POOL_COUNT 4
$ RMU/SHOW LOGICAL_NAMES RDM$BIND_AIJ_BUFFER_POOL_COUNT
"RDM$BIND_AIJ_BUFFER_POOL_COUNT" = "4" (LNM$SYSTEM_TABLE)
$!
$ DEFINE/SYSTEM RDM$BIND_AIJ_BUFFER_POOL_COUNT 16
%DCL-I-SUPERSEDE, previous value of RDM$BIND_AIJ_BUFFER_POOL_COUNT has been
superseded
$ RMU/SHOW LOGICAL_NAMES RDM$BIND_AIJ_BUFFER_POOL_COUNT
"RDM$BIND_AIJ_BUFFER_POOL_COUNT" = "16" (LNM$SYSTEM_TABLE)
$!
$ DEFINE RDM$BIND_AIJ_BUFFER_POOL_COUNT 3
$ RMU/UNLOAD/AFTER_JOURNAL/INCL=ACT=(NOCOMMIT)/STAT=3600/LOG -
/RESTORE METADATA=LOGMINER_METADATA.TXT -
/TABLE=(NAME=TEST_TABLE,OUTPUT=RDB_LOGMINER_OUTPUT_FILE) -
TEST.AIJ
%RMU-I-LMMFRDCNT, Read 13542 objects from metadata file
"DEVICE:[DIRECTORY]LOGMINER_METADATA.TXT;1"
%RMU-I-UNLAIJFL, Unloading table TEST_TABLE to
DEVICE:[DIRECTORY]RDB_LOGMINER_OUTPUT_FILE.DAT;1
%RMU-F-BADBNDPRM, bad bind parameter RDM$BIND_AIJ_BUFFER_POOL_COUNT value "3"
-RMU-F-BADBOUNDS, value not in range 4 to 16
%RMU-F-FTL_RMU, Fatal error for RMU operation at 17-MAR-2019 09:58:08.42
-----
ELAPSED: 0 00:00:00.00 CPU: 0:00:00.00 BUFIO: 5 DIRIO: 0 FAULTS: 46
Table "TEST_TABLE" : 0 records written (0 modify, 0 delete)
Total : 0 records written (0 modify, 0 delete)
$!
$ DEFINE RDM$BIND_AIJ_BUFFER_POOL_COUNT 17
%DCL-I-SUPERSEDE, previous value of RDM$BIND_AIJ_BUFFER_POOL_COUNT has been
superseded
```

```

$ RMU/UNLOAD/AFTER JOURNAL/INCL=ACT=(NOCOMMIT)/STAT=3600/LOG -
/RESTORE_METADATA=LOGMINER_METADATA.TXT -
/TABLE=(NAME=TEST_TABLE,OUTPUT=RDB_LOGMINER_OUTPUT_FILE) -
TEST.AIJ
%RMU-I-LMMFRDCNT, Read 13542 objects from metadata file
"DEVICE:[DIRECTORY]LOGMINER_METADATA.TXT;1"
%RMU-I-UNLAIJFL, Unloading table TEST_TABLE to
DEVICE:[DIRECTORY]RDB_LOGMINER_OUTPUT_FILE.DAT;2
%RMU-F-BADBNDPRM, bad bind parameter RDM$BIND_AIJ_BUFFER_POOL_COUNT value "17"
-RMU-F-BADBOUNDS, value not in range 4 to 16
%RMU-F-FTL_RMU, Fatal error for RMU operation at 17-MAR-2019 09:59:39.02
-----
ELAPSED:      0 00:00:00.00 CPU: 0:00:00.01 BUFIO: 5 DIRIO: 0 FAULTS: 46
Table "TEST_TABLE" : 0 records written (0 modify, 0 delete)
Total : 0 records written (0 modify, 0 delete)
$!

```

12.1.8 RDMSTT Image Optionally Installed

If you plan on using the cluster capability of RMU/SHOW STATISTICS, Oracle recommends that you install the RDMSTT74.EXE image on OpenVMS with the appropriate privileges.

The RMONSTART74.COM command procedure provided by the Oracle Rdb installation allows the system manager to optionally install the RDMSTT74.EXE image at monitor startup time. To take advantage of this, you will need to edit the SYS\$STARTUP:RMONSTART74.COM procedure and remove the comment characters from the following lines:

```

$! CALL INSTALL_IMAGE SYS$COMMON:[SYSEXE]RDMSTT74.EXE
    "/OPEN/HEAD/PROT/PRIV=(CMKRNL,SYSPRV,SHARE) "

```

The command procedure SYS\$MANAGER:RMONSTOP74.COM will automatically remove this image if it is installed.

12.1.9 Some Optional System Tables Can Be Relocated Using a User Defined Storage Map

All system tables are mapped by default to the system storage area RDB\$SYSTEM. If the database is created with the DEFAULT STORAGE AREA clause, then some of these tables will automatically be created in the secondary system area. Additionally, there exists a set of optional system tables (which may not exist in all databases) which may be manually mapped to other storage areas.

To change the mapping for one (or more) of these system tables, you must follow these steps. See Table 12-2 for the list of these optional system tables that allow mapping.

1. Attach to the database. If you are creating a storage map for RDB\$CATALOG_SCHEMA or RDB\$SYNONYMS then you must attach with the option MULTISchema IS OFF. You should not execute any queries on the database as this may cause the system table to be locked and prevent the CREATE and ALTER STORAGE MAP statements from completing.
2. Create a storage map for the optional system table. Note that only those listed here are able to be re-mapped and you must use the specified storage map names.

Table 12–2 Optional System Tables and Their Storage Map Names

Table Name	Storage Map Name	Associated Feature
RDB\$CATALOG_SCHEMA	RDB\$CATALOG_SCHEMA_MAP	Multischema databases
RDB\$CHANGES	RDB\$CHANGES_MAP	Replication Option for Rdb
RDB\$CHANGES_MAX_TSER	RDB\$CHANGES_MAX_TSER_MAP	Replication Option for Rdb
RDB\$SYNONYMS	RDB\$SYNONYMS_MAP	Multischema databases
RDB\$TRANSFERS	RDB\$TRANSFERS_MAP	Replication Option for Rdb
RDB\$TRANSFER_RELATIONS	RDB\$TRANSFER_RELATIONS_MAP	Replication Option for Rdb
RDB\$WORKLOAD	RDB\$WORKLOAD_MAP	Workload Collection was Enabled

3. The storage map must be a simple storage map which simply describes the current state for this table, namely the name of the storage area in which the table resides. See the following example.

```
SQL> create storage map RDB$CHANGES_MAP
cont>   for RDB$CHANGES
cont>   store in RDB$SYSTEM;
```

The following restrictions apply to the created storage map for these special system tables:

- The storage map may not change the defaulted compression attributes
- The storage map may not specify the logical area thresholds
- The storage map may not be placed via an index
- The storage map may not vertically partition the table
- The storage map may only contain one storage area
- And it must be mapped to the default storage area (this may be RDB\$SYSTEM by default or the name of the user specified storage area using the DEFAULT STORAGE AREA clause during CREATE DATABASE)

4. Now that the storage map exists, you may use the ALTER STORAGE MAP statement to move the table to another area.

```
SQL> alter storage map RDB$CHANGES_MAP
cont>   store in RDB_CHANGES_AREA;
```

The following restrictions apply to the altered storage map for these special system tables:

- The storage map may not be placed via an index
- The storage map may only contain one storage area
- The storage map may not vertically partition the table
- The ALTER STORAGE MAP operation may require exclusive access to the database as well as the table while the table data is relocated.

5. These storage map attributes for system tables are not currently exported by the SQL EXPORT DATABASE statement. Therefore, if you EXPORT and IMPORT your database, you will need to repeat these steps to re-map any of these system tables. It is expected that this restriction will be removed in a future version of Oracle Rdb.

12.1.10 Recovering an Oracle Rdb Database After RMU/RESTORE/ONLY_ROOT

Bug 12595718

If it is necessary to recover a database following the RMU/RESTORE/ONLY_ROOT command, be sure that the transaction state of the database is correctly set in the database root by the RMU/RESTORE/ONLY_ROOT command. Otherwise transactions in journal files created before the RMU/RESTORE/ONLY_ROOT command will be ignored by the RMU/RECOVER command, which uses the transaction state of the database stored in the database root file to select which journaled transactions to recover from the After Image Journal (AIJ) files used by the RMU/RECOVER command. Important journaled updates to database parameters, such as the client sequence numbers maintained both in the database root (CLTSEQ) and the database system tables (RDB\$SEQUENCES), may be lost or made inconsistent.

The database should be verified both after the RMU/RESTORE/ONLY_ROOT command completes and after the RMU/RECOVER command completes. Please consult the documentation in the *Oracle RMU Reference Manual* for more information and examples of the options related to AIJ files and setting the database root transaction TSN and CSN values when the Rdb database root file (.rdb) is restored using the RMU/RESTORE/ONLY_ROOT command.

The TSN and CSN values set in the database root restored by the RMU/RESTORE/ONLY_ROOT command are displayed by an informational message if logging is enabled.

```
%RMU-I-SETRTTSNCSN, Setting Root Transaction Sequence TSN to 384,  
Commit Sequence CSN to 384
```

In this example, the /NOSET_TSN qualifier is used so that the TSN and CSN values of the restored root file are set to the values in the backup file used by the RMU/RESTORE/ONLY_ROOT command. As a result, the original journaled client sequence value of "21" is recovered by the RMU/RECOVER command executed after the RMU/RESTORE/ONLY_ROOT command.

```

$ sql$
SQL> attach 'file mf_personnel';
SQL> select RDB$SEQUENCE_NAME, RDB$NEXT_SEQUENCE_VALUE
cont> from rdb$sequences;
  RDB$SEQUENCE_NAME                RDB$NEXT_SEQUENCE_VALUE
  S                                21
1 row selected
exit;
$ delete mf_personnel.rdb;*
$ rmu/restore/only_root/log/NOSET TSN mf_personnel
%RMU-I-AIJRSTBEG, restoring after-image journal "state" information
%RMU-I-AIJRSTJRN, restoring journal "AIJ1" information
%RMU-I-AIJRSTSEQ, journal sequence number is "0"
%RMU-I-AIJRSTSUC, journal "AIJ1" successfully restored from file
  "DEVICE:[DIRECTORY]AIJ_ONE.AIJ;1"
%RMU-I-AIJRSTJRN, restoring journal "AIJ2" information
%RMU-I-AIJRSTNMD, journal has not yet been modified
%RMU-I-AIJRSTSUC, journal "AIJ2" successfully restored from file
  "DEVICE:[DIRECTORY]AIJ_TWO.AIJ;1"
%RMU-I-AIJRSTEND, after-image journal "state" restoration complete
%RMU-I-SETRTSNCSN, Setting Root Transaction Sequence TSN to 384,
  Commit Sequence CSN to 384
%RMU-I-AIJISON, after-image journaling has been enabled
%RMU-W-DOFULLBCK, full database backup should be done to ensure
  future recovery
%RMU-I-AIJRECEND, after-image journal "state" recovery complete
$
$ sql$
SQL> attach 'file mf_personnel';
SQL> select RDB$SEQUENCE_NAME, RDB$NEXT_SEQUENCE_VALUE
cont> from rdb$sequences;
  RDB$SEQUENCE_NAME                RDB$NEXT_SEQUENCE_VALUE
  S                                1
1 row selected
exit;
$ rmu/recover/out=recov.sav AIJ_ONE.aij,AIJ_TWO.aij
$
$ sql$
SQL> attach 'file mf_personnel';
SQL> select RDB$SEQUENCE_NAME, RDB$NEXT_SEQUENCE_VALUE
cont> from rdb$sequences;
  RDB$SEQUENCE_NAME                RDB$NEXT_SEQUENCE_VALUE
  S                                21
1 row selected
exit;
$

```

12.1.11 Oracle Rdb Position on NFS Devices

This release note describes the supported usage of the NFS (Network File System) mounted devices by the Oracle Rdb product. NFS devices appear in most regards as local mounted file systems but do not allow the same level of sharing as provided by local OpenVMS devices. In addition, these files reside on a non-OpenVMS system (for instance a Linux or Windows system) and are therefore outside any scheme used by Rdb to lock buffers and pages of the database.

Active System Files

When Rdb is actively using database files, these files require specific sharing and locking to guarantee database integrity and recovery. Therefore, because of the limitations of the NFS mounted devices, active files such as the database root (.rdb), storage areas (.rda), snapshot files (.snp), row cache work file (.rdc), after image journal files (.aij), and before image recovery journal (.ruj) must not reside on an NFS mounted device.

Archived Data Files

Files that are not part of the active system may be stored on NFS mounted devices. For example, RMU /BACKUP /AFTER_JOURNAL can be used to archive an after image journal to a target on an NFS device. Similarly, RMU /BACKUP can perform a full or incremental backup to an Rdb backup file (.rbf) on an NFS device and RMU /RESTORE can use that NFS mounted source for database recovery, along with archived after image files from an NFS device processed by RMU /RECOVER.

Other Miscellaneous Files

Other files that might be used by an Rdb installation include options files, application procedures and sources, backup journals, record definitions files (.rrd), unloaded database files (.unl), exported databases (.rbr), log files, and so on. These sequential files may be stored on and referenced by RMU and SQL commands from an NFS mounted device.

Setting Up NFS

Complete instructions for setting up an NFS mounted device is beyond the scope of this release note and customers are directed to use system specific documentation for the server platform and for OpenVMS systems. However, during testing with Oracle Rdb we noted the need for the following qualifiers for the TCPIP MOUNT command.

- Use /ADF=CREATE. This ensures that attributes (such as block size and record length) are preserved on the server.
- Use /STRUCTURE=5. This will emulate an ODS-5 device and therefore allow the most complete OpenVMS Files-11 On-Disk Structure emulation.
- Use /TRANSPORT=UDP. For example,

```
$ tcpip mount dnfs1:/host="test.company.com"/path="/scratch"  
/stru=5/serve=unix/adf/vers=2/tran=udp
```

Read Performance Issues

In versions of Oracle Rdb prior to Rdb V7.3.1.2, a significant performance issue exists when reading sequential files from NFS mounted devices. Oracle Rdb uses the RMS read-ahead (RAH) attribute to improve sequential reads but this has an adverse effect when referencing an NFS device. The latest release of Oracle Rdb works around this issue by disabling the use of read-ahead when referencing an NFS device and would be the preferred version when using NFS devices.

Disclaimer

This information is provided to answer customer questions and should not be read as an endorsement or guarantee for NFS systems. Oracle expects configuration, functional testing, performance testing, security and integrity of the NFS data to be performed by our customers.

12.1.12 RDM\$BIND_STAREA_EMERGENCY_DIR Logical Name

Bugs 19545970 and 3682207

RDM\$BIND_STAREA_EMERGENCY_DIR is a HOT STANDBY logical name that can be utilized when replicating the creation of a new storage area from a master database to its standby database.

RDM\$BIND_STAREA_EMERGENCY_DIR provides an alternate device and/or directory specification for the standby that can replace all or part of the master's file specification. Without the logical, the device and directory of the new storage area issued from the master must exist and match exactly on the standby. For example, on the master database we want to create a new starea, \$1SDGA11:[RDB_RANDOM.FOO]A1.RDA. We would issue the following command:

```
SQL> alter database file rdb_random$db
      add storage area a1 filename $1SDGA11:[RDB_RANDOM.FOO]A1.RDA;
```

If the standby did not have a device called \$1SDGA11, the replication would fail and the AIJ Log Roll-Forward Server (LRS) logfile would log the failure.

```
3-SEP-2014 16:22:26.94 - Replicating master FILID 19
3-SEP-2014 16:22:26.94 - Attempting to create starea
"$1SDGA11:[RDB_RANDOM.FOO]A1.RDA;1" ALQ=2808
3-SEP-2014 16:22:26.95 - Unable to create storage area. STATUS: 00DDA89C
3-SEP-2014 16:22:26.95 - No emergency directory defined
3-SEP-2014 16:22:26.95 - Failure reason: LRSSRV$CREATE_AREA_CALLBACK - Could
not create storage area
```

Suppose the target disk on the standby was \$1SDGA109 and we defined the logical RDM\$BIND_STAREA_EMERGENCY_DIR to point to that, as in the following example.

```
$ define/sys RDM$BIND_STAREA_EMERGENCY_DIR "$1SDGA109:"
$ create/dir $1SDGA109:[RDB_RANDOM.FOO]
```

The replication operation would succeed and the LRS logfile would show:

```
3-SEP-2014 15:42:45.65 - Attempting to create starea
"$1SDGA11:[RDB_RANDOM.FOO]A1.RDA;1" ALQ=2808
3-SEP-2014 15:42:45.67 - Unable to create storage area. STATUS: 00DDA89C
3-SEP-2014 15:42:45.67 - Using emergency area "$1SDGA109:[RDB_RANDOM.FOO]A1.RDA"
3-SEP-2014 15:42:45.67 - Attempting to create starea
"$1SDGA109:[RDB_RANDOM.FOO]A1.RDA" ALQ=2808
3-SEP-2014 15:42:45.68 - Starea creation successful
3-SEP-2014 15:42:45.70 - Attempting to create starea
"$1SDGA11:[RDB_RANDOM.FOO]A1.SNP;1" ALQ=404
3-SEP-2014 15:42:45.70 - Unable to create storage area. STATUS: 00DDA89C
3-SEP-2014 15:42:45.70 - Using emergency area "$1SDGA109:[RDB_RANDOM.FOO]A1.SNP"
3-SEP-2014 15:42:45.70 - Attempting to create starea
"$1SDGA109:[RDB_RANDOM.FOO]A1.SNP" ALQ=404
3-SEP-2014 15:42:45.71 - Starea creation successful
```

The RDM\$BIND_STAREA_EMERGENCY_DIR logical must:

- Exist on the standby system prior to the create storage area operation.
- Be defined in the LNM\$SYSTEM_TABLE table.
- Be a valid file specification.

All standby databases on the node where the logical is defined share its use.

This logical was added back in Oracle Rdb Release 7.0.2 but the documentation of the logical was omitted.

12.1.13 RDMS-F-FULLAIJBKUP, Partially-Journaled Changes Made

Bug 7669735

The Oracle Rdb and Oracle CODASYL DBMS Guide to Hot Standby Databases states: "You can stop replication operations by explicitly entering the Replicate After_Journal Stop command on either the standby or master database nodes. Stopping replication on either database terminates replication on both databases."

Although the RMU/REPLICATE AFTER_JOURNAL STOP command may be issued against either Master or Standby to shut down replication, we have determined that there is at least one scenario where the choice is important relating to restarting replication in the future.

If you do the following, the operation will fail with a 'FULLAIJBKUP' error when starting the Master.

1. Stop replication on the Standby.
2. Set the old standby to be the new Master.
3. Set the old Master to be the new Standby.
4. Attempt to restart replication.

This is expected behavior. If the Standby is stopped prior to the Master, Oracle Rdb cannot determine if there has been any network traffic from the Master between the time that the Standby and Master shut down. Since any such information would be lost and may lead to data inconsistencies, replication will not be started.

The workaround for this scenario would be to stop replication on the Master, not the Standby. Consider the following two examples (assuming that replication is currently active):

Example 1: Initially stopping Replication on the Standby.

```
$! Stopping Replication on the Standby:

$ RMU/REPLICATE AFTER STOP/WAIT/LOG STANDBY$DB:STANDBY_PERSONNEL
%RMU-I-HOTSTOPWAIT, stopping database replication, please wait
%RMU-I-LOGMODSTR, stopped master database AIJ Log Replication Server

$! Start Replication of the Standby db (which was previously the Master)

$ RMU/REPLICATE AFTER_JOURNAL START MASTER:MF_PERSONNEL.RDB -
  /CHECKPOINT=10 -
  /LOG -
  /WAIT -
  /BUFFERS=30 -
  /GAP_TIMEOUT=5 -
  /GOVERNOR=DISABLED -
  /MASTER_ROOT=STANDBY$DB:STANDBY_PERSONNEL.RDB -
  /ONLINE
%RMU-I-LOGMODSTR, started standby database AIJ Log Replication Server

$! Start Replication on the Master db (which was previously the Standby)
```



```

$ RMU/REPLICATE AFTER JOURNAL START STANDBY$DB:STANDBY_PERSONNEL.RDB -
  /CHECKPOINT=100 -
  /LOG -
  /WAIT -
  /CONNECT_TIMEOUT=5 -
  /STANDBY_ROOT=MASTER:MF_PERSONNEL.RDB -
  /SYNCHRONIZATION=COLD -
  /QUIET_POINT -
  /TRANSPORT=TCPIP
%RMU-I-LOGMODSTR, started AIJ Log Server
%RDMS-F-CANTSTARTLCS, error starting AIJ Log Catch-Up Server process
-RDMS-F-FULLAIJBKUP, partially-journalled changes made; database may not be
recoverable
%RMU-F-FATALRDB, Fatal error while accessing Oracle Rdb.
%RMU-F-FTL_RMU, Fatal error for RMU operation at 4-AUG-2014 14:19:17.78

```

Example 2: Initially stopping Replication on the Master.

```

$! Stopping Replication on the Master:

$ RMU/REPLICATE AFTER STOP/WAIT/LOG MASTER$DB:MF_PERSONNEL.RDB
%RMU-I-HOTSTOPWAIT, stopping database replication, please wait
%RMU-I-LOGMODSTR, stopped master database AIJ Log Replication Server

$! Start Replication of the Standby db (which was previously the Master)

$ RMU/REPLICATE AFTER JOURNAL START MASTER:MF_PERSONNEL.RDB -
  /CHECKPOINT=10 -
  /LOG -
  /WAIT -
  /BUFFERS=30 -
  /GAP_TIMEOUT=5 -
  /GOVERNOR=DISABLED -
  /MASTER_ROOT=STANDBY$DB:STANDBY_PERSONNEL.RDB -
  /ONLINE
%RMU-I-LOGMODSTR, started standby database AIJ Log Replication Server

$! Start Replication on the Master db (which was previously the Standby)

$ RMU/REPLICATE AFTER JOURNAL START STANDBY$DB:STANDBY_PERSONNEL.RDB -
  /CHECKPOINT=100 -
  /LOG -
  /WAIT -
  /CONNECT_TIMEOUT=5 -
  /STANDBY_ROOT=MASTER:MF_PERSONNEL.RDB -
  /SYNCHRONIZATION=COLD -
  /QUIET_POINT -
  /TRANSPORT=TCPIP
%RMU-I-LOGMODSTR, started AIJ Log Server
%RMU-I-LOGMODSTR, started master database AIJ Log Replication Server

```

The SYS\$HELP:RMU_MSG*.DOC has more information about the FULLAIJBKUP error:

```

FULLAIJBKUP,    partially-journalled changes made; database may not be
recoverable

Explanation:    Partially journalled changes have been made to the
database. This may result in the database being
unrecoverable in the event of database failure; that
is, it may be impossible to roll-forward the
after-image journals, due to a transaction mis-match or
attempts to modify objects that were not journalled.
This condition typically occurs as a result of
replicating database changes using the Hot Standby
feature.

```

User Action: IMMEDIATELY perform a full (not by-sequence) quiet-point AIJ backup to clear the AIJ journals, followed immediately by a full (no-quiet-point allowed) database backup.

12.1.14 Undocumented Hot Standby Logical Names

Bug 3264793

Table 12–3 Hot Standby Logical Names

Logical Name Description	Default Value	Minimum Value	Maximum Value
RDM\$BIND_ALS_LOG_REOPEN_SECS Defines the number of seconds after which the ALS output file will automatically be reopened.	0 seconds (will not be reopened automatically)	0 seconds	31449600 (1 year)
RDM\$BIND_ALS_LOG_REOPEN_SIZE Defines the number of blocks after which the ALS output file will automatically be reopened.	0 blocks (will not be reopened automatically)	0 blocks	Infinite
RDM\$BIND_HOT_ABS_SUSPEND_SHUTDOWN Defines whether or not the AIJ backup server (ABS) should be automatically suspended on graceful shutdown.	0	0	1
RDM\$BIND_HOT_CHECKPOINT Specifies the number of messages per server checkpoint interval. If specified, the first threshold to be exceeded (message count or elapsed time) will cause the checkpoint.	100	1	50000
RDM\$BIND_HOT_CHECKPOINT_INTERVAL Specifies a checkpoint interval, in minutes, to be used in addition to the /CHECKPOINT qualifier specified at Hot Standby startup. If specified, the first threshold to be exceeded (message count or elapsed time) will cause the LRS checkpoint.	0 minutes (don't use elapsed time)	0 minutes	10080 (7 days)
RDM\$BIND_HOT_IGNORE_NET_TIMEOUT Specifies whether or not to ignore network timeout parameters if the LRS process is still active.	0	0	1
RDM\$BIND_HOT_LOG_REOPEN_SECS Defines the number of seconds after which the AIJSERVER output file will automatically be reopened.	0 seconds (will not be reopened automatically)	0 seconds	604800 (1 week)
RDM\$BIND_HOT_LOG_REOPEN_SIZE Defines the number of blocks after which the AIJSERVER output file will automatically be reopened.	0 blocks (will not be reopened automatically)	0	Infinite
RDM\$BIND_HOT_NETWORK_ALT_NODE Defines the secondary network nodename to be used in the event of primary nodename network failure. This logical name allows you to specify an alternate routing pathway to the same standby database.	None		

(continued on next page)

Table 12–3 (Cont.) Hot Standby Logical Names

Logical Name Description	Default Value	Minimum Value	Maximum Value
RDM\$BIND_HOT_NETWORK_RETRY Specifies a network retry timeout interval.	120 seconds	0	1800 (30 minutes)
RDM\$BIND_LCS_AIJ_SCAN_IO_COUNT Defines the number of asynchronous I/O operations to be performed simultaneously during LCS catch-up.	64	1	128
RDM\$BIND_LCS_LOG_REOPEN_SECS Defines the number of seconds after which the LCS output file will automatically be reopened.	0 seconds (will not be reopened automatically)	0 seconds	31449600 (1 year)
RDM\$BIND_LCS_LOG_REOPEN_SIZE Defines the number of blocks after which the LCS output file will automatically be reopened.	0 blocks (will not be reopened automatically)	0 blocks	Infinite
RDM\$BIND_LCS_QUIET_TIMEOUT Defines the number of seconds to wait for the LCS process to obtain the standby database quiet-point.	600 seconds	0 seconds (wait indefinitely)	Infinite
RDM\$BIND_LCS_SYNC_COMMIT_MAX Defines the number of catch-up messages to synchronize with the standby database. A message may contain multiple transactions.	128 messages	32 messages	10000 messages
RDM\$BIND_LRS_LOG_REOPEN_SECS Defines the number of seconds after which the LRS output file will automatically be reopened.	0 seconds (will not be reopened automatically)	0 seconds	31449600 (1 year)
RDM\$BIND_LRS_LOG_REOPEN_SIZE Defines the number of blocks after which the LRS output file will automatically be reopened.	0 blocks (will not be reopened automatically)	0 blocks	Infinite
RDM\$BIND_LRS_QUIET_TIMEOUT Defines the number of seconds to wait for the LRS process to obtain the standby database quiet-point.	600	0 seconds (wait indefinitely)	Infinite
RDM\$BIND_STAREA_EMERGENCY_DIR Defines an alternate device and directory for the creation of storage areas on the standby database. The logical must be defined in the LNM\$SYSTEM_TABLE table and it is shared by all standby databases on that node.			

12.1.15 Corrections to the EDIT STRING Documentation

Bugs 17365476 and 17365597

- The *Oracle Rdb Server SQL Reference Manual*, Volume 1, incorrectly stated that fields following a quoted literal would have leading zeros trimmed if the literal ended with a space. This was incorrect. The trimming only takes place after a space formatting character (B).

This is the corrected text:

Oracle Rdb automatically trims leading zeros from the first numeric field in the output, and any numeric field following a space formatting character (B). The year (Y) and fractional seconds (*) format fields are never trimmed of leading zeros.

- To have SQL represent an OpenVMS date format without removing the leading zero from the Hour field, use the literal string for space rather than the space formatting character (B).

```
edit string 'YYYY-NN-DD" "RR:PP:QQ.**'
```

rather than

```
edit string 'YYYY-NN-DDBRR:PP:QQ.**'
```

- The formatting string ** represents the 100ths of a second field. Prior versions using a narrow field * would erroneously truncate the leading digits. This is corrected in this release, as the trailing digit is truncated.

12.1.16 Changes and Improvements to the Rdb Optimizer and Query Compiler

This release of Oracle Rdb introduces several new capabilities within the query compiler and the query optimizer. These changes fall generally under the title *query rewrite*, and allow the query compiler to present a simplified query for optimization and execution.

- CAST function elimination

In most cases, CAST actions must be executed at runtime to convert from the source data type to that specified by the CAST function. However, in some cases, the Rdb query compiler can eliminate or replace the CAST function with a literal value during query compile. This saves CPU time as the action is performed just once rather than once per row processed.

This replacement includes the following:

- When CAST of DATE (ANSI), DATE (VMS) or TIMESTAMP data types is performed to a compatible type of DATE or TIMESTAMP, then in many cases the CAST operator is not required.
- CAST of string literals to DATE (ANSI), DATE (VMS), TIME, TIMESTAMP and INTERVAL can be processed at compile time. For example, CAST('2013-1-1' AS DATE ANSI) is implicitly converted to a DATE literal DATE'2013-1-1'.
- CAST of small integer values is now done by the compiler. For example, CAST(1 AS SMALLINT) can be performed at compile time.
- CAST of fixed length (CHAR) literal strings to varying length strings (VARCHAR) is now processed by the compiler if the character set is the same and the target VARCHAR is long enough to hold the source string, as seen in the following example:

```
CAST('TABLE' AS VARCHAR(31))
```

- Constant Folding

Simple arithmetic expressions involving integer or floating point literals are evaluated by the query compiler. The overall effect is smaller executable code and some reduced CPU time for queries. FLOAT, REAL, and DOUBLE PRECISION values are combined to produce DOUBLE PRECISION results. Integer literals (with no fractional component) are combined to produce BIGINT results.

The side effect is that some expressions may now return DOUBLE PRECISION or BIGINT results where in prior versions they produced smaller precision results. This should not affect applications which fetch values into different data types as Oracle Rdb will perform an implicit conversion.

This optimization includes the following:

- * Addition (+)
- * Subtraction (-)
- * Multiplication (*)
- * Division (/)

Note that division is not performed at compile time if the divisor is a literal zero (0). Operations which are coded to explicitly divide by zero are probably expected to produce an error at runtime. Although using the SQL SIGNAL statement is now preferred, this technique has been used to terminate procedures when an incorrect input is encountered.

- Algebraic Rules

Additive identity (zero) can be added to an expression without changing the value. The query compiler will eliminate the literal zero (0) from the expression.

Multiply by zero will result in zero if the other operand is a not nullable expression. In this case, the expression will be replaced by zero.

Multiplicative identity (one) can be multiplied by an expression without changing the value. The query compiler will eliminate the literal one (1) from the expression.

The side effect is that some expressions may now return slightly different data types because the literal is no longer considered as part of the data type computation.

- Simple Predicate Elimination

When predicates include comparison of simple expressions, then the query compiler will attempt to eliminate them from the query predicate. For example, WHERE ('A' = 'A') will be replaced by TRUE, WHERE (2 <> 2) will be replaced with FALSE, and so on.

- Not Nullable Aware

The query compiler is now aware of which columns have a NOT NULL NOT DEFERRABLE constraint enabled. Additionally, this attribute is also implied from any PRIMARY KEY NOT DEFERRABLE constraints.

Using this knowledge, the query compiler can reduce (prune) the query expression. This list defines the ways in which this can occur:

- * When IS NULL is applied to a not nullable column or expression, then this predicate is replaced with FALSE.
- * When IS NOT NULL is applied to a not nullable column or expression, then this predicate is replaced with TRUE.

The side effect is that constraints for a table are now loaded for SELECT statements.

This optimization can be disabled using the SET FLAGS statement, or the RDMS\$SET_FLAGS logical name with the value NOREWRITE(IS_NULL). The default is REWRITE(IS_NULL).

- Replace comparisons with NULL

Queries that erroneously compare value expressions with NULL will now be replaced with a simplified UNKNOWN value. For example, a query that uses WHERE EMPLOYEE_ID = NULL will never find matching rows, because the results of the comparison (equals, not equals, greater than, less than, and so on) are always UNKNOWN.

This optimization can be disabled using the SET FLAGS statement, or the RDMS\$SET_FLAGS logical name with the value NOREWRITE(UNKNOWN). The default is REWRITE(UNKNOWN).

- Predicate Pruning

The AND, OR and NOT operators can be simplified if the logical expressions have been reduced to TRUE, FALSE or UNKNOWN expressions. Depending on the operation, the Rdb query compiler might be able to eliminate the Boolean operator and part of the expression.

This optimization can be disabled using the SET FLAGS statement, or the RDMS\$SET_FLAGS logical name with the value NOREWRITE(BOOLEANS). The default is REWRITE(BOOLEANS).

- CASE Expression Pruning

The prior transformation will also be applied to the Boolean WHEN expressions of a conditional expression (CASE, DECODE, NULLIF, COALESCE, NVL, NVL2, SIGN, ABS, and so on).

In some cases, the resulting conditional expression might resolve to an equivalent conditional expression with fewer branches (some WHEN ... THEN clauses being eliminated) or a simple expression with no conditional expression (all WHEN ... THEN clauses are eliminated).

- IN Operator Simplification

The IN operator using a subquery looks similar to the EXISTS boolean expression but it differs in its handling of NULL values. If the query compiler knows that neither source field nor the value set contain NULL, then the EXISTS expression can replace the IN operator. The EXISTS expression generates a better query solution in almost all cases.

This optimization can be disabled using the SET FLAGS statement, or the RDMS\$SET_FLAGS logical name with the value NOREWRITE(IN_CLAUSE). The default is REWRITE(IN_CLAUSE).

In most cases, the results of these optimizations will be transparent to the application. However, database administrators that use SET FLAGS 'STRATEGY,DETAIL' will notice new notations in the displayed strategy.

The following examples show the types of likely results.

In this example, the logical expression (1 = 2) is replaced with FALSE, the logical expression (1 = 1) is replaced with TRUE and the predicate is reduced to just the IS NULL (aka MISSING) check.

```

SQL> select last_name
cont> from employees
cont> where ((1 = 1) and employee_id is null)
cont>          or
cont>          ((1 = 2) and employee_id = '00164');
Tables:
  0 = EMPLOYEES
Conjunct: MISSING (0.EMPLOYEE_ID)
Get      Retrieval sequentially of relation 0:EMPLOYEES
0 rows selected

```

If there existed a NOT NULL NOT DEFERRABLE constraint on the EMPLOYEE_ID column, the expression can be further reduced because the NOT NULL constraint means the IS NULL test is always FALSE.

```

SQL> alter table EMPLOYEES
cont>      alter column EMPLOYEE_ID
cont>      constraint NN_EMPLOYEE_ID
cont>      NOT NULL
cont>      NOT DEFERRABLE
cont> ;
SQL>
SQL> select last_name
cont> from employees
cont> where ((1 = 1) and employee_id is null)
cont>          or
cont>          ((1 = 2) and employee_id = '00164');
Tables:
  0 = EMPLOYEES
Conjunct: FALSE
Get      Retrieval sequentially of relation 0:EMPLOYEES
0 rows selected
SQL>

```

REWRITE Flag

The SET FLAGS statement and the RDMS\$SET_FLAGS logical name can be used to enable or disable some of these rewrite actions. This flag primarily exists for Oracle to test the behavior of the query rewrite changes. It can be used by programmers to revert to pre-V7.3 behavior.

REWRITE enables each rewrite setting and NOREWRITE disables them. Additionally, keywords can be added to REWRITE and NOREWRITE to disable selective rewrite actions.

The following new keywords are added for this release of Oracle Rdb.

- BOOLEANS
- IN_CLAUSE
- IS_NULL
- UNKNOWN

12.1.17 Sorting Capabilities in Oracle Rdb

Oracle Rdb supports both the traditional OpenVMS SORT32 facility as well as simplified internal sort facility called QSORT.

QSORT

Use of QSORT preempts use of all other sorting algorithms. The QSORT algorithm is used if sorting is being done on a single key and if only a small amount of data is involved. The reason for this is that the other sorting algorithms, while handling larger data sets, have a certain amount of setup overhead which can be avoided in simple cases.

QSORT is used by default when:

- There is a single sort key.
- The number of rows to be sorted is 5000 or fewer. Note that this row limit can be adjusted.
- The buffer needed for caching the rows to be sorted is 409600 bytes (800 VMS pagelets) or less.
- The sort key is not floating point (REAL, FLOAT, or DOUBLE PRECISION).

Oracle Rdb supports two controls for the QSORT facility.

- RDMS\$BIND_MAX_QSORT_COUNT

This logical name controls the number of rows which will be buffered before resorting to the SORT32 interface.

The default value is 5000 rows.

- RDMS\$BIND_MAX_QSORT_BUFFER

This logical acts as a governor to RDMS\$BIND_MAX_QSORT_COUNT. If the rows being sorted are long, then the buffer size will reduce the count of candidate rows being used by QSORT. For example, even with the default values, a large row size (say 15304 bytes) will cause the actual limit on rows to be 26. This is displayed by the SORT display when SET FLAGS 'STRATEGY,SORT' is used.

```
SORT(18) SortId# 4, ----- qsort used
Records Input: 4      Record Limit: 26      LogRecLen Input: 15304
```

The default value is 409600 bytes.

How to Disable QSORT

To disable QSORT because of either anomalous or undesirable performance, the user can define either of the following logicals to zero (0):

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

How to Alter QSORT Usage

To change the usage of QSORT to evaluate behavior with other parameters, define new limits for the parameters as follows:

```
$ DEFINE RDMS$BIND_MAX_QSORT_BUFFER n
$ DEFINE RDMS$BIND_MAX_QSORT_COUNT m
```

The relationship between buffer size and row count is that the rows to be sorted are temporarily stored in the buffer until either it is full or the number of rows to be sorted exceeds the limit. For example, if the row to be sorted is 20 bytes long, 5000 of these rows would need 10,000 bytes which would easily fit into a 409600 byte buffer. However, if the rows to be sorted are 1,000 bytes long, only 409 of them can fit into a 409600 byte buffer and SORT32 will be invoked as soon as the sort algorithm is given row number 410.

12.1.18 RDM\$BIND_MAX_DBR_COUNT Documentation Clarification

Bugs 1495227 and 3916606

The *Oracle Rdb7 Guide to Database Performance and Tuning*, Volume 2, page A-18, incorrectly describes the use of the RDM\$BIND_MAX_DBR_COUNT logical.

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

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

The RDM\$BIND_MAX_DBR_COUNT logical name defines the maximum number of database recovery (DBR) processes to be simultaneously invoked by the database monitor for each database during a “node failure” recovery. This logical name applies only to databases that do not have global buffers enabled. Databases that utilize global buffers have only one recovery process started at a time during a “node failure” recovery.

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

Per-Database Value

The RDM\$BIND_MAX_DBR_COUNT logical name specifies the maximum number of database recovery processes to run at once for each database. For example, if there are 10 databases being recovered and the value for the RDM\$BIND_MAX_DBR_COUNT logical name is 8, up to 80 database recovery processes would be started by the monitor after a node failure.

The RDM\$BIND_MAX_DBR_COUNT logical name is translated when the monitor process opens a database. Databases need to be closed and reopened for a new value of the logical to become effective.

12.1.19 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 12–4 shows the TRANS_TPB table.

Table 12–4 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 12–5 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 12–5 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)	

12.1.20 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 12–6 shows the DATABASE table.

Table 12–6 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	

(continued on next page)

Table 12–6 (Cont.) Columns for Table EPC\$1_221_DATABASE

Column Name	Data Type	Domain
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 12–7 shows the REQUEST_ACTUAL table.

Table 12–7 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	

(continued on next page)

Table 12–7 (Cont.) Columns for Table EPC\$1_221_REQUEST_ACTUAL

Column Name	Data Type	Domain
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)	

(continued on next page)

Table 12–7 (Cont.) Columns for Table EPC\$1_221_REQUEST_ACTUAL

Column Name	Data Type	Domain
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	

(continued on next page)

Table 12–7 (Cont.) Columns for Table EPC\$1_221_REQUEST_ACTUAL

Column Name	Data Type	Domain
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 12–8 shows the TRANSACTION table.

Table 12–8 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	

(continued on next page)

Table 12–8 (Cont.) Columns for Table EPC\$1_221_TRANSACTION

Column Name	Data Type	Domain
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	

(continued on next page)

Table 12–8 (Cont.) Columns for Table EPC\$1_221_TRANSACTION

Column Name	Data Type	Domain
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	

(continued on next page)

Table 12–8 (Cont.) Columns for Table EPC\$1_221_TRANSACTION

Column Name	Data Type	Domain
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	

(continued on next page)

Table 12–8 (Cont.) Columns for Table EPC\$1_221_TRANSACTION

Column Name	Data Type	Domain
CROSS_FAC_3_END	INTEGER	
CROSS_FAC_7_END	INTEGER	
CROSS_FAC_14_END	INTEGER	

Table 12–9 shows the REQUEST_BLR table.

Table 12–9 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

12.2 Optimizer Changes

12.2.1 New Optimizations for the SQL IN Predicate

As described in *Section 8.1.7, New Implementation of the SQL IN Predicate*, SQL now uses a more compact representation of the SQL IN predicate. This presents new opportunities for Oracle Rdb 7.4.1.3 and later to improve incoming queries during compile and optimization.

Note

This new functionality can be disabled by defining the logical name `SQL$DISABLE_IN_CLAUSE`.

```
$ define /system SQL$DISABLE_IN_CLAUSE "Y"
```

Use RMU /SHOW LOGICAL_NAME SQL\$DISABLE_IN_CLAUSE
/DESCRIPTION to read more information.

Duplicate Entries in IN list are Eliminated

The new implementation allows the Rdb query compiler to eliminate duplicate literals, parameters, and column references while processing the IN predicate.

This first example shows the prior release's encoding of the query when there are many duplicate values.

```
SQL> set flags 'strategy,detail(2)';
SQL>
SQL> select *
cont> from T
cont> where a in (1,9,2,9,3,9,4,9,5,9,6,9,7,9,8,9)
cont> ;
Tables:
  0 = T
Conjunct: (0.A = 1) OR (0.A = 9) OR (0.A = 2) OR (0.A = 9) OR (0.A = 3) OR (0.A
          = 9) OR (0.A = 4) OR (0.A = 9) OR (0.A = 5) OR (0.A = 9) OR (0.A = 6)
          OR (0.A = 9) OR (0.A = 7) OR (0.A = 9) OR (0.A = 8) OR (0.A = 9)
Get      Retrieval sequentially of relation 0:T
.
.
.
... rows selected
SQL>
```

When this query is run using 7.4.1.3 (or later) shows the reduce query complexity.

```
SQL> select *
cont> from T
cont> where a in (1,9,2,9,3,9,4,9,5,9,6,9,7,9,8,9)
cont> ;
Tables:
  0 = T
Conjunct: (0.A = 1) OR (0.A = 9) OR (0.A = 2) OR (0.A = 3) OR (0.A = 4) OR (0.A
          = 5) OR (0.A = 6) OR (0.A = 7) OR (0.A = 8)
Get      Retrieval sequentially of relation 0:T
.
.
.
... rows selected
SQL>
```

When an index is present then the number of index lookups is minimized.

```

SQL> select *
cont> from T
cont> where a in (1,9,2,9,3,9,4,9,5,9,6,9,7,9,8,9)
cont> ;
Tables:
  0 = T
Conjunct: (0.A = 1) OR (0.A = 9) OR (0.A = 2) OR (0.A = 3) OR (0.A = 4) OR (0.A
          = 5) OR (0.A = 6) OR (0.A = 7) OR (0.A = 8)
Index only retrieval of relation 0:T
Index name  T_NDX [(1:1)9]
Keys: r0: 0.A = 8
      r1: 0.A = 7
      r2: 0.A = 6
      r3: 0.A = 5
      r4: 0.A = 4
      r5: 0.A = 3
      r6: 0.A = 2
      r7: 0.A = 9
      r8: 0.A = 1
.
.
.
... rows selected
SQL>

```

Note

Oracle has observed similar queries from customers that include many duplicate values, most likely generated through front-end query tools. This change was made to address these observed overheads.

Leading value-expression only evaluated once

In prior versions any complex expression used as the left-hand side of the IN predicate would be propagated through the equivalent OR and equal (=) expression and lead to excessive evaluation. External functions, SQL functions, subqueries, and so on would add to the cost of the query execution.

For example, a simple subquery would end up creating a multi-way join expression, one join leg for each value in the IN predicate list.

```

SQL> select *
cont> from T
cont> where (select b from S limit to 1 row) not in (1,9,2,9,3)
cont> ;
Tables:
  0 = T
  1 = S
  2 = S
  3 = S
  4 = S
  5 = S
Cross block of 6 entries  Q1
Cross block entry 1
  Conjunct: <agg0> <> 3
  Aggregate: 0:VIA (5.B) Q6
  Firstn: 1
  Get      Retrieval sequentially of relation 5:S
Cross block entry 2
  Conjunct: <agg1> <> 9
  Aggregate: 1:VIA (4.B) Q5
  Firstn: 1
  Get      Retrieval sequentially of relation 4:S
Cross block entry 3
  Conjunct: <agg2> <> 2
  Aggregate: 2:VIA (3.B) Q4
  Firstn: 1
  Get      Retrieval sequentially of relation 3:S
Cross block entry 4
  Conjunct: <agg3> <> 9
  Aggregate: 3:VIA (2.B) Q3
  Firstn: 1
  Get      Retrieval sequentially of relation 2:S
Cross block entry 5
  Conjunct: <agg4> <> 1
  Aggregate: 4:VIA (1.B) Q2
  Firstn: 1
  Get      Retrieval sequentially of relation 1:S
Cross block entry 6
  Get      Retrieval sequentially of relation 0:T
.
.
.
... rows selected
SQL>

```

The query compiler now generates a single evaluation of this expression thus reducing the query execution cost.

```

SQL> select *
cont> from T
cont> where (select b from S limit to 1 row) in (1,9,2,9,3)
cont> ;
Tables:
  0 = T
  1 = S
Cross block of 2 entries  Q1
Cross block entry 1
  Conjunct: (<agg0> = 1) OR (<agg0> = 9) OR (<agg0> = 2) OR (<agg0> = 3)
  Aggregate: 0:VIA (1.B) Q2
  Firstn: 1
  Get      Retrieval sequentially of relation 1:S
Cross block entry 2
  Get      Retrieval sequentially of relation 0:T

```

```
.  
. .  
. .  
... rows selected  
SQL>
```

Known Problems and Restrictions

This chapter describes problems and restrictions relating to Oracle Rdb and includes workarounds where appropriate.

13.1 Known Problems and Restrictions in All Interfaces

This section describes known problems and restrictions that affect all interfaces. This is not an exhaustive list.

13.1.1 OpenVMS TCPIP Issue May Cause Rdb Kit Installation to Fail

Current releases of OpenVMS may produce this error when executing the SET PASSWORD command from a process that was started using SSH, such as via PUTTY.

```
$ set password
%ACME-F-CONTACTSYSMGR, requested operation has failed; contact the system manager
$
```

This problem affects Oracle Rdb because new accounts are created during installation: RDBSREMOTE74, RDMALJ74 and RDMSTT74. Please refer to the Oracle Rdb for OpenVMS Installation Guide for further details on these accounts.

During the Oracle Rdb product installation, the OpenVMS VMSINSTAL.COM procedure uses SET PASSWORD/GENERATE which may fail during attempts to install Oracle Rdb due to an error in VSI TCPIP for SSH created processes. As a workaround Oracle recommends using SET HOST to perform the installation.

Please reference the following section in the *VSI OpenVMS Alpha Version 8.4-2L2 New Features and Release Notes*, reproduced here with permission. This workaround can be applied to both Alpha and Integrity systems.

VSI OpenVMS Alpha Version 8.4-2L2 New Features and Release Notes.

Order number: DO-DVANFR-01A

June 2017

Copyright (c) 2017 VMS Software, Inc., (VSI), Bolton Massachusetts, USA

2. ACMELOGIN images, which are bundled with the OpenVMS operating system, are a pre-requisite for LDAP authentication. However, with ACMELOGIN installed, the SET PASSWORD command will not work in SSH sessions, even though it works fine in other terminals like Telnet, DECnet, etc. This restriction is documented in the VSI TCP/IP 5.7 ECO5 release notes, as well as the ACME LDAP Installation Guide.

To overcome this limitation, define the following logical before starting SSH server:

```
$ DEFINE /SYSTEM TCPIP$SSH_SERVER_USE_LOGINOUT 1
```

This prompts SSH server to use [SYSEXEC]LOGINOUT.EXE image for authentication (as is done in Telnet and DECnet terminals). Now the "SET PASSWORD" command will work fine in SSH sessions.

13.1.2 The Format of the RMU Analyze Binary_output Files Can Change

Bug 19261529

The /BINARY_OUTPUT qualifier used with the Oracle Rdb RMU Analyze command specifies a binary output *.UNL file and/or a record definition *.RRD file to be created by RMU Analyze to save statistical data created while analysing an Oracle Rdb database and to then load the statistical data into an Oracle Rdb database table. This data can then be used by a user-written management application or procedure. See the *Oracle RMU Reference Manual* for more information on the Binary_output qualifier used with the RMU Analyze command.

The format of the binary and record definition files created by the RMU Analyze Binary_output command can change as changes are made to the statistical data produced by RMU Analyze. Creators and maintainers of user-written management applications or procedures dependent on the format of these files should be aware that Oracle Rdb reserves the right to add, remove or move fields in this record definition file for each major release of Rdb, which can require changes in these user-written management applications or procedures. The release notes for each major version of Oracle Rdb will document any changes in the format of these files. To make users aware of this possibility, the following warning comment has now been added to the record definition file.

```
DEFINE RECORD name
DESCRIPTION IS /* Oracle Rdb V7.3-130
                Oracle reserves the right to add, remove or move fields in
                this record definition file for each major release of Rdb */ .
```

The following example shows a database table named POS being created to hold RMU Analyze Placement statistics. Then an RMU Analyze Placement Binary_output command is executed to save binary placement statistics, for the EMP_EMPLOYEE_ID index in the MF_PERSONNEL database, to the POS.UNL file and the record format for these statistics to the POS.RRD file. The RMU Load command is then executed to load the binary placement statistic records from the POS.UNL file into the POS table using the record field format defined in the POS.RRD file.

```
$ sql
attach 'file mf_personnel';
```



```

create table pos (
    RMU$DATE          DATE VMS,
    RMU$INDEX_NAME    CHAR(32),
    RMU$RELATION_NAME CHAR(32),
    RMU$PARTITION_NAME CHAR(32),
    RMU$AREA_NAME     CHAR(32),
    RMU$LEVEL         INTEGER,
    RMU$FLAGS         INTEGER,
    RMU$COUNT        INTEGER,
    RMU$DUPLICATE_COUNT INTEGER,
    RMU$DUPLICATE_MAP_COUNT INTEGER,
    RMU$KEY_COUNT      INTEGER,
    RMU$DUPLICATE_KEY_COUNT INTEGER,
    RMU$DATA_COUNT     INTEGER,
    RMU$DUPLICATE_DATA_COUNT INTEGER,
    RMU$TOTAL_KEY_PATH INTEGER,
    RMU$TOTAL_PAGE_PATH INTEGER,
    RMU$TOTAL_BUFFER_PATH INTEGER,
    RMU$MAX_KEY_PATH   INTEGER,
    RMU$MAX_PAGE_PATH  INTEGER,
    RMU$MIN_BUF_PATH   INTEGER);
commit;
exit
$ rmu/analyze /placement -
  /binary_output=(file=pos.unl,record_definition=pos.rrd) -
  mf_personnel EMP_EMPLOYEE_ID
$ rmu/load/rms=(file=pos.rrd) mf_personnel POS pos.unl
%RMU-I-DATRECREAD, 1 data records read from input file.
%RMU-I-DATRECSTO, 1 data records stored 3-DEC-2014 10:58:04
$ type pos.rrd

DEFINE RECORD RMU$ANALYZE PLACEMENT
DESCRIPTION IS /* Oracle Rdb V7.3-130
                Oracle reserves the right to add, remove or move fields in
                this record definition file for each major release of Rdb */ .

$

```

13.1.3 RMU Verify Key_values May Fail on Some Indices

In some cases, the RMU Verify Key_values functionality may be unable to perform key value verification failing with an %RDMS-F-OUTLINE_FAILED error. This happens when the generated SQL query and query outline can not be used to ensure index only retrieval from the specified index. This may occur with multi-segment HASHED (ORDERED or SCATTERED) indices, or multi-segment SORTED (and SORTED RANKED) indices with the REVERSE attribute.

The following example shows the reported error.

```

$ rmu/verify/nolog/index=T5_IDENT_IMAGE_NDX /key_values VERIFY_KEY_VALUES_DB
%RDMS-F-OUTLINE_FAILED, could not comply with mandatory query outline directives
%RMU-I-NOTREQVFY, not all requested verifications have been performed
$

```

This problem is a known limitation in this release of Oracle Rdb.

13.1.4 REPLACE Statement Fails With Primary Key Constraint Failure When Used on a View

The REPLACE statement does not show the correct semantics when used to insert into a view defined on a table with a PRIMARY KEY constraint.

The following example shows the problem.

```
SQL> create table SAMPLE
cont>      (ident integer
cont>      ,prod_name char(20)
cont>      ,primary key (ident) not deferrable
cont>      );
SQL>
SQL> create view SAMPLE_VIEW
cont>      (ident, prod_name)
cont>      as select * from SAMPLE
cont> ;
SQL>
SQL> insert into SAMPLE values (1, 'Ajax');
1 row inserted
SQL> insert into SAMPLE_VIEW values (2, 'Mr Clean');
1 row inserted
SQL>
SQL> replace into SAMPLE values (1, 'Borox');
1 row replaced
SQL> replace into SAMPLE_VIEW values (2, 'Mr. Clean');
%RDB-E-INTEG_FAIL, violation of constraint SAMPLE_PRIMARY_IDENT caused operation to fail
-RDB-F-ON_DB, on database RDB$DEFAULT_CONNECTION
SQL>
SQL> select * from SAMPLE order by ident;
      IDENT  PROD_NAME
        1    Borox
        2    Mr Clean
2 rows selected
SQL>
```

This will remain a restriction for this release. Only use REPLACE (or RMU Load Replace) on base tables.

13.1.5 Possible Incorrect Results When Using Partitioned Descending Indexes

Bug 6129797

In the current release of Oracle Rdb, it is possible for some queries using partitioned indexes with segments of mixed ascending and descending order to return incorrect results either on Alpha or IA-64 systems.

The following examples show two problems when using partitioned index with segments of mixed ascending and descending order:

```
create database file foo
  create storage area fooa
  create storage area foob;

create table mesa (id integer, m4 char (1), m5 integer);
create table rasa (id integer, r4 char (1), r5 integer);

insert into mesa (id, m4, m5) values (1, 'm', 1 );
insert into rasa (id, r4, r5) values (1, 'm', 1 );
insert into rasa (id, r4, r5) values (1, 'k', 1 );
insert into rasa (id, r4, r5) values (1, 'e', 1 );

create index x4 on mesa (id asc , m4 asc) ;
```

```

! The following index contains ascending segments followed by descending
! segments and thus causes the query to return the wrong result.
!
! Note that the query works if both segments are either ascending or descending.
!
create index y4 on rasa (id asc , r4 desc)
      store using (id, r4)
      in foaa with limit of (1, 'g' )
      otherwise in foob ;
commit;

! Problem #1:
!
! the following query returns correctly 3 rows on Alpha but 1 row on IA64:
SQL> select m.id, m.m4, r.r4 from
      mesa m inner join rasa r on (m.id = r.id);
          1   m      m
          1   m      k
          1   m      e
3 rows selected

SQL> select m.id, m.m4, r.r4 from mesa m inner join rasa r on (m.id = r.id);
          1   m      e
1 row selected

! Problem #2:
!
! The following query using reverse scan returns 2 rows incorrectly on Alpha
! but 3 rows correctly on IA64:
!

SQL> select id, r4 from rasa where id = 1 and r4 <= 'm' order by id, r4;
Tables:
  0 = RASA
Index only retrieval of relation 0:RASA
  Index name  Y4 [2:1]  Reverse Scan
  Keys: (0.ID = 1) AND (0.R4 <= 'm')
      ID   R4
      1    k
      1    m
2 rows selected

SQL> select id, r4 from rasa where id = 1 and r4 <= 'm' order by id, r4;
Tables:
  0 = RASA
Index only retrieval of relation 0:RASA
  Index name  Y4 [2:1]  Reverse Scan
  Keys: (0.ID = 1) AND (0.R4 <= 'm')
      ID   R4
      1    e
      1    k
      1    m
3 rows selected

```

This problem is related to the construction and comparison of the descending key values while processing the index partitions.

The problem will be corrected in a future version of Oracle Rdb.

13.1.6 Application and Oracle Rdb Both Using SYS\$HIBER

In application processes that use Oracle Rdb and the SYS\$HIBER system service (possibly via RTL routines such as LIB\$WAIT), it is very important that the application ensures that the event being waited for has actually occurred. Oracle Rdb utilizes \$HIBER/\$WAKE sequences for interprocess communication and synchronization.

Because there is just a single process-wide "hibernate" state along with a single process-wide "wake pending" flag, Oracle Rdb must assume that it "shares" use of the hibernate/wake state with the user's application code. To this end, Oracle Rdb generally will re-wake the process via a pending wake request after using a hibernate sequence.

Oracle Rdb's use of the \$WAKE system service will interfere with other users of \$HIBER (such as the routine LIB\$WAIT) that do not check for event completion, possibly causing a \$HIBER to be unexpectedly resumed without waiting at all.

To avoid these situations, applications that use HIBER/WAKE facilities must use a code sequence that avoids continuing without a check for the operation (such as a delay or a timer firing) being complete.

The following pseudo-code shows one example of how a flag can be used to indicate that a timed-wait has completed correctly. The wait does not complete until the timer has actually fired and set TIMER_FLAG to TRUE. This code relies on ASTs being enabled.

```
ROUTINE TIMER_WAIT:
  BEGIN
    ! Clear the timer flag
    TIMER_FLAG = FALSE

    ! Schedule an AST for sometime in the future
    STAT = SYS$SETIMR (TIMADR = DELTATIME, ASTRTN = TIMER_AST)
    IF STAT <> SS$_NORMAL THEN LIB$SIGNAL (STAT)

    ! Hibernate. When the $HIBER completes, check to make
    ! sure that TIMER_FLAG is set indicating that the wait
    ! has finished.
    WHILE TIMER_FLAG = FALSE
      DO SYS$HIBER()
    END

ROUTINE TIMER_AST:
  BEGIN
    ! Set the flag indicating that the timer has expired
    TIMER_FLAG = TRUE

    ! Wake the main-line code
    STAT = SYS$WAKE ()
    IF STAT <> SS$_NORMAL THEN LIB$SIGNAL (STAT)
  END
```

Starting with OpenVMS V7.1, the LIB\$WAIT routine includes a FLAGS argument (with the LIB\$K_NOWAKE flag set) to allow an alternate wait scheme (using the \$SYNCH system service) that can avoid potential problems with multiple code sequences using the \$HIBER system service. See the OpenVMS RTL Library (LIB\$) Manual for more information about the LIB\$WAIT routine.

In order to prevent application hangs, inner-mode users of SYS\$HIBER must take explicit steps to ensure that a pending wake is not errantly "consumed". The general way of accomplishing this is to issue a SYS\$WAKE to the process after the event is complete if a call to SYS\$HIBER was done. Rdb takes this step

and therefore application programs must be prepared for cases where a wakeup might appear unexpectedly.

13.1.7 Unexpected RCS Termination

It has been observed in internal testing of Rdb Release 7.2.2 that if the Record Cache Server (the RCS) terminates in an uncontrolled fashion this may, under some conditions, cause corruption of the database and/or the After Image Journal file.

When the RCS terminates, the database is shut down and a message like the following is written to the monitor log:

```
6-DEC-2007 15:04:17.02 - Received Record Cache Server image termination from
22ED5144:1
- database name "device:[directory]database.RDB;1" [device] (1200,487,0)
- abnormal Record Cache Server termination detected
- starting delete-process shutdown of database:
  - %RDMS-F-RCSABORTED, record cache server process terminated abnormally
- sending process deletion to process 22ED10F9
- sending process deletion to process 22ECED59
- sending process deletion to process 22EC0158
- sending process deletion to process 22EB9543 (AIJ Log server)
- database shutdown waiting for active users to terminate
```

A future attempt to roll forward the AIJ following a restore of a database backup might fail with a bugcheck dump if this problem has happened.

The only currently known situation where this problem has been observed is if the logical name RDM\$BIND_RCS_VALIDATE_SECS is defined to some value and the logical name RDM\$BIND_RCS_LOG_FILE at the same time is undefined or defined incorrectly.

To prevent this problem, Oracle recommends any customer using the Row Cache feature to either avoid defining the logical name RDM\$BIND_RCS_VALIDATE_SECS or if this logical name for any reason needs to be defined, to make sure RDM\$BIND_RCS_LOG_FILE is correctly defined (i.e. defined with the /SYSTEM and /EXECUTIVE qualifiers and is pointing to a valid file name in an existing directory on a cluster accessible device with sufficient free space).

This recommendation applies to all versions of Oracle Rdb.

13.1.8 External Routine Images Linked with PTHREAD\$RTL

The OpenVMS Guide to the POSIX Threads Library describes that it is not supported to dynamically activate the core run-time library shareable image PTHREAD\$RTL. Oracle has found in testing that a shareable image supplied for use as an External Routine that is linked with PTHREAD\$RTL can be expected to cause a hang during dynamic image activation on OpenVMS IA-64 systems. This problem has not been observed on OpenVMS Alpha systems.

To avoid this problem, in any case where the shareable image used for an Rdb External Routine is linked with PTHREAD\$RTL, the main program image must likewise be linked with PTHREAD\$RTL. This requirement applies to customer built application main programs as well as the main interactive SQL image.

The shareable image RDB\$NATCONN_FUNC73.EXE supplied with OCI Services for Oracle Rdb (part of SQL/Services) is one such shareable image that is linked with PTHREAD\$RTL. Customer built applications that utilize External Routines from the RDB\$NATCONN_FUNC73.EXE image must ensure that the main image is linked with PTHREAD\$RTL. The external routines that a user may call that use functions from RDB\$NATCONN_FUNC73.EXE include:

- TO_CHAR
- TO_NUMBER
- TO_DATE

You can use the OpenVMS command ANALYZE/IMAGE to determine whether an image depends upon PTHREAD\$RTL. For more information, see the OpenVMS documentation.

13.1.9 ILINK-E-INVORINI Error on IA-64

When linking an application with multiple modules, the following error message may be returned:

```
%ILINK-E-INVORINI, incompatible multiple initializations for overlaid section
      section: VMSRDB
      module: M1
      file: DKA0:[BLD]M1.OBJ;1
      module: M2
      file: DKA0:[BLD]SYS.OLB;1
```

On IA-64 systems, it is not allowed to have a program section that attempts to be initialized a subsequent time where the non-zero portions of the initializations do not match. This is a difference from OpenVMS Alpha and VAX systems where the linker permitted such initializations.

If the modules specified are SQL module language or precompiler produced, the application build procedures usually need to be modified. Typically, the solution is to initialize the database handles in only one of the modules. The SQLMOD command line qualifiers /NOINITIALIZE_HANDLES and /INITIALIZE_HANDLES are used to specify whether or not alias definitions are coerced into alias references.

13.1.10 SYSTEM-F-INSFMEM Fatal Error With SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED in Galaxy Environment

When using the GALAXY SUPPORT IS ENABLED feature in an OpenVMS Galaxy environment, a %SYSTEM-F-INSFMEM, insufficient dynamic memory error may be returned when mapping record caches or opening the database. One source of this problem specific to a Galaxy configuration is running out of Galaxy Shared Memory regions. For Galaxy systems, GLX_SHM_REG is the number of shared memory region structures configured into the Galaxy Management Database (GMDB).

While the default value (for OpenVMS versions through at least V7.3-1) of 64 regions might be adequate for some installations, sites using a larger number of databases or row caches, when the SHARED MEMORY IS SYSTEM or LARGE MEMORY IS ENABLED features are enabled, may find the default insufficient.

If a %SYSTEM-F-INSFMEM, insufficient dynamic memory error is returned when mapping record caches or opening databases, Oracle Corporation recommends that you increase the GLX_SHM_REG parameter by 2 times the sum of the number of row caches and number of databases that might be accessed in the Galaxy at one time. As the Galaxy shared memory region structures are not very large, setting this parameter to a higher than required value does not consume a significant amount of physical memory. It also may avoid a later reboot of the Galaxy environment. This parameter must be set on all nodes in the Galaxy.

Changing the GLX_SHM_REG system parameter requires that the OpenVMS Galaxy environment be booted from scratch. That is, all nodes in the Galaxy must be shut down and then the Galaxy reformed by starting each instance.

13.1.11 Oracle Rdb and OpenVMS ODS-5 Volumes

OpenVMS Version 7.2 introduced an Extended File Specifications feature, which consists of two major components:

- A new, optional, volume structure, ODS-5, which provides support for file names that are longer and have a greater range of legal characters than in previous versions of OpenVMS.
- Support for “deep” directory trees.

ODS-5 was introduced primarily to provide enhanced file sharing capabilities for users of Advanced Server for OpenVMS 7.2 (formerly known as PATHWORKS for OpenVMS), as well as DCOM and JAVA applications.

In some cases, Oracle Rdb performs its own file and directory name parsing and explicitly requires ODS-2 (the traditional OpenVMS volume structure) file and directory name conventions to be followed. Because of this knowledge, Oracle does not support any Oracle Rdb database file components (including root files, storage area files, after image journal files, record cache backing store files, database backup files, after image journal backup files, etc.) that utilize any non-ODS-2 file naming features. For this reason, Oracle recommends that Oracle Rdb database components not be located on ODS-5 volumes.

Oracle does support Oracle Rdb database file components on ODS-5 volumes provided that all of these files and directories used by Oracle Rdb strictly follow the ODS-2 file and directory name conventions. In particular, all file names must be specified entirely in uppercase and “special” characters in file or directory names are forbidden.

13.1.12 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, a user has two tables T1 and T2, both with one column, and wishes to ensure that all values in table T1 exist in T2. Both tables have an index on the referenced field. The user could use a PRIMARY KEY constraint on T2 and a FOREIGN KEY constraint on T1.

```
SQL> alter table t2 alter column f2 primary key not deferrable;  
SQL> alter table t1 alter column f1 references t2 not deferrable;
```

When deleting from the PRIMARY KEY table, Rdb will only check for rows in the FOREIGN KEY table where the FOREIGN KEY has the deleted value. This can be seen as an index lookup on T1 in the retrieval strategy.


```

SQL> delete from t2 where f2=1;
Get      Temporary relation      Retrieval by index of relation T2
      Index name  I2 [1:1]
Index only retrieval of relation T1
      Index name  I1 [1:1]
%RDB-E-INTEG_FAIL, violation of constraint T1_FOREIGN1 caused operation to fail

```

The failure of the constraint is not important. What is important is that Rdb efficiently detects that only those rows in T1 with the same values as the deleted row in T2 can be affected.

It is necessary sometimes to define this type of relationship using CHECK constraints. This could be necessary because the presence of NULL values in the table T2 precludes the definition of a primary key on that table. This could be done with a CHECK constraint of the form:

```

SQL> alter table t1 alter column f1
cont>  check (f1 in (select * from t2)) not deferrable;
SQL> delete from t2 where f2=1;
Get      Temporary relation      Retrieval by index of relation T2
      Index name  I2 [1:1]
Cross block of 2 entries
  Cross block entry 1
    Index only retrieval of relation T1
      Index name  I1 [0:0]
  Cross block entry 2
    Conjunct      Aggregate-F1      Conjunct
    Index only retrieval of relation T2
      Index name  I2 [0:0]
%RDB-E-INTEG_FAIL, violation of constraint T1_CHECK1 caused operation to fail

```

The cross block is for the constraint evaluation. This retrieval strategy indicates that to evaluate the constraint, the entire index on table T1 is being scanned and for each key, the entire index in table T2 is being scanned. The behavior can be improved somewhat by using an equality join condition in the select clause of the constraint:

```

SQL> alter table t1 alter column f1
cont>  check (f1 in (select * from t2 where f2=f1)) not deferrable;
or:

```

```

SQL> alter table t1 alter column f1
cont>  check (f1=(select * from t2 where f2=f1)) not deferrable;

```

In both cases, the retrieval strategy will look like this:

```

SQL> delete from t2 where f2=1;
Get      Temporary relation      Retrieval by index of relation T2
      Index name  I2 [1:1]
Cross block of 2 entries
  Cross block entry 1
    Index only retrieval of relation T1
      Index name  I1 [0:0]
  Cross block entry 2
    Conjunct      Aggregate-F1      Conjunct
    Index only retrieval of relation T2
      Index name  I2 [1:1]
%RDB-E-INTEG_FAIL, violation of constraint T1_CHECK1 caused operation to fail

```

While the entire T1 index is scanned, at least the value from T1 is used to perform an index lookup on T2.

These restrictions result from semantic differences in the behavior of the "IN" and "EXISTS" operators with respect to null handling, and the complexity of dealing with non-equality join conditions.

To improve the performance of this type of integrity check on larger tables, it is possible to use a series of triggers to perform the constraint check. The following triggers perform a similar check to the constraints above.

```
SQL> create trigger t1_insert after insert on t1
cont> when (not exists (select * from t2 where f2=f1))
cont> (error) for each row;
SQL> create trigger t1_update after update on t1
cont> when (not exists (select * from t2 where f2=f1))
cont> (error) for each row;
SQL> ! A delete trigger is not needed on T1.
SQL> create trigger t2_delete before delete on t2
cont> when (exists (select * from t1 where f1=f2))
cont> (error) for each row;
SQL> create trigger t2_modify after update on t2
cont> referencing old as t2o new as t2n
cont> when (exists (select * from t1 where f1=t2o.f2))
cont> (error) for each row;
SQL> ! An insert trigger is not needed on T2.
```

The strategy for a delete on T2 is now:

```
SQL> delete from t2 where f2=1;
Aggregate-F1      Index only retrieval of relation T1
      Index name  I1 [1:1]
Temporary relation      Get      Retrieval by index of relation T2
      Index name  I2 [1:1]
%RDB-E-TRIG_INV_UPD, invalid update; encountered error condition defined for
trigger
-RDMS-E-TRIG_ERROR, trigger T2_DELETE forced an error
```

The trigger strategy is the index only retrieval displayed first. You will note that the index on T1 is used to examine only those rows that may be affected by the delete.

Care must be taken when using this workaround as there are semantic differences in the operation of the triggers, the use of "IN" and "EXISTS", and the use of referential integrity constraints.

This workaround is useful where the form of the constraint is more complex, and cannot be phrased using referential integrity constraints. For example, if the application is such that the value in table T1 may be spaces or NULL to indicate the absence of a value, the above triggers could easily be modified to allow for these semantics.

13.1.13 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.

13.1.14 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.

13.1.15 Row Cache Not Allowed While Hot Standby Replication is Active

The row cache feature may not be enabled on a hot standby database while replication is active. The hot standby feature will not start if row cache is enabled.

This restriction exists because rows in the row cache are accessed via logical dbkeys. However, information transferred to the standby database via the after image journal facility only contains physical dbkeys. Because there is no way to maintain rows in the cache via the hot standby processing, the row cache must be disabled when the standby database is open and replication is active.

A new command qualifier, ROW_CACHE=DISABLED, has been added to the RMU Open command. To open the hot standby database prior to starting replication, use the ROW_CACHE=DISABLED qualifier on the RMU Open command.

13.1.16 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 13–1 Sort Memory Logicals

Logical	Definition
RDMSS\$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 very 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.
RDMSS\$BIND_SORT_MEMORY_MAX_BYTES	Specifies an absolute limit to be used when allocating sort memory for the built-in SORT32 package. If not defined, the default value is unlimited (up to 1GB), the maximum value is 2147483647 and the minimum value is 32768.

13.1.17 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 cursor's subject table will be updated. For example, if the SQL syntax UPDATE ... WHERE CURRENT OF is used to perform updates of target rows, or the RDO/RDML MODIFY statement uses the context variable for the stream. Then the optimizer will choose an alternate access strategy if an update can occur which may cause the Halloween problem. This can be seen in the access strategy in Example 2-2 as a "Temporary relation" being created to hold the result of the cursor query.

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

The following example shows that the EMP_LAST_NAME index is used for retrieval. Any update performed will possibly be subject to the Halloween problem.

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

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

```
SQL> set flags 'strategy';
SQL> declare emp2 cursor for
cont> select * from employees where last_name >= 'M'
cont> order by last_name for update of last_name;
SQL> open emp2;
Temporary relation      Conjunct      Get
Retrieval by index of relation 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 cursor's 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.

13.1.18 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.

13.1.19 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.

13.1.20 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.

13.1.21 RMU Convert Fails When Maximum Relation ID is Exceeded

If, when relation IDs are assigned to new system tables during an RMU Convert to a V7.2 database, the maximum relation ID of 8192 allowed by Oracle Rdb is exceeded, the fatal error %RMU-F-RELMAXIDBAD is displayed and the database is rolled back to the prior database version. Contact your Oracle support representative if you get this error. Note that when the database is rolled back, the fatal error %RMU-F-CVTRLSUC 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-CVTRLSUC error message when the database has been rolled back to V7.0 since it cannot be converted to V7.2:

```
$rmu/convert mf_personnel
%RMU-I-RMUTXT 000, Executing RMU for Oracle Rdb V7.2
Are you satisfied with your backup of
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 and your backup of
  any associated .aij files [N]? Y
%RMU-I-LOGCONVRT, database root converted to current structure level
%RMU-F-RELMAXIDBAD, ROLLING BACK CONVERSION - Relation ID exceeds maximum
  8192 for system table RDB$RELATIONS
%RMU-F-CVTRLSUC, CONVERT rolled-back for
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 to version V7.0
```

The following example shows the normal case when the maximum allowed relation ID is not exceeded:

```
$rmu/convert mf_personnel
%RMU-I-RMUTXT 000, Executing RMU for Oracle Rdb V7.2
Are you satisfied with your backup of
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 and your backup of
  any associated .aij files [N]? Y
%RMU-I-LOGCONVRT, database root converted to current structure level
%RMU-S-CVTDBSUC, database DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1
  successfully converted from version V7.0 to V7.2
%RMU-I-CVTCOMSUC, CONVERT committed for
  DEVICE:[DIRECTORY]MF_PERSONNEL.RDB;1 to version V7.2
```

13.1.22 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=optionsfile qualifier option file can be used with the TYPE qualifier. For example, to repair the EMPLOYEES table of the MF_PERSONNEL database, you would create an options file that contains the following line:

```
EMPLOYEES /TYPE=TABLE
```

For partitioned logical areas, the AREA=name qualifier can be used to identify the specific storage areas that are to be updated. For example, to repair the EMPLOYEES table of the MF_PERSONNEL database for the EMPID_OVER storage area only, you would create an options file that contains the following line:

```
EMPLOYEES /AREA=EMPID_OVER /TYPE=TABLE
```

The TYPE qualifier specifies the type of a logical area. The following keywords are allowed:

- **TABLE**
Specifies that the logical area is a data table. This would be a table created using the SQL CREATE TABLE syntax.
- **B-TREE**
Specifies that the logical area is a B-tree index. This would be an index created using the SQL CREATE INDEX TYPE IS SORTED syntax.
- **HASH**
Specifies that the logical area is a hash index. This would be an index created using the SQL CREATE INDEX TYPE IS HASHED syntax.
- **SYSTEM**
Specifies that the logical area is a system record that is used to identify hash buckets. Users cannot explicitly create these types of logical areas.

Note

This type should NOT be used for the RDB\$SYSTEM logical areas. This type does NOT identify system relations.

- **BLOB**
Specifies that the logical area is a BLOB repository.

There is no explicit error checking of the type specified for a logical area. However, an incorrect type may cause the RMU Unload After_Journal command to be unable to correctly return valid, logical dbkeys.

13.1.23 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.

13.1.24 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*, *Oracle Rdb7 Guide to Database Design and Definition*, and the *Oracle RMU Reference Manual* for additional information about Oracle Rdb backup and restore operations.

13.1.25 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.

13.1.26 Converting Single-File Databases

Because of a substantial increase in the database root file information for V7.0, you should ensure that you have adequate disk space before you use the RMU Convert command with single-file databases and V7.0 or higher.

The size of the database root file of any given database increases a maximum of about 600 disk blocks. The actual increase depends mostly on the maximum number of users specified for the database.

13.1.27 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.

13.1.28 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.

13.1.29 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 illustrates the behavior:

```
SQL> create table T1
cont>      (id integer
cont>        ,last_name char(12)
cont>        ,first_name char(12)
cont>        );
SQL> create storage map M for T1
cont>      partitioning not updatable
cont>      store using (id)
cont>          in EMPIDS_LOW with limit of (200)
cont>          in EMPIDS_MID with limit of (400)
cont>          otherwise in EMPIDS_OVER;
SQL> insert into T1 values (150,'Boney','MaryJean');
1 row inserted
SQL> insert into T1 values (350,'Morley','Steven');
1 row inserted
SQL> insert into T1 values (300,'Martinez','Nancy');
1 row inserted
SQL> insert into T1 values (450,'Gentile','Russ');
1 row inserted
SQL>
SQL> set flags 'EXECUTION(100),STRATEGY,DETAIL(2),INDEX_PARTITIONS';
SQL>
SQL> select * from T1 where ID > 400;
~S#0001
Tables:
  0 = T1
Conjunct: 0.ID > 400
Get      Retrieval sequentially of relation 0:T1          (partitioned scan#1)
~E#0001.1: Strict Partitioning using 2 areas
```

```

partition 2 (larea=60) "EMPIDS_MID"
partition 3 (larea=61) "EMPIDS_OVER" otherwise
      ID  LAST_NAME  FIRST_NAME
      450  Gentile    Russ
1 row selected
SQL>

```

In this example, partition 2 does not need to be scanned but is still accessed due to the structure of the generated key values. This does not affect the correctness of the result. Users can avoid the extra scan by using values other than the boundary values.

13.1.30 Restriction When Adding Storage Areas with Users Attached to Database

If you try to interactively add a new storage area where the page size is less than the smallest existing page size and the database has been manually opened or users are active, the add operation fails with the following errors:

```
%RDMS-F-NOEUACCESS, unable to acquire exclusive access to database
```

or

```
%RDB-F-SYS REQUEST, error from system services request
-RDMS-F-FILACCERR, error opening database root DKA0:[RDB]TEST.RDB;1
-SYSTEM-W-ACCONFLICT, file access conflict
```

You can make this change only when no users are attached to the database and, if the database is set to OPEN IS MANUAL, the database is closed. Several internal Oracle Rdb data structures are based on the minimum page size and these structures cannot be resized if users are attached to the database.

Furthermore, because this particular change is not recorded in the AIJ, any recovery scenario fails. Note also that if you use .aij files, you must backup the database and restart after-image journaling because this change invalidates the current AIJ recovery.

13.1.31 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.

13.1.32 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.

13.1.33 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 RDMSS\$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.

13.1.34 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 13–2 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

13.1.35 Side Effect When Calling Stored Routines

When calling a stored routine, you must not use the same routine to calculate argument values by a stored function. For example, if the routine being called is also called by a stored function during the calculation of an argument value, passed arguments to the routine may be incorrect.

The following example shows a stored procedure P being called during the calculation of the arguments for another invocation of the stored procedure P:


```

SQL> create module M
cont>     language SQL
cont>
cont>     procedure P (in :a integer, in :b integer, out :c integer);
cont>     begin
cont>         set :c = :a + :b;
cont>     end;
cont>
cont>     function F () returns integer
cont>         comment is 'expect F to always return 2';
cont>     begin
cont>         declare :b integer;
cont>         call P (1, 1, :b);
cont>         trace 'returning ', :b;
cont>         return :b;
cont>     end;
cont> end module;
SQL>
SQL> set flags 'TRACE';
SQL> begin
cont> declare :cc integer;
cont> call P (2, F(), :cc);
cont> trace 'Expected 4, got ', :cc;
cont> end;
~Xt: returning 2
~Xt: Expected 4, got 3

```

The result as shown above is incorrect. The routine argument values are written to the called routine's parameter area before complex expression values are calculated. These calculations may (as in the example) overwrite previously copied data.

The workaround is to assign the argument expression (in this example calling the stored function F) to a temporary variable and pass this variable as the input for the routine. The following example shows the workaround:

```

SQL> begin
cont> declare :bb, :cc integer;
cont> set :bb = F();
cont> call P (2, :bb, :cc);
cont> trace 'Expected 4, got ', :cc;
cont> end;
~Xt: returning 2
~Xt: Expected 4, got 4

```

This problem will be corrected in a future version of Oracle Rdb.

13.1.36 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.