

Volume

4

ORACLE CORPORATION

**RMU SHOW STATISTIC
DBA HANDBOOK**

(Second Edition)

SCREEN INFORMATION

**RICK ANDERSON
ORACLE RDB ENGINEERING
REV. 12/15/98**

Oracle is a registered trademark of Oracle Corporation, Redwood City, California.

Oracle Rdb, Oracle RMU and Rdb7 are trademarks of Oracle Corporation, Redwood City, California.

Open VMS is a trademark of Digital Equipment Corporation.

All other company or product names are used for identification purposes only and may be trademarks of their respective owners.

This software/documentation contains proprietary information of Oracle Corporation; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited. If this software is delivered to a U.S. Government Agency of the Department of Defense, then it is delivered with Restricted Rights and the following legend is applicable:

Restricted Rights Legend

Use, duplication or disclosure by the Government is subject to restrictions for commercial computer software and shall be deemed to be Restricted Rights software under Federal law, and as set forth in subparagraph (c) (1) (ii) of DFARS 252.227-7013, Rights in Technical Data and Computer Software (October 1988).

This material or any portion of it may not be copied in any form or by any means without the express prior written permission of Oracle Education Group of Oracle Corporation. Any other copying is a violation of copyright law and may result in civil and/or criminal penalties.

If this software/documentation is delivered to a U.S. Government Agency not within the Department of Defense, then it is delivered with "Restricted Rights," as defined in FAR 52.227-14, Rights in Data - General, including Alternate III (June 1987).

The information in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error free.

Oracle7, Oracle8, Oracle Parallel Server, Oracle Application Display Manager, Oracle Applications, Oracle Applications Window Manager, Oracle Engineering, Oracle Forms, Oracle Reports, Oracle Browser, Oracle Data Query, Oracle Document Manager, DDE Manager, Oracle Business Manager, Oracle Human Resources, Oracle Office Administrator, Oracle Office Scheduler, Oracle Office Directory, Oracle Assets, Oracle Payables, Oracle Project Accounting, Oracle Purchasing, Oracle Personnel, Oracle Open Gateway, Oracle Order Entry, Oracle Payroll, Oracle Receivables, Oracle Sales Analysis, Oracle Imaging, Oracle Inventory, Oracle*Mail, Oracle Office Mail, Oracle Manufacturing, Oracle Master Scheduling, Oracle MRP, Oracle SQL*Tutor, Oracle*Terminal, Oracle Toolkit, Oracle Work In Progress, SQL*TextRetrieval, PL/SQL, Pro*C, Report Card, Oracle Graphics, The Oracle Network Builder, Oracle Bill of Materials, Oracle BookBatch, Oracle BookBuilder, Oracle CASE, Hyper*SQL, and CASE*Generator are trademarks of Oracle Corporation.

Table of Contents

PREFACE	1
PURPOSE OF THIS HANDBOOK	1
INTENDED AUDIENCE	1
HANDBOOK CONTENTS	1
ORACLE CODASYL DBMS	2
OPERATING SYSTEM INFORMATION.....	2
RELATED MANUALS	2
ACKNOWLEDG-MENTS	2
AUTHOR.....	3
RMU SHOW STATISTIC ROW CACHE (ONE CACHE) SCREENS	4
ROW CACHE SCREEN	5
SCREEN LOCATION	5
SCREEN EXAMPLE.....	5
SCREEN FIELDS	5
ON-SCREEN MENU OPTIONS.....	7
SCREEN INFORMATION.....	8
CONFIGURATION OPTIONS.....	8
RMU SHOW STATISTIC ROW CACHE (ONE FIELD) SCREENS	9
ROW CACHE SCREEN	10
SCREEN LOCATION	10
SCREEN EXAMPLE.....	10
ON-SCREEN MENU OPTIONS.....	11
SCREEN INFORMATION.....	11
CONFIGURATION OPTIONS.....	11
RMU SHOW STATISTIC ROW CACHE INFORMATION SCREENS	12
ROW CACHE UTILIZATION SCREEN	13
SCREEN LOCATION	13
SCREEN EXAMPLE.....	13
SCREEN FIELDS	14
ON-SCREEN MENU OPTIONS.....	14
SCREEN INFORMATION.....	14
CONFIGURATION OPTIONS.....	15
HOT ROW INFORMATION SCREEN	16

SCREEN LOCATION	16
SCREEN EXAMPLE.....	16
SCREEN FIELDS	17
ON-SCREEN MENU OPTIONS	18
SCREEN INFORMATION.....	18
CONFIGURATION OPTIONS.....	19
ROW CACHE STATUS SCREEN.....	20
SCREEN LOCATION	20
SCREEN EXAMPLE.....	20
SCREEN FIELDS	20
ON-SCREEN MENU OPTIONS	22
SCREEN INFORMATION.....	22
CONFIGURATION OPTIONS.....	23
ROW CACHE QUEUE LENGTH SCREEN.....	24
SCREEN LOCATION	24
SCREEN EXAMPLE.....	24
SCREEN FIELDS	25
ON-SCREEN MENU OPTIONS	25
SCREEN INFORMATION.....	25
CONFIGURATION OPTIONS.....	25
ROW LENGTH DISTRIBUTION SCREEN.....	27
SCREEN LOCATION	27
SCREEN EXAMPLE.....	27
SCREEN FIELDS	28
ON-SCREEN MENU OPTIONS	28
SCREEN INFORMATION.....	29
CONFIGURATION OPTIONS.....	29
RCS STATISTICS SCREEN.....	30
SCREEN LOCATION	30
SCREEN EXAMPLE.....	30
SCREEN FIELDS	30
ON-SCREEN MENU OPTIONS	30
SCREEN INFORMATION.....	31
CONFIGURATION OPTIONS.....	31
RMU SHOW STATISTIC INDEX INFORMATION SCREENS	32
INDEX STATISTICS (RETRIEVAL) SCREEN	33
SCREEN LOCATION	33
SCREEN EXAMPLE.....	33
SCREEN FIELDS	33
ON-SCREEN MENU OPTIONS	35
SCREEN INFORMATION.....	35
CONFIGURATION OPTIONS.....	36
INDEX STATISTICS (INSERTION) SCREEN.....	37
SCREEN LOCATION	37

SCREEN EXAMPLE.....	37
SCREEN FIELDS	37
ON-SCREEN MENU OPTIONS	39
SCREEN INFORMATION.....	39
CONFIGURATION OPTIONS.....	40
INDEX STATISTICS (REMOVAL) SCREEN.....	41
SCREEN LOCATION	41
SCREEN EXAMPLE.....	41
SCREEN FIELDS	41
ON-SCREEN MENU OPTIONS	42
SCREEN INFORMATION.....	43
CONFIGURATION OPTIONS.....	43
HASH INDEX STATISTICS SCREEN	44
SCREEN LOCATION	44
SCREEN EXAMPLE.....	44
SCREEN FIELDS	44
ON-SCREEN MENU OPTIONS	45
SCREEN INFORMATION.....	46
CONFIGURATION OPTIONS.....	46
INDEX STATISTICS SCREEN.....	47
SCREEN LOCATION	47
SCREEN EXAMPLE.....	47
SCREEN FIELDS	47
ON-SCREEN MENU OPTIONS	49
SCREEN INFORMATION.....	50
CONFIGURATION OPTIONS.....	50
RMU SHOW STATISTIC GENERAL INFORMATION SCREENS	51
NAME TRANSLATION SCREEN	52
SCREEN LOCATION	52
SCREEN EXAMPLE.....	52
SCREEN FIELDS	52
ON-SCREEN MENU OPTIONS	52
SCREEN INFORMATION.....	53
CONFIGURATION OPTIONS.....	53
RMU SHOW STATISTIC OBJECTS (ONE STAT TYPE) SCREENS	54
KROOT OBJECT SCREEN.....	55
SCREEN LOCATION	55
SCREEN EXAMPLE.....	55
SCREEN FIELDS	55
ON-SCREEN MENU OPTIONS	56
SCREEN INFORMATION.....	56
CONFIGURATION OPTIONS.....	57
FILID OBJECT SCREEN.....	58

SCREEN LOCATION	58
SCREEN EXAMPLE.....	58
SCREEN FIELDS	58
ON-SCREEN MENU OPTIONS	59
SCREEN INFORMATION.....	59
CONFIGURATION OPTIONS.....	60
SEQBLK OBJECT SCREEN	61
SCREEN LOCATION	61
SCREEN EXAMPLE.....	61
SCREEN FIELDS	61
ON-SCREEN MENU OPTIONS	62
SCREEN INFORMATION.....	62
CONFIGURATION OPTIONS.....	63
TSNBLK OBJECT SCREEN	64
SCREEN LOCATION	64
SCREEN EXAMPLE.....	64
SCREEN FIELDS	64
ON-SCREEN MENU OPTIONS	65
SCREEN INFORMATION.....	65
CONFIGURATION OPTIONS.....	66
AIJDB OBJECT SCREEN.....	67
SCREEN LOCATION	67
SCREEN EXAMPLE.....	67
SCREEN FIELDS	67
ON-SCREEN MENU OPTIONS	68
SCREEN INFORMATION.....	68
CONFIGURATION OPTIONS.....	69
AIJFB OBJECT SCREEN	70
SCREEN LOCATION	70
SCREEN EXAMPLE.....	70
SCREEN FIELDS	70
ON-SCREEN MENU OPTIONS	71
SCREEN INFORMATION.....	71
CONFIGURATION OPTIONS.....	72
RTUPB OBJECT SCREEN	73
SCREEN LOCATION	73
SCREEN EXAMPLE.....	73
SCREEN FIELDS	73
ON-SCREEN MENU OPTIONS	74
SCREEN INFORMATION.....	74
CONFIGURATION OPTIONS.....	75
ACTIVE OBJECT SCREEN	76
SCREEN LOCATION	76
SCREEN EXAMPLE.....	76
SCREEN FIELDS	76

ON-SCREEN MENU OPTIONS	77
SCREEN INFORMATION.....	77
CONFIGURATION OPTIONS.....	78
CPT OBJECT SCREEN.....	79
SCREEN LOCATION	79
SCREEN EXAMPLE.....	79
SCREEN FIELDS	79
ON-SCREEN MENU OPTIONS	80
SCREEN INFORMATION.....	80
CONFIGURATION OPTIONS.....	81
RCACHE OBJECT SCREEN	82
SCREEN LOCATION	82
SCREEN EXAMPLE.....	82
SCREEN FIELDS	82
ON-SCREEN MENU OPTIONS	83
SCREEN INFORMATION.....	83
CONFIGURATION OPTIONS.....	84
CLIENT OBJECT SCREEN.....	85
SCREEN LOCATION	85
SCREEN EXAMPLE.....	85
SCREEN FIELDS	85
ON-SCREEN MENU OPTIONS	86
SCREEN INFORMATION.....	86
CONFIGURATION OPTIONS.....	87
CLTSEQ OBJECT SCREEN	88
SCREEN LOCATION	88
SCREEN EXAMPLE.....	88
SCREEN FIELDS	88
ON-SCREEN MENU OPTIONS	89
SCREEN INFORMATION.....	89
CONFIGURATION OPTIONS.....	90
UTILITY OBJECT SCREEN.....	91
SCREEN LOCATION	91
SCREEN EXAMPLE.....	91
SCREEN FIELDS	91
ON-SCREEN MENU OPTIONS	92
SCREEN INFORMATION.....	92
CONFIGURATION OPTIONS.....	93
RMU SHOW STATISTIC OBJECTS (ONE STAT FIELD) SCREENS.....	94
OBJECTS FETCH SHRD SCREEN.....	95
SCREEN LOCATION	95
SCREEN EXAMPLE.....	95
SCREEN FIELDS	95
ON-SCREEN MENU OPTIONS	96

SCREEN INFORMATION.....	97
CONFIGURATION OPTIONS.....	97
OBJECTS FETCH EXCL SCREEN.....	98
SCREEN LOCATION	98
SCREEN EXAMPLE.....	98
SCREEN FIELDS	98
ON-SCREEN MENU OPTIONS	99
SCREEN INFORMATION.....	100
CONFIGURATION OPTIONS.....	100
OBJECTS REFRESHED SCREEN	101
SCREEN LOCATION	101
SCREEN EXAMPLE.....	101
SCREEN FIELDS	101
ON-SCREEN MENU OPTIONS	102
SCREEN INFORMATION.....	103
CONFIGURATION OPTIONS.....	103
OBJECTS MODIFIED SCREEN.....	104
SCREEN LOCATION	104
SCREEN EXAMPLE.....	104
SCREEN FIELDS	104
ON-SCREEN MENU OPTIONS	105
SCREEN INFORMATION.....	106
CONFIGURATION OPTIONS.....	106
OBJECTS WRITTEN SCREEN	107
SCREEN LOCATION	107
SCREEN EXAMPLE.....	107
SCREEN FIELDS	107
ON-SCREEN MENU OPTIONS	108
SCREEN INFORMATION.....	109
CONFIGURATION OPTIONS.....	109
OBJECTS RELEASED SCREEN.....	110
SCREEN LOCATION	110
SCREEN EXAMPLE.....	110
SCREEN FIELDS	110
ON-SCREEN MENU OPTIONS	111
SCREEN INFORMATION.....	112
CONFIGURATION OPTIONS.....	112
RMU SHOW STATISTIC DATABASE DASHBOARD SCREENS	113
IO DASHBOARD SCREEN.....	114
SCREEN LOCATION	114
SCREEN EXAMPLE.....	115
SCREEN FIELDS	115
ON-SCREEN MENU OPTIONS	116
SCREEN INFORMATION.....	116

CONFIGURATION OPTIONS.....	117
LOCKING DASHBOARD SCREEN.....	118
SCREEN LOCATION	118
SCREEN EXAMPLE.....	119
SCREEN FIELDS	119
ON-SCREEN MENU OPTIONS	120
SCREEN INFORMATION.....	120
CONFIGURATION OPTIONS.....	120
AIJ DASHBOARD SCREEN.....	122
SCREEN LOCATION	122
SCREEN EXAMPLE.....	123
SCREEN FIELDS	123
ON-SCREEN MENU OPTIONS	124
SCREEN INFORMATION.....	124
CONFIGURATION OPTIONS.....	125
CHECKPOINT DASHBOARD SCREEN.....	126
SCREEN LOCATION	126
SCREEN EXAMPLE.....	127
SCREEN FIELDS	127
ON-SCREEN MENU OPTIONS	127
SCREEN INFORMATION.....	128
CONFIGURATION OPTIONS.....	128
HOT STANDBY DASHBOARD SCREEN.....	130
SCREEN LOCATION	130
SCREEN EXAMPLE.....	131
SCREEN FIELDS	131
ON-SCREEN MENU OPTIONS	133
SCREEN INFORMATION.....	133
CONFIGURATION OPTIONS.....	133
ROW CACHE DASHBOARD SCREEN.....	135
SCREEN LOCATION	135
SCREEN EXAMPLE.....	136
SCREEN FIELDS	136
ON-SCREEN MENU OPTIONS	136
SCREEN INFORMATION.....	137
CONFIGURATION OPTIONS.....	137
RUJ DASHBOARD SCREEN.....	139
SCREEN LOCATION	139
SCREEN EXAMPLE.....	140
SCREEN FIELDS	140
ON-SCREEN MENU OPTIONS	140
SCREEN INFORMATION.....	140
CONFIGURATION OPTIONS.....	141
MONITOR DASHBOARD SCREEN.....	142

SCREEN LOCATION	142
SCREEN EXAMPLE.....	143
SCREEN FIELDS	143
ON-SCREEN MENU OPTIONS	144
SCREEN INFORMATION.....	144
CONFIGURATION OPTIONS.....	144
ABS DASHBOARD SCREEN	146
SCREEN LOCATION	146
SCREEN EXAMPLE.....	147
SCREEN FIELDS	147
ON-SCREEN MENU OPTIONS	148
SCREEN INFORMATION.....	148
CONFIGURATION OPTIONS.....	148
ALS DASHBOARD SCREEN.....	150
SCREEN LOCATION	150
SCREEN EXAMPLE.....	151
SCREEN FIELDS	151
ON-SCREEN MENU OPTIONS	152
SCREEN INFORMATION.....	152
CONFIGURATION OPTIONS.....	152
DBR DASHBOARD SCREEN.....	154
SCREEN LOCATION	154
SCREEN EXAMPLE.....	155
SCREEN FIELDS	155
ON-SCREEN MENU OPTIONS	155
SCREEN INFORMATION.....	155
CONFIGURATION OPTIONS.....	156
RCS DASHBOARD SCREEN	157
SCREEN LOCATION	157
SCREEN EXAMPLE.....	158
SCREEN FIELDS	158
ON-SCREEN MENU OPTIONS	159
SCREEN INFORMATION.....	159
CONFIGURATION OPTIONS.....	160
PER-PROCESS IO DASHBOARD SCREEN.....	161
SCREEN LOCATION	161
SCREEN EXAMPLE.....	162
SCREEN FIELDS	162
ON-SCREEN MENU OPTIONS	163
SCREEN INFORMATION.....	164
CONFIGURATION OPTIONS.....	164
PER-PROCESS JOURNAL DASHBOARD SCREEN.....	166
SCREEN LOCATION	166
SCREEN EXAMPLE.....	167
SCREEN FIELDS	167

ON-SCREEN MENU OPTIONS	168
SCREEN INFORMATION.....	169
CONFIGURATION OPTIONS.....	169
PER-PROCESS ROW CACHE DASHBOARD SCREEN.....	171
SCREEN LOCATION	171
SCREEN EXAMPLE.....	172
SCREEN FIELDS	172
ON-SCREEN MENU OPTIONS	172
SCREEN INFORMATION.....	173
CONFIGURATION OPTIONS.....	173
RMU SHOW STATISTIC ONLINE ANALYSIS & INFORMATION SCREENS.....	175
BUFFER ANALYSIS SCREEN.....	176
SCREEN LOCATION	176
SCREEN EXAMPLE.....	176
SCREEN FIELDS	176
ON-SCREEN MENU OPTIONS	180
SCREEN INFORMATION.....	180
CONFIGURATION OPTIONS.....	180
TRANSACTION ANALYSIS SCREEN.....	182
SCREEN LOCATION	182
SCREEN EXAMPLE.....	182
SCREEN FIELDS	182
SCREEN INFORMATION.....	185
CONFIGURATION OPTIONS.....	185
AIJ ANALYSIS SCREEN.....	187
SCREEN LOCATION	187
SCREEN EXAMPLE.....	187
SCREEN FIELDS	187
ON-SCREEN MENU OPTIONS	197
ON-SCREEN MENU OPTIONS	197
SCREEN INFORMATION.....	197
CONFIGURATION OPTIONS.....	198
RUJ ANALYSIS SCREEN	199
SCREEN LOCATION	199
SCREEN EXAMPLE.....	199
SCREEN FIELDS	199
ON-SCREEN MENU OPTIONS	201
SCREEN INFORMATION.....	201
CONFIGURATION OPTIONS.....	201
RECOVERY ANALYSIS SCREEN.....	203
SCREEN LOCATION	203
SCREEN EXAMPLE.....	203
SCREEN FIELDS	203
SCREEN INFORMATION.....	206

CONFIGURATION OPTIONS.....	206
RECORD ANALYSIS SCREEN	207
SCREEN LOCATION	207
SCREEN EXAMPLE.....	207
SCREEN FIELDS	207
ON-SCREEN MENU OPTIONS	210
SCREEN INFORMATION.....	211
CONFIGURATION OPTIONS.....	211
AREA ANALYSIS SCREEN.....	212
SCREEN LOCATION	212
SCREEN EXAMPLE.....	212
SCREEN FIELDS	212
ON-SCREEN MENU OPTIONS	214
SCREEN INFORMATION.....	214
CONFIGURATION OPTIONS.....	214
LOCKING ANALYSIS SCREEN	215
SCREEN LOCATION	215
SCREEN EXAMPLE.....	215
SCREEN FIELDS	215
ON-SCREEN MENU OPTIONS	216
SCREEN INFORMATION.....	216
CONFIGURATION OPTIONS.....	217
INDEX ANALYSIS SCREEN	218
SCREEN LOCATION	218
SCREEN EXAMPLE.....	218
SCREEN FIELDS	218
ON-SCREEN MENU OPTIONS	220
SCREEN INFORMATION.....	220
CONFIGURATION OPTIONS.....	220
ROW CACHE ANALYSIS SCREEN	222
SCREEN LOCATION	222
SCREEN EXAMPLE.....	222
SCREEN FIELDS	222
ON-SCREEN MENU OPTIONS	223
SCREEN INFORMATION.....	223
CONFIGURATION OPTIONS.....	224
PROCESS ANALYSIS SCREEN	225
SCREEN LOCATION	225
SCREEN EXAMPLE.....	225
SCREEN FIELDS	225
ON-SCREEN MENU OPTIONS	227
SCREEN INFORMATION.....	227
CONFIGURATION OPTIONS.....	227

Volume 4 Preface

Preface

Oracle Rdb is a general-purpose high-end database management system based on the relational data model. Applications operate in a client-server environment when accessing the database. The RMU Show Statistic utility is used by DBAs to monitor the operation of the application and analyze performance characteristics of the database.

PURPOSE OF THIS
HANDBOOK

This handbook describes the various screens available for the RMU Show Statistic utility. The information contained in this handbook is an amalgamation of various sources of information on the utility, combined into a common frame of reference.

This handbook is not intended to be a reference manual nor should the handbook be construed as a tutorial on how to use the utility. Rather, the information contained in this handbook should be useful for anyone interested in detecting, analyzing and correcting performance problems with the database or application.

INTENDED
AUDIENCE

This handbook is intended for experienced database administrators familiar with the RMU Show Statistic utility. You should also be familiar with database management procedures and terminology.

HANDBOOK
CONTENTS

The RMU Show Statistic “DBA Handbook” is comprised of five volumes. These are the following:

- Volume 1 - Methods and Internals.
 - Volume 2 - Screen Chapters 1 through 7.
 - Volume 3 - Screen Chapters 8 Through 13.
-

- Volume 4 - Screen Chapters 14 Through 22.
- Volume 5 - Appendix Information.

ORACLE CODASYL
DBMS

This handbook discusses the Oracle Rdb product. However, the RMU Show Statistic utility is almost identical for the Oracle CODASYL DBMS product DBO Show Statistic utility; in fact, it uses a common same source code. Therefore, the few statistic screens that are unique to the Oracle DBMS product are included in this handbook.

OPERATING
SYSTEM
INFORMATION

The information contained in this handbook is, for the most part, operating system independent. However, on the rare occasion when an operating system specific item is discussed, the section will be appropriately highlighted.

RELATED MANUALS

For more information on the RMU Show Statistic utility, and database performance tuning information in general, please refer to the following Oracle Rdb documentation:

- Oracle Rdb RMU Reference Manual
- Oracle Rdb Guide to Database Maintenance
- Oracle Rdb Guide to Database Performance and Tuning

In addition, there are several “white papers” discussing RMU Show Statistic issues available in **MetaLink**. For information on the MetaLink system, please contact Oracle Client Relations or visit the Oracle Support webpage at

<http://www.oracle.com/support/>

ACKNOWLEDG -
MENTS

A special word of thanks is extended to Simon Pickering for writing the initial version of the Handbook back in 1995, upon which this new version is loosely based.

A word of thanks is also extended to Kathy Oakey, Bill Gettys, Lilian Hobbs and all the other Oracle Rdb consultants who begged and pleaded for more detailed performance analysis screens to make their job easier. Without these suggestions and ideas, the RMU Show Statistic utility would not be as powerful as it is today.

Finally, I have to thank Anna Logan, who wrote the *Guide to Database Performance and Tuning* manual. She diligently kept after me to keep the “help” documentation up-to-

date, even when I was adding screens to the RMU Show Statistic was utility on a daily basis.

A U T H O R

If you have any comments, suggestions or corrections for this handbook, please forward them to the author at the following email address:

rjanders@us.oracle.com

Chapter

14

Row Cache (One Cache) Screens

RMU Show Statistic Row Cache (One Cache) Screens

The “Row Cache (One Cache)” menu screens display information about specific row caches.

Row Cache Screen

This screen monitors a specific row cache.

The name in parenthesis in the header region of the screen reflects your row cache selection.

This screen resides in the "Row Cache (One Cache)" menu.

SCREEN LOCATION

The following is an example of the "Row Cache (cache)" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor 23-JAN-1998 14:02:07.09
Rate: 1.00 Second        Row Cache (EMPIDS_OVER)           Elapsed: 00:29:31.21
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
```

```
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
latch requests	0	0	0.0	0	0.0	0.0
retried	0	0	0.0	0	0.0	0.0
cache searches	0	0	0.0	0	0.0	0.0
found in workset	0	0	0.0	0	0.0	0.0
found in cache	0	0	0.0	0	0.0	0.0
found too big	0	0	0.0	0	0.0	0.0
insert cache	0	0	0.0	0	0.0	0.0
row too big	0	0	0.0	0	0.0	0.0
cache full	0	0	0.0	0	0.0	0.0
collision	0	0	0.0	0	0.0	0.0
VLM requests	0	0	0.0	0	0.0	0.0
window turns	0	0	0.0	0	0.0	0.0
skipped dirty slot	0	0	0.0	0	0.0	0.0
skipped inuse slot	0	0	0.0	0	0.0	0.0
hash misses	0	0	0.0	0	0.0	0.0
cache unmark	0	0	0.0	0	0.0	0.0

```
-----
```

Exit Graph Help Menu Options Reset Set_rate Time_plot Unreset Write X_plot !

SCREEN FIELDS

latch requests This field displays the total number of latch requests that have occurred. A latch is requested whenever an internal data structure needs to be atomically modified, and

indicates a potential stall point. The duration of the latch is essentially non-measurable.

- retried This field indicates the latch could not be immediately acquired and the latch had to be requested again. This field provides an indication of the contention on the row cache. Ideally, this field should be as close to zero as possible.
- cache searches This field displays the total number of times the row cache was searched for a particular row DBKEY. Values within this field are subdivided into the following fields:
 - found in workset
 - found in cache
 - found too big
- found in workset This field indicates the particular row was found in the process' working set and no additional work was required to fetch the row. This is the ideal situation.
- found in cache This field indicates the particular row was not found in the process' working set, but was found in the global row cache. When this occurs, it is possible that the process will have to make room in the working set by removing an existing row to the global cache.
- found too big This field indicates the particular row was found in the global cache, but the row is now too large to fit into the specified cache buffer. At one time, the row fit into the cache, but it no longer does. Therefore, the cache effectiveness is reduced because a cache entry is reserved for the row but no row caching is possible.
- insert cache This field indicates the total number of times new rows have been inserted into the row cache. Values within this field are subdivided into the following fields
 - row too big
 - cache full
 - collision
- row too big This field indicates the row was initially too large to fit into the specified row cache buffer.
- cache full This field indicates that all row cache entries were modified and could not be flushed to disk in order to make space for the new row. This is an indication that the row cache checkpoint intervals specified for the database may be too large.
- collision This field displays the total number of row cache hash table collisions that occurred when inserting new rows. This field provides an indication of the effectiveness of the cache size.

- VLM requests This field displays the number of VLM requests made.
- window turns This field displays the number of VLM requests made for which the VLM address range was not immediately available.
- skipped dirty slot This field displays the total number of cache entries that could not be replaced because the slot contained modified data rows. This field provides an indication of the relative amount of work required to find space in order to insert a new row into the cache.
- skipped inuse slot This field displays the total number of cache entries that could not be replaced because the slot was referenced by other processes. This field provides an indication of the relative amount of work required to find space in order to insert a new row into the cache.
- hash misses This field displays the total number of times the row cache hash table overflowed.
- cache unmark This field displays the total number of times a row in the cache was flushed to disk, thus making it eligible to be replaced. However, this field does not indicate whether or not the row was emptied from the cache.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN

INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not available if row cache disabled or INPUT qualifier specified unless the binary file was created with OUTPUT=row_cache qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

Chapter

15

Row Cache (One Field) Screens

RMU Show Statistic Row Cache (One Field) Screens

The “Row Cache (One Field)” menu screens display information on a specific statistic field for all row caches.

Row Cache Screen

This screen monitors a specific statistic field for all available row caches.

The name in parenthesis in the header region of the screen reflects your statistic field selection.

SCREEN LOCATION

This screen resides in the "Row Cache (One Field)" menu.

SCREEN EXAMPLE

The following is an example of the "Row Cache (field)" screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  4-FEB-1998 14:37:02.28
Rate: 1.00 Second        Row Cache (latch requests)      Elapsed: 00:27:59.62
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
statistic..... rate.per.second..... total..... ave rage.....
name..... max..... cur..... avg..... count..... per.trans....
REAL_BIG_NAME_TOO_LO      0      0      0.0      0      0.0
EMPIDS_LOW                 0      0      0.0      0      0.0
EMPIDS_MID                 0      0      0.0      0      0.0
EMPIDS_OVER                0      0      0.0      0      0.0
EMPLOYEES                  0      0      0.0      0      0.0
DEPARTMENTS                0      0      0.0      0      0.0
JOB_HISTORY                0      0      0.0      0      0.0
SALARY_HISTORY             0      0      0.0      0      0.0
CANDIDATES                 0      0      0.0      0      0.0
TEST6                      0      0      0.0      0      0.0
TEST7                      0      0      0.0      0      0.0
TEST8                      0      0      0.0      0      0.0
TEST9                      0      0      0.0      0      0.0
TEST10                     0      0      0.0      0      0.0
TEST11                     0      0      0.0      0      0.0
TEST12                     0      0      0.0      0      0.0
TEST13                     0      0      0.0      0      0.0
TEST14                     0      0      0.0      0      0.0
TEST15                     0      0      0.0      0      0.0
```

```
-----
Exit Graph Help Menu >next_page <prev_page Options Reset Set_rate Time_plot Writ
```

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Reset This on-screen menu option temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is not available if row cache disabled or INPUT qualifier specified unless the binary file was created with OUTPUT=row_cache qualifier.
- Cluster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.
- Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
 OPTIONS

This screen does not have any configuration options.

Row Cache Information Screens

RMU Show Statistic Row Cache Information Screens

The “Row Cache Information” menu screens display detailed summary information for all row caches.

Row Cache Utilization Screen

This screen provides utilization information in a graphical format for each row in a specific row cache.

This screen is organized in units of cache entries. Each character of the screen identifies a single cache entry.

The number of users sharing a particular row cache entry is known as the share-count. This screen is based on a display-threshold that is compared against the number of users sharing a particular row. Initially, the display threshold value is "1". The display threshold can be configured, as described below.

Cache entries are identified as follows:

A row of "X" marks delineates the end of the row cache. All blank lines following the row of "X" marks can be ignored.

SCREEN LOCATION

This screen resides in the "Row Cache Information" menu.

SCREEN EXAMPLE

The following is an example of the "Row Cache Utilization" screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  4-FEB-1998 15:25:10.34
Rate: 1.00 Second        Row Cache Utilization           Elapsed: 01:16:07.68
Page: 1 of 2            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
-----
                                For Cache: EMPIDS_OVER
Legend: cursor:_ <thrshld:. >thrshld:#, =thrshld:=, modified: , hot:H toobig:T
Cache ID: 4      #Rows: 3000      RowSize: 300
                =      ===
                2=
==
2
==
==
==
H      H
=====
H      =
H      =
-----
Config Exit Help Menu >next_page <prev_page Options Set_rate Write !

```

SCREEN FIELDS

- Space (" ") Cache entry is empty.
- Dot (.) A cache entry with a share-count less than the display-threshold.
- Equal sign (=) A cache entry with a share-count identical to the display-threshold.
- Digit from 2 through 9 A cache entry with share-count that exceeds the display-threshold but is less than ten users, identified by the respective digit.
- Asterisk (*) A cache entry with share-count that exceeds the display-threshold by 10 users or more.
- Highlighted entry Cache entry has been modified.
- T** Cache entry is too large for the buffer.

ON-SCREEN MENU
OPTIONS

- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION



As a general measure of row cache utilization when using the default display-threshold, a display containing a lot of blanks or dots is very bad; this probably indicates that the row cache is sized too large for the current number of users accessing the database. A display containing a lot of equal signs (=) is bad since this indicates non-sharing of rows; this probably indicates that using a row cache may not be appropriate. A display containing a lot of numeric digits is good, since this indicates good row sharing. A display containing a lot of asterisks(*) is great, since this indicates tremendous row sharing.

- Page Navigation If there is more information than can fit on one page, the notation "Page 1 of n" appears in the header region, where *n* is the total number of pages. You can display successive pages by pressing the right angle bracket (>) key or the Next Screen (↓) key. To display a previous page, press the left angle bracket (<) key or the Previous Screen (↑) key.
- Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

ster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option to filter the “Row Cache Utilization” screen. Select this option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu contains the following options:

Filter Utilization Display Selecting this option will display the prompt “Enter row cache threshold (0+) [current=1]: ”.

The threshold indicates the number of users sharing a row to be displayed; for instance, if you entered “5” then only those rows being accessed by 5 or more users will be displayed. Also, those rows shared by exactly 5 users would be identified using an equal sign (=) and those more than 5 would be identified using the digits 6 through 9 or the asterisk (*) if more than nine users are sharing a row.

Hot Row Information Screen

This screen displays a list of the most frequently accessed rows for a specific row cache.

The “Hot Row Information” screen contains a list of the hottest, or most heavily shared, rows in the row cache. The list is sorted in descending shared frequency and ascending DBKEY sequence.

The sorting algorithm ensures that the first displayed page contains the set of hottest rows while the last displayed page contains the coldest rows.

In order to display the largest possible amount of information, the “Hot Row Information” screen is divided into two columns. The information should be viewed from left-to-right, and top-to-bottom. For example, the left column of row 1 contains the first hottest row, while the right column of row 1 contains the second hottest row, and so on.

SCREEN LOCATION

This screen resides in the “Row Cache Information” menu.

SCREEN EXAMPLE

The following is an example of the “Hot Row Information” screen:

Node: MYNODE (1/1/1) Oracle Rdb X7.0-00 Perf. Monitor 4-FEB-1998 15:29:31.79
 Rate: 1.00 Second Hot Row Information Elapsed: 01:20:29.13
 Page: 1 of 40 DISK\$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1 Mode: Online

For Cache: EMPIDS_OVER (unsorted)

Area:Page:Ln	#Users	State	Length	SlotNo	Area:Page:Ln	#Users	State	Length	SlotNo
55:481:14	0		32	0	55:470:6	0		51	38
55:673:9	0		32	1	55:875:14	0		32	39
55:426:20	0		32	2	55:423:23	0		32	40
55:87:18	0		51	3	55:421:23	0		32	41
55:900:27	0		32	4	55:412:2	0		51	42
55:669:18	0		51	5	55:718:1	0		32	43
55:656:7	0		89	6	55:489:1	0		32	44
55:1096:4	0		32	7	55:458:3	0		32	45
55:480:21	0		32	8	55:715:27	0		32	46
55:677:28	0		32	9	55:685:26	0		51	47
55:466:22	0		32	10	55:477:3	0		70	48
55:488:26	0		32	11	55:889:16	0		32	49
55:1095:24	0		32	12	55:686:16	0		32	50
55:886:16	0		51	13	55:716:22	0		32	51
55:478:7	0		32	14	55:679:8	0		51	52
55:488:18	0 M		51	15	55:716:20	0		32	53
55:673:21	0		51	16	55:533:5	0		32	54
55:117:18	0		51	17	55:495:2	0		51	55
55:1095:22	0		32	18	55:1089:11	0		32	56
55:682:24	0		32	19	55:475:16	0		32	57
55:680:28	0		32	20	55:673:4	0		51	58
55:476:8	0		51	21	55:664:3	0		32	59
55:405:20	0		32	22	55:407:17	0		32	60
55:298:20	0		51	23	55:256:22	0		51	61
55:655:21	0		32	24	55:533:3	0		32	62
55:132:13	0		70	25	55:884:16	0		51	63
55:131:23	0		32	26	55:675:11	0		32	64
55:133:18	0		127	27	55:1096:6	0		32	65
55:666:23	0		32	28	55:689:2	0		32	66
55:446:22	0		32	29	55:480:13	0		32	67
55:716:19	0		32	30	55:472:7	0		70	68
55:422:15	0		32	31	55:871:21	0		32	69
55:417:7	0		70	32	55:672:2	0 MH		32	70
55:184:23	0		51	33	55:169:21	0		32	71
55:716:17	0		32	34	55:530:27	0		32	72
55:690:8	0		32	35	55:682:20	0		32	73
55:886:14	0		51	36	55:901:16	0		32	74
55:680:25	0 M		32	37	55:679:17	0		32	75

Config Exit Help Menu >next_page <prev_page Options Set_rate Write Zoom !

SCREEN FIELDS

Area:Page:Ln This field identifies the DBKEY of the row. These are logical area DBKEYS. In the above example, the area 55 refers to logical area 55. For more information on how to identify the logical area name, please refer to Appendix G.

The value “Reserved” indicates a cache entry that is reserved for future use.

The value “Empty:*n*” indicates an empty cache entry with its corresponding slot number for identification.

#Users This field indicates the number of users sharing the row. Ideally, this field should not contain the value “0” since this indicates wasted row cache entries.

State This field indicates the various states of the row in the cache. Values are the following:

<u>Keyword</u>	<u>Description</u>
H	The row is “hot”.
M	The row is “modified”.
T	The row is “too big”.

SlotNo This field identifies the internal slot number of the cache entry.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation If there is more information than can fit on one page, the notation “Page 1 of *n*” appears in the header region, where *n* is the total number of pages. You can display successive pages by pressing the right angle bracket (>) key or the Next Screen (↓) key. To display a previous page, press the left angle bracket (<) key or the Previous Screen (↑) key.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen The zoom screen displays row cache information on its zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option to configure the “Hot Row Information” screen. Select this option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu contains the following options:

Unsorted display This configuration option displays the “Hot Row Information” screen in an unsorted order.

Sort by frequency This configuration option displays the “Hot Row Information” screen sorted in descending frequency order.

Sort by DBKEY This configuration option displays the “Hot Row Information” screen sorted in ascending DBKEY order.



Note that “Reserved” entries are displayed first since they represent potential DBKEYs.

Sort by “Marked” This configuration option displays the “Hot Row Information” screen sorted by modified rows in ascending DBKEY order.

Sort by “Hot” This configuration option displays the “Hot Row Information” screen by hottest rows in ascending DBKEY order.

Sort by “Too Big” This configuration option displays the “Hot Row Information” screen by too large rows in ascending DBKEY order.

Row Cache Status Screen

This screen provides overall status information for a specific row cache. The information displayed is extremely dynamic and changes constantly on a high-throughput system.

The "Row Cache Status" screen can also be configured to fine-tune the information displayed.

This screen resides in the "Row Cache Information" menu.

SCREEN LOCATION

The following is an example of the "Row Cache Status" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  5-FEB-1998 12:35:36.30
Rate: 1.00 Second        Row Cache Status              Elapsed: 00:00:23.08
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

 For Cache: EMPIDS_OVER


Statistic.Name Stat.Value Percent

```
Total slots:      3000 100.0% Slot Length: 300 Hash slots: 4096
Slots full:       0    0.0% Use:           0    0.0%
Slots empty:     3000 100.0% Rsv:          0    0.0%
Marked Slots:    0    0.0% Hot:           0    0.0% Cold:           0    0.0%
Clean Slots:     3000 100.0% Hot:           0    0.0% Cold:           3000 100.0%
Used Space:       0k   0.0%
Wasted Space:    0k   0.0%
Free Space:      900k 100.0%
```

```
Hash Que Lengths: Empty:4096 1:0      2:0      3:0      4+:0
Cursor position:  0 of 3000 wrapped 0 times
Cache latched:   No
Cache is full:   No
Cache modified:  No
Section name:    RDM70R_2H73LA80 (2075 pages each 512 bytes)
Number of checkpoints: None
Cache Recovery:  157:2
```

 Config Exit Help Menu Options Set_rate Write !

SCREEN FIELDS

- Statistic Name** This field identifies the row cache statistic located directly underneath this column header.
- Stat Value** This field identifies the corresponding statistic field value.
- Percent** This field identifies the corresponding statistic field value expressed as a percentage of the corresponding maximum value.
- Total slots** This field identifies the total number of rows that can be stored in the row cache.
- Slot Length** This field identifies the maximum record length, including internal overhead, that can be stored in a cache entry.
-  Note that this means the actual maximum user record length will be at least five bytes smaller than this value.
- Hash slots** This field identifies the total number of internal hash table slots are used to map records stored in the row cache.
- Slots full** This field identifies the total number of slots that are full.
- Use field**
This field identifies the total number of slots that are full and referenced by one or more users.
- Slots empty** This field identifies the total number of slots that are empty.
- Rsv** This field identifies the total number of slots that are reserved.
- Marked Slots** This field identifies the total number of slots that are currently modified.
- Hot** This field identifies the total number of slots that are currently modified and hot, which is after-image journal information has been written since the last checkpoint.
- Cold field**
This field identifies the total number of slots that are currently modified and cold.
- Clean Slots** This field identifies the total number of slots that are currently un-modified.
- Hot** This field identifies the total number of slots that are currently un-modified and hot, which is after-image journal information has been written since the last checkpoint.
- Cold field**
This field identifies the total number of slots that are currently modified and cold
- Used Space** This field identifies the total size of all rows actually stored in the row cache.

Wasted Space	This field identifies the total wasted space in the cache, which is the total size of the cache minus the total used space (including empty rows).
Free Space	This field identifies the total space wasted by empty rows in the cache
Hash Que Lengths	This field identifies the hash table queue lengths. The starting hash queue length can be configured using the “Define minimum hash chain length” configuration option.
Cursor position	This field identifies the row cache slot number where the cursor resides.
Cache latched	This field indicates whether or not the row cache is latched.
Cache is full	This field indicates whether or not the row cache is completely full.
Cache modified	This field indicates whether or not the row cache has been modified.
Section name	This field indicates the row cache global section name.
Number of checkpoints	This field indicates the total number of checkpoints that have occurred for this row cache.
Cache Recovery	This field identifies the current row cache checkpoint location in the after-image journal.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options	This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
---------	---

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation	This screen does <u>not</u> contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
-----------------	---

Binary File Support	This screen is <u>not</u> recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.
---------------------	---

Cluster Statistic Collection Support	This screen is <u>not</u> integrated into the “Cluster Statistic Collection” facility.
--------------------------------------	--

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option to configure the “Row Cache Status” screen. Select this option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu contains the following options:

ine minimum hash chain length This configuration option specifies the starting hash table queue length to be displayed by the “Hash Que Lengths” field.

Row Cache Queue Length Screen

This screen helps you determine the relative CPU performance impact of row caching. Poor CPU performance can result if many hash table collisions occur.

Each character position of the screen identifies a specific hash table entry. The character, if any, identifies the length of that entry's queue. The asterisk (*) denotes a queue length of 10 or more.

The line of "X" marks indicates the end of the display. All blank lines following the row of "X" marks can be ignored.

The number of rows in the row cache affects the size of the corresponding hash table. Hash table queue lengths can be controlled by increasing or decreasing the number of row cache entries.

SCREEN LOCATION

This screen resides in the "Row Cache Information" menu.

SCREEN EXAMPLE

The following is an example of the "Row Cache Queue Length" screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  5-FEB-1998 15:16:19.25
Rate: 1.00 Second        Row Cache Queue Length          Elapsed: 00:09:53.29
Page: 1 of 2            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
                          For Cache: EMPIDS_OVER

Cache ID: 4    #Rows: 3000    Table Size: 4096
1 1           1 1           1 1           1 1           1 1           1 1           1 1           1 1
      1 1           1 1           1 1           1 1           1 1           1 1           1 1           1 1
    111         111         111         111         111         111         111         111
  111         111         111         111         111         111         111         111
11          111         111         111         111         111         1111 111         111
 111         111         111         111         111         111         111         111
111         111         111         111         111         111         111         111
1          111         111         111         111         111         111         111
      111         111         111         111         111         111         111         111
    111         111         111111         111         111         111         111         111
  111         111         111         111         111         111         111         111
111         111         111         1 11         111         111         111         111
      111         111         111         111         11         11         11         11
1          11         11         11         11         11         11         11         11
      11         11         11         11         11         11         11         11
11         11         11         11         11         11         1111         11
11         11         11         11         11         11         11         11
      11         11         11         11         11         11         11         11
-----
Config Exit Help Menu >next_page <prev_page Options Set_rate Write !
  
```

SCREEN FIELDS


- Cache ID This field identifies the internal identifier for the row cache.
- #Rows This field identifies the total number of rows in the row cache.
- Table Size This field identifies the total number of hash table entries for the row cache.


ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

 Ideally, every space of the screen should contain a the value “1”. A display containing long queues and a lot of blanks indicates poor CPU utilization.

 Also note that a spiral effect, as displayed in the above screen example, usually indicates poor hashing and the size of the cache should potentially be adjusted larger or smaller to influence the hash table size.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.
- Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.
- Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option to filter the “Row Cache Queue Length” screen. Select this option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu contains the following options:

Filter Utilization Display Selecting this option will display the prompt “Enter row cache threshold (0+) [current=1]: ”.

Any queue length which is less than the specified utilization threshold will not be displayed. By default, the utilization threshold is “0”, so that all entries on the screen are displayed, even those of length 0.


Row Length Distribution Screen

This screen graphically describes the distribution of the various row lengths within a particular row cache. It is necessary to know the distribution of the various row lengths in the row cache to determine how well sized the row cache is. For instance, knowing that each cache entry is wasting even 5 bytes is significant if there are 100,000 entries in the cache.

The summary region displays three lines containing information about the row cache and boundary row lengths. The first line identifies the selected row cache name.

The information displayed is based on a scale that represents the row lengths to be analyzed. Initially, the screen's scale is based on the cache's specified row size. This scale can be changed using the "Config" on-screen menu option.

Each vertical line of the screen identifies a row length bucket, based on the selected screen scale. Each asterisk (*) in the bucket represents a particular number of rows in the cache.

 Note that a value of "10", for instance, means that up to 10 rows are represented by that asterisk (*).

Two sliding indicators are displayed along the bottom horizontal axis; "A" indicates the average row size and "M" indicates the median row size.

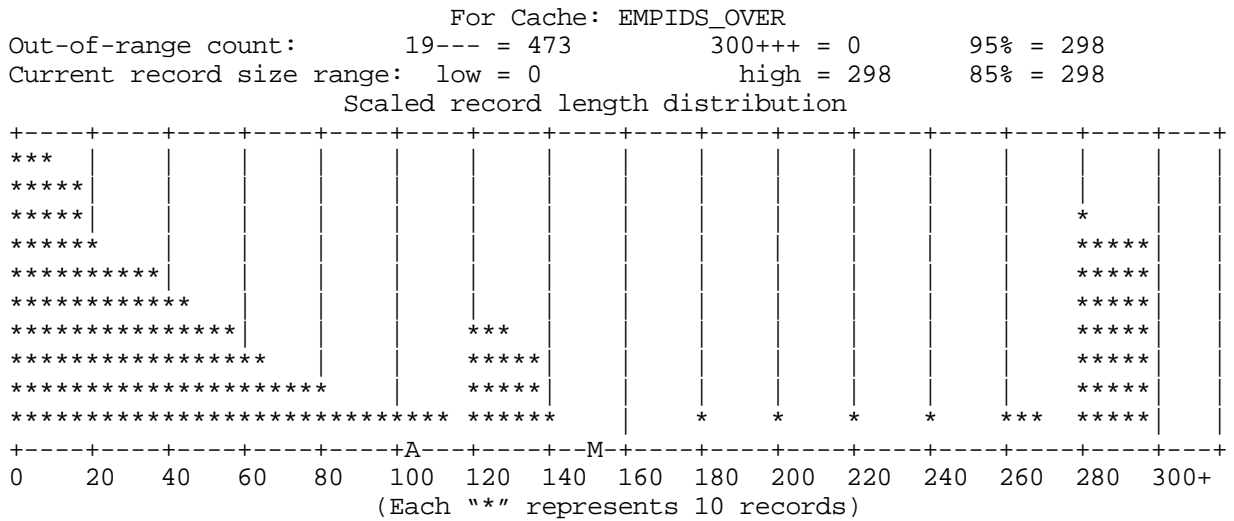
SCREEN LOCATION

This screen resides in the "Row Cache Information" menu.

SCREEN EXAMPLE

The following is an example of the "Row Length Distribution" screen:

Node: MYNODE (1/1/1) Oracle Rdb X7.0-00 Perf. Monitor 5-FEB-1998 15:27:17.29
 Rate: 1.00 Second Row Length Distribution Elapsed: 00:20:51.33
 Page: 1 of 1 DISK\$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1 Mode: Online



 Config Exit Help Menu Options Set_rate Write !

SCREEN FIELDS

Out-of-range count This field identifies the number of rows whose lengths were below the minimum selected row length (XXX--) or above the maximum selected row length (XXX+++). Also, the 95th percentile row length is displayed; this value indicates the row size that 95% of the rows in the cache are less than or equal to.

Current record size range This field identifies the lowest and highest row length in the cache. Also, the 85th percentile row length is displayed; this value indicates the row size that 85% of the row in the cache are less than or equal to.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION



The number of collection buckets is based on the display width of the terminal. By default, each bucket comprises 5 columns of the display; the number of columns can be configured by the user using the “Config” menu option. For example, setting the terminal width to 132 columns will allow you to display more information than with an 80-column terminal width. Decreasing the number of columns per bucket will achieve the same effect on an 80-column terminal width, but with a corresponding loss of precision.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not available if row cache disabled or INPUT qualifier specified unless the binary file was created with OUTPUT=row_cache qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option to filter the “Row Length Distribution” screen. Select this option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu contains the following options:

the number of bucket columns Decreasing the number of columns per bucket allows you to display more row sizes. Conversely, increasing the number of columns per bucket allows you to display fewer row sizes, but with more accurate size reporting (i.e. better scaling) per bucket. The minimum number of columns is 2 and the maximum is 10.

Set the row length range values Setting a specific row size range values allows you to determine the actual row size display range, which is by default based on the row cache database parameters.

Eliminate empty slots This configuration option allows you to eliminate the empty slots from the display. This gives you a distribution of just used slots.

Include empty slots This configuration option allows you to include the empty slots in the display. This gives you an overview row length distribution analysis.

RCS Statistics Screen

This screen provides information about the runtime operation of the Row Cache Server (RCS) process.

SCREEN LOCATION

This screen resides in the “Row Cache Information” menu.

SCREEN EXAMPLE

The following is an example of the “RCS Statistics” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  5-FEB-1998 15:36:13.38
Rate: 1.00 Second        RCS Statistics                Elapsed: 00:29:47.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....      rate.per.second..... total..... ave rage.....
name.....           max..... cur..... avg..... count..... per.trans....
Find buffer stall           0         0         0.0           1           0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

Find buffer stall This field gives the length of time (in hundredths of a second) required to allocate a checkpoint buffer. Typically, this value should be extremely small.

ON-SCREEN MENU OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph** This on-screen menu option displays a histogram graph instead of numbers.
- Numbers** This on-screen menu option displays numeric statistics instead of histogram graph.
- Options** This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

- Reset** This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot** This on-screen menu option plots a specific field's value by time.
- Unreset** This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot** This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
INFORMATION

This section discusses screen-specific issues.

- Page Navigation** This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support** This screen is not available if row cache disabled or INPUT qualifier specified unless the binary file was created with OUTPUT=row_cache qualifier.
- Cluster Statistic Collection Support** This screen is integrated into the “Cluster Statistic Collection” facility.
- Zoom Screen** This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

Chapter

17

Index Information Screens

RMU Show Statistic Index Information Screens

The “Index Information” menu screens display summary statistic information about btree and hash indexes.

Index Statistics (Retrieval) Screen

This screen monitors how much retrieval activity is taking place in a database's sorted indexes.

Oracle Rdb often uses direct index lookups and index scans to access records in the database. This screen monitors these operations as well as the number of index nodes fetched.

This screen resides in the "Index Information" menu.

SCREEN LOCATION

The following is an example of the "Index Statistics (Retrieval)" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 15:45:33.67
Rate: 1.00 Second        Index Statistics (Retrieval)      Elapsed: 00:02:24.84
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
```

```
-----
statistic.....      rate.per.second..... total..... ave rage.....
name.....           max..... cur..... avg..... count..... per.trans....

transactions          0          0          0.0          0          0.0
verb successes        0          0          0.0          0          0.0
verb failures         0          0          0.0          0          0.0


node fetches          0          0          0.0          0          0.0
  leaf fetches        0          0          0.0          0          0.0
  dup. fetches         0          0          0.0          0          0.0

index lookups         0          0          0.0          0          0.0
index scans           0          0          0.0          0          0.0
  primary entries     0          0          0.0          0          0.0
  dup. entries        0          0          0.0          0          0.0
```

```
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
```

SCREEN FIELDS

transactions This field gives the number of completed database transactions. This is the count of the COMMIT and ROLLBACK statements that have executed.

- verb successes This field gives the number of completed verbs that returned a successful status code.
- A verb is an atomic SQL statement or action. For example, a record insert is a verb, as is a record deletion.
- Also, within a compound statement each individual statement is atomic and Oracle Rdb performs a verb-success operation after processing each one. To avoid this overhead, you can use the SQL BEGIN atomic statement to treat the entire block as a single verb.
- verb failures This field gives the number of completed verbs that returned an error status code. Errors include end-of-collection and deadlocks, as well as all other exception conditions.
- A verb is an atomic SQL statement or action. For example, a record insert is a verb, as is a record deletion.
- Excessive verb failures are usually an indication of a failed constraint, such as uniqueness criteria, or an invalid DDL statement.
-  Note that in the case of cursors and scans, reaching the end-of-stream always results in a verb failure.
- Note that SQL performs its own internal queries to identify metadata, such as relation or index names.
- Oracle Rdb rarely issues a verb-failure unless there is an exception of some kind, such as a constraint failure.
- node fetches This field gives the number of times Oracle Rdb fetched an index node during index retrievals. This number includes the number of leaf nodes and duplicate nodes fetched. Therefore, the calculation for the number of upper-level index nodes accessed is: this “node fetches” field minus the sum of the leaf and duplicate node fetches. The result can indicate the depth of the database indexes.
- leaf fetches This field gives the number of times Oracle Rdb fetched bottom level (leaf) nodes during index retrievals. This number, along with the “index scans” field, can indicate the length of scans in terms of index nodes accessed. There is one leaf node fetch for each “index lookup” retrieval.
- dup fetches This field gives the number of times Oracle Rdb fetched a duplicate node (as opposed to a leaf node) during index retrievals. This number can indicate the lengths of duplicate node chains in the database indexes. When a duplicate node is retrieved, the operation always includes one leaf fetch.
- index lookups This field gives the number of direct single-key retrievals performed on the database indexes. This statistic shows up only on unique key retrievals and not on duplicate key retrievals.

- index scans This field gives the number of scans, or range retrievals, performed on the database indexes. In an index scan, Oracle Rdb searches an index from top to bottom to find the starting point (low value) of the retrieval. Oracle Rdb then searches the bottom level nodes of the index, including duplicate nodes, until the scan's end condition is met.
- primary entries This field gives the number of unique keys found during the index scan.
- dup entries This field gives the number of duplicate keys found during the index scans. If an index has two entries with the same key value, the first one is a primary entry and the second is a duplicate entry.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

This screen does not have any configuration options.

C O N F I G U R A T I O N
O P T I O N S

Index Statistics (Insertion) Screen

This screen monitors the update activity of a database's sorted indexes during insertions; that is, when you store or modify an index key field or when you use the SQL CREATE INDEX statement on a table. This screen also indicates in which type of index node the insertions occur and displays node creations by node type.

By examining this screen, you can monitor how a database balances its sorted indexes after insertions into the database.

SCREEN LOCATION

This screen resides in the "Index Information" menu.

SCREEN EXAMPLE

The following is an example of the "Index Statistics (Insertion)" screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 15:45:34.81
Rate: 1.00 Second        Index Statistics (Insertion)      Elapsed: 00:02:25.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....              max..... cur..... avg..... count..... per.trans....

node insertions          0          0          0.0          0          0.0
root insertions          0          0          0.0          0          0.0
leaf insertions          0          0          0.0          0          0.0
dup. insertions          0          0          0.0          0          0.0

node creations           0          0          0.0          0          0.0
root splits              0          0          0.0          0          0.0
leaf creations           0          0          0.0          0          0.0
dup. creations           0          0          0.0          0          0.0

index creations          0          0          0.0          0          0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

node insertions This field gives the number of index entries inserted into all index nodes. This number includes root, leaf, and duplicate entries within user-and system-defined indexes.

This number is greater than the number of records being stored in the database because it usually takes one to two insertions into an index for each record for each index. The calculation of node insertions minus the sum of the root, leaf, and duplicate insertions yields the number of entries inserted into mid-level nodes. This number and the “root insertions” field indicate sorted balancing activity.

- root insertions This field gives the number of entries inserted into the root (top-level) index nodes. The number of insertions should be small except for when you load the database. Also, if an index consists of only one node, insertions into this node are not included in this field, but are included in the “leaf insertions” field.
- leaf insertions This field gives the number of unique keys inserted into the database's indexes. This field indicates the number of entries inserted into the leaf (bottom-level) index nodes.
- dup insertions This field gives the number of duplicate index keys inserted into the database's indexes. There should be a one-to-one correspondence to the number of duplicate records being stored in the tables.
- node creations This field gives the total number of index nodes created during insertion of index entries into the index trees. This includes root, leaf, and duplicate nodes created within user- and system-defined indexes. Nodes are created three ways:
- When an index is first defined
 - When a node cannot accommodate an insertion, causing it to overflow into a new node (node splitting)
 - When the first duplicate for a particular key is inserted into an index, causing a duplicate node to be created

The total number of nodes created and the associated fields should be relatively small, except for an initial load of the database with indexes already defined, or for creation of indexes on already-stored data.

- root splits This field gives the number of times the root nodes have split because they overflowed after an insertion. A root node split causes the index to grow by one level—a parent node must be created to point to the two halves of the overflowed root node. Therefore, two nodes are created—the parent node and the node for the second half of the root node. Increasing the number of tree levels means Oracle Rdb must search more index nodes to access a data row; this can result in additional I/O operations.
- leaf creations This field gives the number of times a leaf (bottom level) node was created because an existing leaf node had become full and needed to accommodate another unique index key entry.
- dup creations This field gives the number of times a duplicate node was created to accommodate more duplicated entries within the duplicate index node or on the first store of a duplicate key entry.

index creations This field gives the number of times an index was created on a particular table. This count is the number of CREATE INDEX statements. Also, if an index is partitioned over three areas, for example, there will be a count of three index creations.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Graph This on-screen menu option displays a histogram graph instead of numbers.

Numbers This on-screen menu option displays numeric statistics instead of histogram graph.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.

Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.

Time_plot This on-screen menu option plots a specific field's value by time.

Unreset This on-screen menu option restores the original statistic values after using "Reset" option.

X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

Index Statistics (Removal) Screen

This screen monitors the update activity of a database's sorted indexes when you perform any removal operation; that is, erase, alter, or modify an index key field or drop or delete an index. This screen indicates from which type of index node the removals occur. It also shows node deletions by node type.

This screen lets you monitor how a database balances its sorted indexes when nodes are removed from the indexes.

This screen resides in the "Index Information" menu.

SCREEN LOCATION

The following is an example of the "Index Statistics (Removal)" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 15:45:35.75
Rate: 1.00 Second        Index Statistics (Removal)      Elapsed: 00:02:26.92
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
statistic.....      rate.per.second..... total..... ave rage.....
name.....           max..... cur..... avg..... count..... per.trans....

node removals          0      0      0.0      0      0.0
root removals         0      0      0.0      0      0.0
leaf removals         0      0      0.0      0      0.0
dup. removals         0      0      0.0      0      0.0

node deletions         0      0      0.0      0      0.0
leaf deletions        0      0      0.0      0      0.0
dup. deletions        0      0      0.0      0      0.0

index destructions    0      0      0.0      0      0.0
```

```
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
```

SCREEN FIELDS

node removals This field gives the total number of index entries within the root, leaf, and duplicate nodes that have been removed. This removal can be triggered by erasing rows, deleting tables, or deleting indexes. The calculation of node removals minus the sum of

the root, leaf, and duplicate node removals yields the number of entries removed from mid-level nodes. A node is not deleted until all its entries are removed.

- root removals This field gives the number of index entries removed from a root node due to deletion of entries within lower-level nodes. If an index consists of only one node, removals from this node are not included in this field, but are included in the “leaf removals” field.
- leaf removals This field gives the number of unique index keys removed from the leaf nodes during an SQL DELETE operation.
- dup removals This field gives the number of duplicate index keys removed from duplicate nodes due to the deletion of duplicate records. This should be a one-to-one correspondence to the number of erased duplicate records within the database.
- node deletions This field gives the total number of index nodes deleted due to an SQL DROP. When an index is deleted, this number should be equal to the total number of index nodes within the index. This field minus the sum of leaf and duplicate node deletions yields the number of mid-level node deletions.
- leaf deletions This field gives the number of leaf (bottom level) nodes deleted from the database's indexes. A leaf node is deleted only when it becomes empty.
- dup deletions This field gives the number of duplicate node deletions within the indexes.
- index destructions This field gives the number of indexes deleted with an SQL DROP INDEX statement. This count will be 1 if the index is not partitioned. If an index that is partitioned over three areas is deleted, for example, then the count will be 3. This count also gives the number of root node deletions.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.

Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.

Time_plot This on-screen menu option plots a specific field's value by time.

Unreset This on-screen menu option restores the original statistic values after using “Reset” option.

X_plot This on-screen menu option plots a specific field’s value using scatter-based display.

This section discusses screen-specific issues.

SCREEN
INFORMATION

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

This screen does not have any configuration options.

CONFIGURATION
OPTIONS

Hash Index Statistics Screen

This screen monitors the update and retrieval activity of a database's hashed indexes. It indicates the total number of key insertions and deletions. It also indicates the number of scans that were opened.

For retrievals (successful fetches), the screen indicates the total number of nodes (either bucket fragments or duplicate nodes) that were fetched.

This screen resides in the "Index Information" menu.

SCREEN LOCATION

The following is an example of the "Hash index statistics" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 15:45:36.80
Rate: 1.00 Second        Hash index statistics           Elapsed: 00:02:27.97
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
statistic.....      rate.per.second..... total..... ave rage.....
name.....           max..... cur..... avg..... count..... per.trans....

hash insertions      0          0          0.0         0          0.0
  duplicates         0          0          0.0         0          0.0

hash deletions       0          0          0.0         0          0.0
  duplicates         0          0          0.0         0          0.0

hash scans           0          0          0.0         0          0.0

hash index fetches   0          0          0.0         0          0.0
  bucket fragments   0          0          0.0         0          0.0
  duplicate nodes    0          0          0.0         0          0.0
```

```
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
```

SCREEN FIELDS

hash insertions This field gives the number of hash key insertions in the database's hashed indexes. It includes unique key insertions as well as duplicate key insertions.

duplicates This field gives the number of duplicate key updates in the database's hashed indexes.

- hash deletions This field gives the number of hash key deletions from the database's hashed indexes. It includes unique key deletions as well as duplicate key deletions.
- duplicates This field gives the number of duplicate key deletions in the database's hashed indexes.
- hash scans This field gives the number of hashed index scans, including both retrieval and update scans, that were opened on the database's hashed indexes. A scan is defined as the sequential processing of the records that meet the search criteria of a query. Hashed scans then refer to the case where duplicate records are returned that meet the search criteria of a query from a scan of the hashed index.
- hash index fetches This field gives the number of hashed index nodes that were fetched on a successful search of the database's hashed indexes. This includes fetches of duplicate nodes as well as bucket fragment nodes.
- bucket fragments This field gives the number of bucket fragments that were fetched on a successful search of the database's hashed indexes.
- duplicate nodes This field gives the number of duplicate nodes that were fetched on a successful search of the database's hashed indexes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN

INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

Oracle
 DBMS
 only

Index Statistics Screen

This screen monitors the database b-tree (index) activity and also gives an indication of the transaction and verb activity. Both retrieval and update operations are monitored.

This screen resides in the "Index Information" menu.

SCREEN LOCATION

The following is an example of the "Verb Statistics (1)" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/16)      Oracle CODASYL DBMS X7.0-00      30-JAN-1998 08:30:30.61
Rate: 1.00 Second          Verb Statistics (1)          Elapsed: 00:00:14.67
Page: 1 of 1              DISK$:[DB_HOMEDIR.WORK.STATS]MUT.ROO;1      Mode: Online
```

```
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....
	max.....	cur.....	avg.....		
transactions	0	0	0.0	0	0.0
verb successes	0	0	0.0	0	0.0
verb failures	0	0	0.0	0	0.0
search index	0	0	0.0	0	0.0
create index	0	0	0.0	0	0.0
destroy index	0	0	0.0	0	0.0
insert entry	0	0	0.0	0	0.0
modify entry	0	0	0.0	0	0.0
remove entry	0	0	0.0	0	0.0
balance index	0	0	0.0	0	0.0
one-way scroll	0	0	0.0	0	0.0
two-way scroll	0	0	0.0	0	0.0
three-way scroll	0	0	0.0	0	0.0

```
-----
```

Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !

SCREEN FIELDS

transactions This field gives the number of completed database transactions. This is the count of the COMMIT and ROLLBACK statements that have executed.

verb successes This field gives the number of completed verbs that returned a successful status code.

A verb is an atomic SQL statement or action. For example, a record insert is a verb, as is a record deletion.

Also, within a compound statement each individual statement is atomic and Oracle Rdb performs a verb-success operation after processing each one. To avoid this overhead, you can use the SQL BEGIN atomic statement to treat the entire block as a single verb.

verb failures This field gives the number of completed verbs that returned an error status code. Errors include end-of-collection and deadlocks, as well as all other exception conditions.

A verb is an atomic SQL statement or action. For example, a record insert is a verb, as is a record deletion.

Excessive verb failures are usually an indication of a failed constraint, such as uniqueness criteria, or an invalid DDL statement.



Note that in the case of cursors and scans, reaching the end-of-stream always results in a verb failure.

Note that SQL performs its own internal queries to identify metadata, such as relation or index names.

Oracle Rdb rarely issues a verb-failure unless there is an exception of some kind, such as a constraint failure.

search index This field gives the number of times an index has been searched. Although an index search may be performed as a result of a FIND statement, it may also be performed as a side-effect of many other operations (like a CONNECT or DISCONNECT statement).

create index This field gives the number of times an index has been created. An index creation occurs when the first member record is connected into an empty index-mode set occurrence.

destroy index This field gives the number of times an index has been destroyed. An index destruction occurs when the final member record is disconnected from an index-mode set occurrence, causing that set to become empty.

insert entry This field gives the number of times an entry has been added to an index. An index entry is inserted for each member record that has a unique sort key value. An index insertion may require an index rebalance operation.

modify entry This field gives the number of times an entry has been modified in an index. An index entry is modified when a new member is added to the set, but the sort key value for that new member is a duplicate of a sort key value that is already in the set, and the duplicates order clause specifies DUPLICATES ARE FIRST. An index entry is also

modified when the first member with a given sort key value is removed from the set, but there are other (duplicate) members still in the set. An entry modification may require an index rebalance operation.

- remove entry This field gives the number of times an entry has been removed from an index. An index entry is removed when the final member record with a given sort key value is removed from the set. An entry removal may require an index rebalance operation.
- balance index This field gives the number of times an index has been dynamically re-balanced. An index balance occurs when a node gets too large or too small by adding, modifying or removing entries.
- one way scroll This field gives the number of times that two index nodes were merged into one during a balance operation when the nodes became nearly empty.
- two way scroll This field gives the number of times that entry shuffling occurred between two nodes during a balance operation.
- three way scroll This field gives the number of times that two nodes were split into three during a balance operation.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN

INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

Chapter

18

General Information Screens

RMU Show Statistic General Information Screens

The “General Information” menu screens display miscellaneous statistic information.

Name Translation Screen

This screen shows statistics on database dashboard updates and logical name translation. This screen provides a measure of how many logical names are actually being used instead of using the default values stored in the database.

SCREEN LOCATION

This screen resides in the "General Information" menu.

SCREEN EXAMPLE

The following is an example of the "Name translation" screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:12.33
Rate: 1.00 Second          Name translation           Elapsed: 00:01:04.84
Page: 1 of 1      DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....      rate.per.second..... total..... ave rage.....
name.....           max..... cur..... avg..... count..... per.trans....

Dashboard updated           0           0           0.0           0           0.0
Name Translated             8           0           0.5           34          0.0
  Defaulted                 8           0           0.5           33          0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- dashboard updated This field displays the number of times a user has received notification that the database dashboard has been updated.
- name translated This field displays the number of times a logical name has been translated.
- defaulted This field displays the number of times the default value was used for a logical name.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.

- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot This on-screen menu option plots a specific field’s value using scatter-based display.

SCREEN
INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.
- Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.



Objects (One Stat Type) Screens

RMU Show Statistic Objects (One Stat Type) Screens

The “Objects (One Stat Type)” menu screens display information for a specific statistic field for all database rootfile objects.

KROOT Object Screen

This screen displays the statistics for the KROOT object. The KROOT object contains the database control information that describes all of the other database objects.

The majority KROOT updates occur as the result of the SQL ALTER DATABASE statement.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (KROOT object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (KROOT object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....              max..... cur..... avg..... count..... per.trans....

objects fetch shrd           0         0         0.0         0         0.0
objects fetch excl          0         0         0.0         0         0.0
objects refreshed           0         0         0.0         0         0.0
objects modified            0         0         0.0         0         0.0
objects written              0         0         0.0         0         0.0
objects released            0         0         0.0         0         0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.

- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.
- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

FILID Object Screen

This screen displays the statistics for the FILID object. The FILID object contains the storage area information.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (FILID object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (FILID object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....      rate.per.second.....  total.....  ave rage.....
name.....           max.....  cur.....  avg.....  count.....  per.trans....

objects fetch shrd           0           0           0.0           0           0.0
objects fetch excl           0           0           0.0           0           0.0
objects refreshed            0           0           0.0           0           0.0
objects modified             0           0           0.0           0           0.0
objects written              0           0           0.0           0           0.0
objects released             0           0           0.0           0           0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

SEQBLK Object Screen

This screen displays the statistics for the SEQBLK object. The SEQBLK object contains the information on the allocation of sequence numbers such as transaction sequence numbers (TSN) and commit sequence numbers (CSN)..

The SEQBLK is also used to serialize access to the TSNBLK locks.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (SEQBLK object)” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (SEQBLK object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
objects fetch shrd	0	0	0.0	0	0.0	0.0
objects fetch excl	0	0	0.0	0	0.0	0.0
objects refreshed	0	0	0.0	0	0.0	0.0
objects modified	0	0	0.0	0	0.0	0.0
objects written	0	0	0.0	0	0.0	0.0
objects released	0	0	0.0	0	0.0	0.0

```
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.

- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.
- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

TSNBLK Object Screen

This screen displays the statistics for the TSNBLK object. The TSNBLK object contains the information on the last committed transaction.

The number of TSNBLK objects is a function of the maximum number of users in the database; there is one TSNBLK object for every 28 database users (rounded up). For example, a database containing a maximum of 512 users would contain 19 TSNBLK objects.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (TSNBLK object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (TSNBLK object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1   Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....              max..... cur..... avg..... count..... per.trans....

objects fetch shrd           0         0         0.0         0         0.0
objects fetch excl           0         0         0.0         0         0.0
objects refreshed           0         0         0.0         0         0.0
objects modified            0         0         0.0         0         0.0
objects written             0         0         0.0         0         0.0
objects released            0         0         0.0         0         0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.

- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.
- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

- This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.
- Graph This on-screen menu option displays a histogram graph instead of numbers.
 - Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
 - Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
 - Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
 - Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
 - Time_plot This on-screen menu option plots a specific field's value by time.
 - Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
 - X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

AIJDB Object Screen

This screen displays the statistics for the AIJDB object. The AIJDB object contains the after-image journal control information.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (AIJDB object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (AIJDB object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....              max..... cur..... avg..... count..... per.trans....

objects fetch shrd           0         0         0.0         0         0.0
objects fetch excl          0         0         0.0         0         0.0
objects refreshed           0         0         0.0         0         0.0
objects modified            0         0         0.0         0         0.0
objects written              0         0         0.0         0         0.0
objects released            0         0         0.0         0         0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

AJFB Object Screen

This screen displays the statistics for the AJFB object. The AJFB object contains the after-image journal information.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (AJFB object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (AJFB object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....      rate.per.second..... total..... ave rage.....
name.....           max..... cur..... avg..... count..... per.trans....

objects fetch shrd          0          0          0.0          0          0.0
objects fetch excl          0          0          0.0          0          0.0
objects refreshed           0          0          0.0          0          0.0
objects modified            0          0          0.0          0          0.0
objects written              0          0          0.0          0          0.0
objects released            0          0          0.0          0          0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

RTUPB Object Screen

This screen displays the statistics for the RTUPB object. The RTUPB object contains information on active users.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (RTUPB object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (RTUPB object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....      rate.per.second.....  total.....  ave rage.....
name.....           max.....  cur.....  avg.....  count.....  per.trans....

objects fetch shrd          0          0          0.0          0          0.0
objects fetch excl          0          0          0.0          0          0.0
objects refreshed           0          0          0.0          0          0.0
objects modified            0          0          0.0          0          0.0
objects written              0          0          0.0          0          0.0
objects released            0          0          0.0          0          0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

ACTIVE Object Screen

This screen displays the statistics for the ACTIVE object. The ACTIVE object contains information on active transactions.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (ACTIVE object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (ACTIVE object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....              max..... cur..... avg..... count..... per.trans....

objects fetch shrd           0           0           0.0           0           0.0
objects fetch excl           0           0           0.0           0           0.0
objects refreshed            0           0           0.0           0           0.0
objects modified             0           0           0.0           0           0.0
objects written              0           0           0.0           0           0.0
objects released             0           0           0.0           0           0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

CPT Object Screen

This screen displays the statistics for the CPT object. The CPT object contains information on the corrupt page table.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (CPT object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (CPT object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....      rate.per.second.....  total.....  ave rage.....
name.....           max.....  cur.....  avg.....  count.....  per.trans....

objects fetch shrd           0           0           0.0           0           0.0
objects fetch excl          0           0           0.0           0           0.0
objects refreshed           0           0           0.0           0           0.0
objects modified            0           0           0.0           0           0.0
objects written              0           0           0.0           0           0.0
objects released            0           0           0.0           0           0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

RCACHE Object Screen

This screen displays the statistics for the RCACHE object. The RCACHE object contains the row cache information.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (RCACHE object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (RCACHE object)      Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....              max..... cur..... avg..... count..... per.trans....

objects fetch shrd           0           0           0.0           0           0.0
objects fetch excl           0           0           0.0           0           0.0
objects refreshed            0           0           0.0           0           0.0
objects modified             0           0           0.0           0           0.0
objects written              0           0           0.0           0           0.0
objects released             0           0           0.0           0           0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

CLIENT Object Screen

This screen displays the statistics for the CLIENT object. The CLIENT object contains client-specific information.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (CLIENT object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (CLIENT object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....                max..... cur..... avg..... count..... per.trans....

objects fetch shrd           0           0           0.0           0           0.0
objects fetch excl           0           0           0.0           0           0.0
objects refreshed            0           0           0.0           0           0.0
objects modified             0           0           0.0           0           0.0
objects written              0           0           0.0           0           0.0
objects released             0           0           0.0           0           0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

CLTSEQ Object Screen

This screen displays the statistics for the CLTSEQ object. The CLTSEQ object contains the client "sequence" information.

SCREEN LOCATION

This screen resides in the "Objects (One Stat Type)" menu.

SCREEN EXAMPLE

The following is an example of the "Object (CLTSEQ object)" screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (CLTSEQ object)           Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....      rate.per.second.....  total.....  ave rage.....
name.....           max.....  cur.....  avg.....  count.....  per.trans....

objects fetch shrd          0          0          0.0          0          0.0
objects fetch excl          0          0          0.0          0          0.0
objects refreshed          0          0          0.0          0          0.0
objects modified           0          0          0.0          0          0.0
objects written            0          0          0.0          0          0.0
objects released           0          0          0.0          0          0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

UTILITY Object Screen

This screen displays the statistics for the UTILITY object. The UTILITY object contains information used by the Oracle RMU utility.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Type)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (UTILITY object)” screen:

```

Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:22:21.47
Rate: 1.00 Second        Object (UTILITY object)      Elapsed: 00:01:13.98
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
statistic.....          rate.per.second..... total..... ave rage.....
name.....              max..... cur..... avg..... count..... per.trans....

objects fetch shrd           0           0           0.0           0           0.0
objects fetch excl           0           0           0.0           0           0.0
objects refreshed            0           0           0.0           0           0.0
objects modified             0           0           0.0           0           0.0
objects written              0           0           0.0           0           0.0
objects released             0           0           0.0           0           0.0
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
  
```

SCREEN FIELDS

- objects fetch shrd This field displays the number of objects that are fetched for the shared retrieval process.
- objects fetch excl This field displays the number of objects that are fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object actually was updated. This object is sometimes acquired in exclusive mode to serialize access to an operation.
- objects refreshed This field displays the number of objects whose information in the global section was detected as being stale, so the information was read again from the database root file.

- objects modified This field displays the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.
- objects written This field displays the number of objects whose information was written back to the database root file.
- objects released This field displays the number of objects whose shared or exclusive access was released to other processes.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the "Pause" on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the "Unreset" on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using "Reset" option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.
- ster Statistic Collection Support This screen is integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

Objects (One Stat Field) Screens

RMU Show Statistic Objects (One Stat Field) Screens

The “Objects (One Stat Field)” menu screens display information for a specific database rootfile object.

Objects Fetch Shrd Screen

This screen displays the statistics for the “objects fetch shrd” collection category. This category shows the number of objects that are fetched for the shared retrieval operation.

This screen resides in the “Objects (One Stat Field)” menu.

SCREEN LOCATION

The following is an example of the “Object (objects fetch shrd)” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:29:11.65
Rate: 1.00 Second        Object (objects fetch shrd)      Elapsed: 00:00:13.06
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
KROOT object	0	0	0.0	0	0	0.0
FILID object	0	0	0.0	0	0	0.0
SEQBLK object	0	0	0.0	1	1	0.0
TSNBLK object	0	0	0.1	2	2	0.0
AIJDB object	0	0	0.0	1	1	0.0
AIJFB object	2	0	0.8	11	11	0.0
RTUPB object	0	0	0.0	0	0	0.0
ACTIVE object	0	0	0.0	0	0	0.0
CPT object	0	0	0.0	0	0	0.0
RCACHE object	0	0	0.0	0	0	0.0
CLIENT object	0	0	0.0	0	0	0.0
CLTSEQ object	0	0	0.0	0	0	0.0
UTILITY object	0	0	0.0	0	0	0.0

```
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
```

SCREEN FIELDS

KROOT object This field displays the database control information that describes all of the other database objects.

FILID object This field displays the storage area information.

SEQBLK object This field displays the information on the allocation of transaction sequence numbers.

- TSNBLK object This field displays the information on the last committed transaction.
- AIJDB object This field displays the after-image journal control information.
- AIJFB object This field displays the after-image journal information.
- RTUPB object This field displays information on active users.
- ACTIVE object This field displays information on active transactions.
- CPT object This field displays information on the corrupt page table.
- RCACHE object This field displays information on row caches.
- CLIENT object This field displays client-specific information.
- CLTSEQ object This field displays client “sequence” information.
- UTILITY object This field displays Oracle RMU utility information.

ON - SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

Objects Fetch Excl Screen

This screen displays the statistics for the “objects fetch excl” collection category. This category shows the number of objects fetched for exclusive access with the intention of subsequently being updated. This statistic does not indicate that the object was actually updated.

This screen resides in the “Objects (One Stat Field)” menu.

SCREEN LOCATION

The following is an example of the “Object (objects fetch excl)” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:29:11.65
Rate: 1.00 Second        Object (objects fetch excl)      Elapsed: 00:00:13.06
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
KROOT object	0	0	0.0	0	0.0	
FILID object	0	0	0.0	0	0.0	
SEQBLK object	0	0	0.0	1	0.0	
TSNBLK object	0	0	0.1	2	0.0	
AIJDB object	0	0	0.0	1	0.0	
AIJFB object	2	0	0.8	11	0.0	
RTUPB object	0	0	0.0	0	0.0	
ACTIVE object	0	0	0.0	0	0.0	
CPT object	0	0	0.0	0	0.0	
RCACHE object	0	0	0.0	0	0.0	
CLIENT object	0	0	0.0	0	0.0	
CLTSEQ object	0	0	0.0	0	0.0	
UTILITY object	0	0	0.0	0	0.0	

```
-----
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !
```

SCREEN FIELDS

KROOT object This field displays the database control information that describes all of the other database objects.

FILID object This field displays the storage area information.

- SEQBLK object This field displays the information on the allocation of transaction sequence numbers.
- TSNBLK object This field displays the information on the last committed transaction.
- AIJDB object This field displays the after-image journal control information.
- AIJFB object This field displays the after-image journal information.
- RTUPB object This field displays information on active users.
- ACTIVE object This field displays information on active transactions.
- CPT object This field displays information on the corrupt page table.
- RCACHE object This field displays information on row caches.
- CLIENT object This field displays client-specific information.
- CLTSEQ object This field displays client “sequence” information.
- UTILITY object This field displays Oracle RMU utility information.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN

INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

Objects Refreshed Screen

This screen displays the statistics for the “objects refreshed” collection category. This category shows the number of objects whose information was detected as being stale, so the information was read again from the database root file.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Field)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (objects refreshed)” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:29:11.65
Rate: 1.00 Second        Object (objects refreshed)      Elapsed: 00:00:13.06
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
KROOT object	0	0	0.0	0	0.0	0.0
FILID object	0	0	0.0	0	0.0	0.0
SEQBLK object	0	0	0.0	1	0.0	0.0
TSNBLK object	0	0	0.1	2	0.0	0.0
AIJDB object	0	0	0.0	1	0.0	0.0
AIJFB object	2	0	0.8	11	0.0	0.0
RTUPB object	0	0	0.0	0	0.0	0.0
ACTIVE object	0	0	0.0	0	0.0	0.0
CPT object	0	0	0.0	0	0.0	0.0
RCACHE object	0	0	0.0	0	0.0	0.0
CLIENT object	0	0	0.0	0	0.0	0.0
CLTSEQ object	0	0	0.0	0	0.0	0.0
UTILITY object	0	0	0.0	0	0.0	0.0

```
-----
```

Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !

SCREEN FIELDS

KROOT object This field displays the database control information that describes all of the other database objects.

FILID object This field displays the storage area information.

SEQBLK object This field displays the information on the allocation of transaction sequence numbers.

- TSNBLK object This field displays the information on the last committed transaction.
- AIJDB object This field displays the after-image journal control information.
- AIJFB object This field displays the after-image journal information.
- RTUPB object This field displays information on active users.
- ACTIVE object This field displays information on active transactions.
- CPT object This field displays information on the corrupt page table.
- RCACHE object This field displays information on row caches.
- CLIENT object This field displays client-specific information.
- CLTSEQ object This field displays client “sequence” information.
- UTILITY object This field displays Oracle RMU utility information.

ON - SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

Objects Modified Screen

This screen displays the statistics for the “objects modified” collection category. This category shows the number of objects whose information was modified. Only objects fetched for exclusive access can be modified.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Field)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (objects modified)” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:29:11.65
Rate: 1.00 Second        Object (objects modified)      Elapsed: 00:00:13.06
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
KROOT object	0	0	0.0	0	0.0	0.0
FILID object	0	0	0.0	0	0.0	0.0
SEQBLK object	0	0	0.0	1	0.0	0.0
TSNBLK object	0	0	0.1	2	0.0	0.0
AIJDB object	0	0	0.0	1	0.0	0.0
AIJFB object	2	0	0.8	11	0.0	0.0
RTUPB object	0	0	0.0	0	0.0	0.0
ACTIVE object	0	0	0.0	0	0.0	0.0
CPT object	0	0	0.0	0	0.0	0.0
RCACHE object	0	0	0.0	0	0.0	0.0
CLIENT object	0	0	0.0	0	0.0	0.0
CLTSEQ object	0	0	0.0	0	0.0	0.0
UTILITY object	0	0	0.0	0	0.0	0.0

```
-----
```

Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !

SCREEN FIELDS

KROOT object This field displays the database control information that describes all of the other database objects.

FILID object This field displays the storage area information.

SEQBLK object This field displays the information on the allocation of transaction sequence numbers.

- TSNBLK object This field displays the information on the last committed transaction.
- AIJDB object This field displays the after-image journal control information.
- AIJFB object This field displays the after-image journal information.
- RTUPB object This field displays information on active users.
- ACTIVE object This field displays information on active transactions.
- CPT object This field displays information on the corrupt page table.
- RCACHE object This field displays information on row caches.
- CLIENT object This field displays client-specific information.
- CLTSEQ object This field displays client “sequence” information.
- UTILITY object This field displays Oracle RMU utility information.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN

INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION

OPTIONS

This screen does not have any configuration options.

Objects Written Screen

This screen displays the statistics for the “objects written” collection category. This category shows the number of objects whose information was written back to the database root file.

This screen resides in the “Objects (One Stat Field)” menu.

SCREEN LOCATION

The following is an example of the “Object (objects written)” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:29:11.65
Rate: 1.00 Second        Object (objects written)      Elapsed: 00:00:13.06
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
KROOT object	0	0	0.0	0	0	0.0
FILID object	0	0	0.0	0	0	0.0
SEQBLK object	0	0	0.0	1	1	0.0
TSNBLK object	0	0	0.1	2	2	0.0
AIJDB object	0	0	0.0	1	1	0.0
AIJFB object	2	0	0.8	11	11	0.0
RTUPB object	0	0	0.0	0	0	0.0
ACTIVE object	0	0	0.0	0	0	0.0
CPT object	0	0	0.0	0	0	0.0
RCACHE object	0	0	0.0	0	0	0.0
CLIENT object	0	0	0.0	0	0	0.0
CLTSEQ object	0	0	0.0	0	0	0.0
UTILITY object	0	0	0.0	0	0	0.0

```
-----
```

Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !

SCREEN FIELDS

KROOT object This field displays the database control information that describes all of the other database objects.

FILID object This field displays the storage area information.

SEQBLK object This field displays the information on the allocation of transaction sequence numbers.

- TSNBLK object This field displays the information on the last committed transaction.
- AIJDB object This field displays the after-image journal control information.
- AIJFB object This field displays the after-image journal information.
- RTUPB object This field displays information on active users.
- ACTIVE object This field displays information on active transactions.
- CPT object This field displays information on the corrupt page table.
- RCACHE object This field displays information on row caches.
- CLIENT object This field displays client-specific information.
- CLTSEQ object This field displays client “sequence” information.
- UTILITY object This field displays Oracle RMU utility information.

ON - SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

Objects Released Screen

This screen displays the statistics for the “objects released” collection category. This category shows the number of objects whose shared or exclusive access was released to other processes.

SCREEN LOCATION

This screen resides in the “Objects (One Stat Field)” menu.

SCREEN EXAMPLE

The following is an example of the “Object (objects released)” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  6-FEB-1998 17:29:11.65
Rate: 1.00 Second        Object (objects released)      Elapsed: 00:00:13.06
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
```

statistic..... name.....	rate.per.second.....			total..... count.....	ave rage..... per.trans....	
	max.....	cur.....	avg.....			
KROOT object	0	0	0.0	0	0.0	0.0
FILID object	0	0	0.0	0	0.0	0.0
SEQBLK object	0	0	0.0	1	0.0	0.0
TSNBLK object	0	0	0.1	2	0.0	0.0
AIJDB object	0	0	0.0	1	0.0	0.0
AIJFB object	2	0	0.8	11	0.0	0.0
RTUPB object	0	0	0.0	0	0.0	0.0
ACTIVE object	0	0	0.0	0	0.0	0.0
CPT object	0	0	0.0	0	0.0	0.0
RCACHE object	0	0	0.0	0	0.0	0.0
CLIENT object	0	0	0.0	0	0.0	0.0
CLTSEQ object	0	0	0.0	0	0.0	0.0
UTILITY object	0	0	0.0	0	0.0	0.0

```
-----
```

Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Write X_plot Yank !

SCREEN FIELDS

KROOT object This field displays the database control information that describes all of the other database objects.

FILID object This field displays the storage area information.

SEQBLK object This field displays the information on the allocation of transaction sequence numbers.

- TSNBLK object This field displays the information on the last committed transaction.
- AIJDB object This field displays the after-image journal control information.
- AIJFB object This field displays the after-image journal information.
- RTUPB object This field displays information on active users.
- ACTIVE object This field displays information on active transactions.
- CPT object This field displays information on the corrupt page table.
- RCACHE object This field displays information on row caches.
- CLIENT object This field displays client-specific information.
- CLTSEQ object This field displays client “sequence” information.
- UTILITY object This field displays Oracle RMU utility information.

ON - SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Graph This on-screen menu option displays a histogram graph instead of numbers.
- Numbers This on-screen menu option displays numeric statistics instead of histogram graph.
- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Pause This on-screen menu option temporarily pauses the presentation of the statistic information. Pressing the “Pause” on-screen menu option will resume the statistic presentation.
- Reset This on-screen menu option resets temporarily resets the statistic information. Use the “Unreset” on-screen menu option to restore the original statistic information.
- Time_plot This on-screen menu option plots a specific field's value by time.
- Unreset This on-screen menu option restores the original statistic values after using “Reset” option.
- X_plot This on-screen menu option plots a specific field's value using scatter-based display.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is available during replay of a binary file specified by the INPUT qualifier.

Cluster Statistic Collection Support This screen is integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

This screen does not have any configuration options.

Chapter

21

Database Dashboard Screens

RMU Show Statistic Database Dashboard Screens

The “Database Dashboard” menu screens display database dashboard information.

IO Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_BUFFERS logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the "IO Dashboard" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  7-FEB-1998 06:56:26.54
Rate: 1.00 Second          IO Dashboard                Elapsed: 00:00:41.43
Page: 1 of 1      DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Buffer Count	20	20	20	20	20	0
APF Enabled	1	1	1	1	1	0
APF Depth	5	5	5	5	5	0
DAPF Enabled	1	1	1	1	1	0
DAPF Depth Count	4	4	4	4	4	0
DAPF Start Count	4	4	4	4	4	0
ABW Enabled	1	1	1	1	1	0
ABW Clean BufCount	5	5	5	5	5	0
ABW Batch Max	4	4	4	4	4	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- Buffer Count** This field displays the default buffer count used by database processes.
- APF Enabled** This field indicates whether asynchronous pre-fetch operations are enabled. The default value 1 indicates that asynchronous pre-fetch operations are enabled and the value "0" indicates that they are disabled. You can override the default value with the RDM\$BIND_APF_ENABLED logical name.
- APF Depth** This field displays the asynchronous pre-fetch depth setting. You can override the default value of 5 buffers with the RDM\$BIND_APF_DEPTH logical name.
- DAPF Enabled** This field indicates whether detected asynchronous pre-fetch operations are enabled. The default value 1 indicates that detected asynchronous pre-fetch operations are enabled and the value "0" indicates that they are disabled. You can override the default value with the RDM\$BIND_DAPF_ENABLED logical name.
- DAPF Depth Count** This field displays the detected asynchronous pre-fetch depth count setting.
- DAPF Start Count** This field displays the detected asynchronous pre-fetch start count.

ABW Enabled This field indicates whether asynchronous batch write operations are enabled. The default value 1 indicates that asynchronous batch write operations are enabled and the value “0” indicates that they are disabled. You can override the default value with the RDM\$BIND_ABW_ENABLED logical name.

ABW Clean BufCount This field displays the asynchronous batch write clean buffer count setting. You can override the default value of 5 buffers with the RDM\$BIND_CLEAN_BUF_CNT logical name.

ABW Batch Max This field displays the asynchronous batch write maximum batch size setting.


ON-SCREEN MENU
 OPTIONS


This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
 INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change

This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change

This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes

Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Locking Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the `RDM$BIND_LOCK_TIMEOUT_INTERVAL` logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the `OPTIONS=UPDATE` qualifier and you need OpenVMS `WORLD`, `BYPASS`, and `SYSNAM` privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the “Locking Dashboard” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:18.67
Rate: 1.00 Second        Locking Dashboard           Elapsed: 00:00:14.52
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Lock Timeout Intvl	2147483647	2147483647	2147483647	2147483647	2147483647	0
Ready AreaSerially	0	0	0	0	0	0
Snap Quiet Point	1	1	1	1	1	0
Hold Retrvl Locks	0	0	0	0	0	0
Coarse Buf Lockng	0	0	0	0	0	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS


- Lock Timeout Intvl** This field displays the number of seconds for a process to wait during a lock conflict before timing out. The default value is 2147483647 seconds. You can override the default value with the RDM\$BIND_LOCK_TIMEOUT_INTERVAL logical name.
- Ready AreaSerially** This field indicates whether Oracle Rdb grants lock requests for logical and physical areas in the order that the lock requests were made. The default value “0” indicates that lock requests are not granted serially. You can override the default value with the RDM\$BIND_READY_AREA_SERIALLY logical name.
- Snap Quiet Point** This field indicates whether the snapshot locks acquire a quiet-point lock. The default value 1 indicates that a quiet-point lock is acquired and the value “0” indicates that a quiet-point lock is not required. You can override the default value with the RDM\$BIND_SNAP_QUIET_POINT logical name.
- Hold Retrvl Locks** This field indicates whether hold retrieval locks are enabled. The default value “0” indicates that hold retrieval locks are disabled and the value 1 indicates that they are enabled. You can override the default value with the RDM\$BIND_HRL_ENABLED logical name.
- Coarse Buf Lockng** This field indicates whether coarse buffer locking is enabled. The default value “0” indicates that coarse buffer locking is disabled and the value 1 indicates that it is enabled. You can override the default value with the RDM\$BIND_CBL_ENABLED logical name.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Update Allows the updating of selected dashboard items.

SCREEN
INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.
- Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.
- Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

- Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.
- Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.
- Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the

user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

AIJ Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_AIJ_IO_MIN logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

This screen resides in the “Database Dashboard” menu.

SCREEN LOCATION

The following is an example of the "AIJ Dashboard" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:20.77
Rate: 1.00 Second        AIJ Dashboard                Elapsed: 00:00:16.62
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Min IO Blocks	0	0	0	0	0	0
Max IO Blocks	127	127	127	127	127	0
AIJ Stall Interval	50	50	50	50	50	0
Root Stall Intervl	50	50	50	50	50	0
Switch Global Ckpt	1	1	1	1	1	0
Check Control Recs	1	1	1	1	1	0
Cache Shuffle Dsbl	0	0	0	0	0	0
ARB Count	300	300	300	300	300	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- Min IO Blocks** This field displays the minimum after-image journal group commit I/O buffer size, in blocks. The default value is 8 blocks. You can override the default value with the RDM\$BIND_AIJ_IO_MIN logical name.
- Max IO Blocks** This field displays the maximum after-image journal group commit I/O buffer size, in blocks. The default value is 127 blocks. You can override the default value with the RDM\$BIND_AIJ_IO_MAX logical name.
- AIJ Stall Interval** This field displays the after-image journal group commit stall time, in milliseconds. The stall time permits a larger number of transactions in the group commit operation. You can override the default value with the RDM\$BIND_AIJ_STALL logical name.
- Root Stall Intervl** This field displays the TSNBLK group commit stall time, in milliseconds. You can override the default value with the RDM\$BIND_COMMIT_STALL logical name.
- Switch Global Ckpt** This field indicates whether to perform a global checkpoint after an after-image journal switch-over has occurred. The default value 1 indicates that a global checkpoint will be performed and the value "0" indicates that a global checkpoint will not be performed. You can override the default value with the RDM\$BIND_AIJ_SWITCH_GLOBAL_CKPT logical name.


- Check Control Recs** This field indicates whether to check for control records during after-image journal cache formatting. The default value “1” indicates that Oracle Rdb will check for control records and the value “0” indicates that Oracle Rdb will not check for control records. You can override the default value with the RDM\$BIND_AIJ_CHECK_CONTROL_RECS logical name.
- Cache Shuffle Dsbl** This field indicates whether or not the group commit buffer “cache shuffle” mechanism is disabled. The default value “0” indicates that the cache shuffle mechanism is enabled, and the value “1” indicates that the cache shuffle mechanism is disabled. You can override the default value with the RDM\$BIND_AIJ_SHUFFLE_DISABLED logical name.
- ARB Count** This field displays the number of after-image journal request blocks (ARBs) that are in use for the database. Changing this field has no effect since it is evaluated only when the database is initially opened.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Options** This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Update** Allows the updating of selected dashboard items.

SCREEN
 INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

- Page Navigation** This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support** This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.
- Cluster Statistic Collection Support** This screen is not integrated into the “Cluster Statistic Collection” facility.
- Zoom Screen** This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change

This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change

This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes

Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Checkpoint Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_CKPT_BLOCKS logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the "Checkpoint Dashboard" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:21.82
Rate: 1.00 Second        Checkpoint Dashboard           Elapsed: 00:00:17.67
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Ckpt Blocks	256	256	256	256	256	0
Ckpt Time Interval	600	600	600	600	600	0
Ckpt Tx Interval	2147483647	2147483647	2147483647	2147483647	2147483647	0
CTJ Tx Interval	256	256	256	256	256	0

Config Exit Help Menu Options Set_rate Update Write !

SCREEN FIELDS

- Ckpt Blocks** This field indicates the number of after-image journal blocks after which a checkpoint will occur. The default value is 0 blocks. You can override the default value with the RDMSBIND_CKPT_BLOCKS logical name.
- Ckpt Time Interval** This field indicates the amount of time, in seconds, after which a checkpoint will occur. The default value is 0. You can override the default value with the RDMSBIND_CKPT_TIME logical name.
- Ckpt Tx Interval** This field indicates the number of committed transactions after which a checkpoint will occur. By default, if you have fast commit processing enabled, a process checkpoints when the after-image journal block size limit is reached or the time interval limit is exceeded, whichever occurs first. You can specify the number of transactions as the checkpoint trigger by defining the RDMSBIND_CKPT_TRANS_INTERVAL logical name.
- CTJ Tx Interval** This field indicates the number of transactions to be allocated as a single batch when the "Commit To Journal optimization" (CTJ) feature is enabled. You can specify the number of transactions to be allocated by defining the RDMSBIND_TSN_INTERVAL logical name.


ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.



 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.



You can use the Config menu option, by typing the letter "C", to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used.

The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Hot Standby Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the `RDM$BIND_HOT_NETWORK_TIMEOUT` logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the `OPTIONS=UPDATE` qualifier and you need OpenVMS `WORLD`, `BYPASS`, and `SYSNAM` privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the “Hot Standby Dashboard” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:23.92
Rate: 1.00 Second        Hot Standby Dashboard           Elapsed: 00:00:19.77
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Network Timeout	120	120	120	120	120	0
Connect Timeout	5	5	5	5	5	0
Sync. Commit Freq	0	0	0	0	0	0
Data Sync Mode	0	0	0	0	0	0
Server Checkpoint	100	100	100	100	100	0
Gap Timeout	5	5	5	5	5	0
Commit Queue Min	0	0	0	0	0	0
Commit Queue Cur	0	0	0	0	0	0
Commit Queue Max	0	0	0	0	0	0
Governor Enabled	1	1	1	1	1	0
Suspend ABS	0	0	0	0	0	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- Network Timeout This field displays the amount of time, in seconds, after which the Hot Standby network times out. The default is 120 seconds. You can override the default with the RDM\$BIND_HOT_NETWORK_TIMEOUT logical name.
- Connect Timeout This field displays the amount of time, in minutes, to wait for the connection to be made between the master and the standby database. The default is 5 minutes. You can override the default with the RDM\$BIND_LCS_CONNECT_TIMEOUT logical name.
- Sync. Commit Freq¹ This field displays the Log Catch-up Server (LCS) process message synchronization count, expressed as the number of messages after which the LCS waits for the Log Rollforward Server (LRS) process on the standby database to finish processing all previously transmitted messages. The default is 64 messages. You can override the default value with the RDM\$BIND_LCS_SYNC_COMMIT_MAX logical name.
- Data Sync Mode This field displays the current database synchronization mode.

¹ Available in Oracle Rdb 7.0.2.0.

<u>Value</u>	<u>Description</u>
0	cold
1	warm
2	hot
3	commit

You can override the default with the `RDMSBIND_HOT_DATA_SYNC_MODE` logical name.

Server Checkpoint This field displays the number of messages per server checkpoint interval. The default is 100 messages. You can override the default with the `RDMSBIND_HOT_CHECKPOINT` logical name.

Gap Timeout This field displays the amount of time, in minutes, to wait for stalled message sequence number (MSN) resolution. The default is 5 minutes. You can override the default value with the `RDMSBIND_LRS_GAP_TIMEOUT` logical name.

Commit Queue Min² This field displays the minimum number of committed transactions the Log Rollforward Server (LRS) process on the standby database will process in a single batch. Decreasing this value may cause asynchronous pre-fetch stalls; increasing this value may cause buffer turn-over. The default is 10 transactions. You can override the default value with the `RDMSBIND_LRS_COMMIT_QUEUE_MIN` logical name.

Commit Queue Cur³ This field displays the current number of committed transactions the Log Rollforward Server (LRS) process on the standby database will process in a single batch. This value dynamically floats between the minimum and maximum dashboard values. The default is 10 transactions. You can override the default value with the `RDMSBIND_LRS_COMMIT_QUEUE_CUR` logical name.

Commit Queue Max⁴ This field displays the maximum number of committed transactions the Log Rollforward Server (LRS) process on the standby database will process in a single batch. Decreasing this value may cause asynchronous pre-fetch stalls; increasing this value may cause buffer turn-over. The default is 500 transactions. You can override the default value with the `RDMSBIND_LRS_COMMIT_QUEUE_MAX` logical name.

² Available in Oracle Rdb 7.0.2.0.

³ Available in Oracle Rdb 7.0.2.0.

⁴ Available in Oracle Rdb 7.0.2.0.

Governor Enabled This field indicates whether the governor is enabled. The default value 1 indicates that the governor is enabled and the value “0” indicates that it is disabled. You can override the default value with the RDM\$BIND_LRS_GOVERNOR_ENABLED logical name.

Suspend ABS This field indicates whether the ABS process is suspended as a result of a Hot Standby (replication) failure. The default value 1 indicates that the ABS process is suspended and the value “0” indicates that the ABS process is not suspended. You can override the default value with the RDM\$BIND_HOT_ABS_SUSPEND logical name.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
INFORMATION



Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.



Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

- Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.
- Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.
- Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Row Cache Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the `RDM$BIND_RCACHE_INSERT_ENABLED` logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the `OPTIONS=UPDATE` qualifier and you need OpenVMS `WORLD`, `BYPASS`, and `SYSNAM` privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the “Row Cache Dashboard” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor   8-FEB-1998 06:59:27.06
Rate: 1.00 Second        Row Cache Dashboard           Elapsed: 00:00:22.91
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1   Mode: Online
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	0	0	0	0	0	0
Insert Enabled	1	1	1	1	1	0
RCRL Count	0	0	0	0	0	0
Latch Spin Count	1024	1024	1024	1024	1024	0

Config Exit Help Menu Options Set_rate Update Write !

SCREEN FIELDS

Insert Enabled This field indicates whether or not rows can be inserted into the row cache. The default value 1 indicates that rows can be inserted into the cache and the value “0” indicates that rows cannot be inserted into the cache. You can override the default value with the RDMSBIND_RCACHE_INSERT_ENABLED logical name.

RCRL Count This field displays the number of reserved row cache slots. You can override the default value of 20 slots with the RDMSBIND_RCACHE_RCRL_COUNT logical name.

Latch Spin Count This field indicates the number of times a process trying to acquire a latch should spin before hibernating.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
INFORMATION



Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.



Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

RUJ Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_RUJ_ALLOC_BLKCNT logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

This screen resides in the “Database Dashboard” menu.

SCREEN LOCATION

The following is an example of the "RUJ Dashboard" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:30.20
Rate: 1.00 Second        RUJ Dashboard                Elapsed: 00:00:26.05
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
```

```
-----
Database..... Current... Previous.. Lowest.... Highest... Original.. Chng
Attribute.Name... Value..... Value..... Value..... Value..... Value..... Cnt.
RUJ Alloc Blocks          127          127          127          127          127          0
RUJ Extend Blocks        127          127          127          127          127          0
```

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

RUJ Alloc Blocks This field indicates the number of blocks the .RUJ file is created with. The default value is 127 blocks. You can override the default value with the RDM\$BIND_RUJ_ALLOC_BLKCNT logical name.

RUJ Extend Blocks This field indicates the number of blocks the .RUJ file is extended by. The default value is 127 blocks. You can override the RDM\$BIND_RUJ_EXTEND_BLKCNT logical name.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
INFORMATION



Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.



Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Monitor Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the `RDM$BIND_RUJ_ALLOC_BLKCNT` logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the `OPTIONS=UPDATE` qualifier and you need OpenVMS `WORLD`, `BYPASS`, and `SYSNAM` privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the "Monitor Dashboard" screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:32.34
Rate: 1.00 Second        Monitor Dashboard           Elapsed: 00:00:28.19
Page: 1 of 1             DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Max DBR Count	50	50	50	50	50	0
ABS Priority	4	4	4	4	4	0
ALS Priority	4	4	4	4	4	0
DBR Priority	4	4	4	4	4	0
LCS Priority	4	4	4	4	4	0
LRS Priority	4	4	4	4	4	0
RCS Priority	4	4	4	4	4	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- Max DBR Count** This field displays the maximum number of database recovery (DBR) processes that will be invoked following node failure. This field can be initialized using the RDM\$BIND_MAX_DBR_COUNT logical name.
- ABS Priority** This field displays the default priority that the detached AIJ Backup Server (ABS) server process will be invoked with. This field can be initialized using the RDM\$BIND_ABS_PRIORITY logical name.
- ALS Priority** This field displays the default priority that the detached AIJ Log Server (ALS) server process will be invoked with. This field can be initialized using the RDM\$BIND_ALS_PRIORITY logical name.
- DBR Priority** This field displays the default priority that the detached database recovery (DBR) server process will be invoked with. This field can be initialized using the RDM\$BIND_DBR_PRIORITY logical name.
- LCS Priority** This field displays the default priority that the detached AIJ Log Catch-up Server (LCS) server process will be invoked with. This field can be initialized using the RDM\$BIND_LCS_PRIORITY logical name.
- LRS Priority** This field displays the default priority that the detached AIJ Log Rollforward Server (LRS) server process will be invoked with. This field can be initialized using the RDM\$BIND_LRS_PRIORITY logical name.

RCS Priority This field displays the default priority that the detached Row Cache Server (RCS) process will be invoked with. This field can be initialized using the RDM\$BIND_RCS_PRIORITY logical name


ON-SCREEN MENU
OPTIONS


This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

ABS Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the `RDM$BIND_ABS_QUIET_POINT` logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “**U**”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the `OPTIONS=UPDATE` qualifier and you need OpenVMS `WORLD`, `BYPASS`, and `SYSNAM` privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

SCREEN EXAMPLE

The following is an example of the “ABS Dashboard” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:35.57
Rate: 1.00 Second        ABS Dashboard                Elapsed: 00:00:31.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Quiet Point	0	0	0	0	0	0
Overwrite Allowed	0	0	0	0	0	0
OverwriteImmediate	0	0	0	0	0	0

```
-----
```

Config Exit Help Menu Options Set_rate Update Write !

SCREEN FIELDS

Quiet Point This field indicates whether the After-image journal Backup Server (ABS) will perform a quiet-point after-image journal backup.

The default value “0” indicates that a no-quiet-point backup will be performed while 1 indicates that a quiet-point backup will be performed. You can override the default value with the RDM\$BIND_ABS_QUIET_POINT logical name.

Overwrite Allowed This field indicates whether the after-image journal backup server (ABS) is allowed to reset overwritten after-image journals.

The default value “0” indicates that the ABS cannot reset overwritten journals while the value 1 indicates that the ABS can reset overwritten journals. You can override the default value with the RDM\$BIND_ABS_OVERWRITE_ALLOWED logical name.

OverwriteImmediate This field indicates whether journals should be immediately reset if RDM\$BIND_ABS_OVERWRITE_ALLOWED logical name is enabled.

The default value “0” indicates that after-image journals should not be immediately reset. The value 1 indicates that after-image journals should be immediately reset depending on the value of the RDM\$BIND_ABS_OVERWRITE_ALLOWED logical name.


You can override the default value with the RDM\$BIND_ABS_OVERWRITE_IMMEDIATE logical name.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Update Allows the updating of selected dashboard items.

SCREEN
INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

- Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.
- Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.
- Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

- Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.
- Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.
- Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the

user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

ALS Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_AIJ_SWITCH_GLOBAL_CKPT logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the “ALS Dashboard” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:37.65
Rate: 1.00 Second        ALS Dashboard                Elapsed: 00:00:33.50
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
AIJ Hiber Time	50	50	50	50	50	0
Switch Global Ckpt	1	1	1	1	1	0
Check Control Recs	1	1	1	1	1	0
Emergency AIJ	0	0	0	0	0	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- AIJ Hiber Time** This field displays the amount of time (in hundredths of a second) that processes have spent hibernating while the AIJ Log Server (ALS) processes their requests.
- Switch Global Ckpt** This field indicates whether to perform a global checkpoint after an after-image journal switch-over has occurred. The default value 1 indicates that a global checkpoint will be performed and the value “0” indicates that a global checkpoint will not be performed. You can override the default value with the RDM\$BIND_AIJ_SWITCH_GLOBAL_CKPT logical name.
- Check Control Recs** This field indicates whether to check for control records during after-image journal cache formatting. The default value 1 indicates that Oracle Rdb will check for control records and the value “0” indicates that Oracle Rdb will not check for control records. You can override the default value with the RDM\$BIND_AIJ_CHECK_CONTROL_RECS logical name.
- Emergency AIJ** This field indicates whether the ALS process creates an emergency after-image journal if the after-image journal switch-over operation enters the suspended state. The default value “0” indicates that the creation of emergency journals is disabled and the value 1 indicates that it is enabled. You can override the default value with the RDM\$BIND_ALS_CREATE_AIJ logical name.


ON-SCREEN MENU
 OPTIONS


This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
 INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
 OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

DBR Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_DBR_BUFFER_CNT logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

SCREEN EXAMPLE

The following is an example of the “DBR Dashboard” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:38.70
Rate: 1.00 Second        DBR Dashboard                Elapsed: 00:00:34.55
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
Database..... Current... Previous.. Lowest.... Highest... Original.. Chng
Attribute.Name.... Value..... Value..... Value..... Value..... Value..... Cnt.
Buffer Count                20          20          20          20          20          0
```

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS


Buffer Count This field displays the default buffer count used by database processes. This field can be initialized using the RDMSBIND_DBR_BUFFER_CNT logical name.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
 INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

C O N F I G U R A T I O N
O P T I O N S

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

RCS Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the `RDM$BIND_RCS_CKPT_BUFFER_CNT` logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the `OPTIONS=UPDATE` qualifier and you need OpenVMS `WORLD`, `BYPASS`, and `SYSNAM` privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

SCREEN EXAMPLE

The following is an example of the "RCS Dashboard" screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 06:59:40.89
Rate: 1.00 Second        RCS Dashboard                Elapsed: 00:00:36.74
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
Ckpt Buffer Count	15	15	15	15	15	0
Batch Count	3000	3000	3000	3000	3000	0
Checkpoint	0	0	0	0	0	0
Ckpt Time Interval	600	600	600	600	600	0
Sweep Interval	1	1	1	1	1	0
Low Cold Thshld	1	1	1	1	1	0
High Cold Thshld	1000	1000	1000	1000	1000	0
Cold Record Count	0	0	0	0	0	0
Abort Sweep	0	0	0	0	0	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- Ckpt Buffer Count** This field indicates the number of buffers to be examined as a single batch by the row cache server (RCS) process during a checkpoint operation. You can override the default value of 15 buffers with the RDM\$BIND_RCS_CKPT_BUFFER_CNT logical name.
- Batch Count** This field displays the number of rows to be flushed to disk.
- Checkpoint** This field indicates whether the row cache server (RCS) performs a checkpoint. The default value 1 indicates a checkpoint is performed and the value "0" indicates that a checkpoint is not performed. You can override the default value with the RDM\$BIND_RCS_CHECKPOINT logical name.
- Ckpt Time Interval** This field indicates the amount of time, in seconds, after which a checkpoint will occur. The default value is 0. You can override the default value with the RDM\$BIND_CKPT_TIME logical name.
- Sweep Interval** This field indicates the amount of time, in minutes between sweeps. The default sweep interval is 1 minute. You can override the default value with the RDM\$BIND_RCS_SWEEP_INTERVAL logical name.


- Low Cold Thshld** This field indicates the number of unmarked records below which the row cache server (RCS) sweep completes. You can override the default value with the RDM\$BIND_RCS_MIN_COLD logical name.
- High Cold Thshld** This field indicates the number of marked records above which the row cache server (RCS) sweep starts. You can override the default value with the RDM\$BIND_RCS_MAX_COLD logical name.
- Cold Record Count** This field displays the number of unmodified rows in the row cache.
- Abort Sweep** This field indicates whether the sweep should be aborted. The default value "0" indicates that the sweep should be continued and the value 1 indicates that the sweep should be aborted.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Options** This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Update** Allows the updating of selected dashboard items.

SCREEN
 INFORMATION

 Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.

 Database attributes are updated on the current node only.

- Page Navigation** This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.
- Binary File Support** This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.
- Cluster Statistic Collection Support** This screen is not integrated into the "Cluster Statistic Collection" facility.
- Zoom Screen** This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

- Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.
- Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.
- Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Per-Process IO Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_BUFFERS logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

SCREEN EXAMPLE

The following is an example of the "Per-Process I/O Dashboard" screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor   8-FEB-1998 07:00:33.99
Rate: 1.00 Second        Per-Process I/O Dashboard      Elapsed: 00:01:29.84
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
For Process: 38211CB8:1 - RICK7						
Buffer Count	500	500	500	500	500	0
APF Enabled	1	1	1	1	1	0
APF Depth	5	5	5	5	5	0
DAPF Enabled	1	1	1	1	1	0
DAPF Depth Count	4	4	4	4	4	0
DAPF Start Count	4	4	4	4	4	0
ABW Enabled	1	1	1	1	1	0
ABW Clean BufCount	5	5	5	5	5	0
ABW Batch Max	4	4	4	4	4	0
Lock Timeout Intvl	2147483647	2147483647	2147483647	2147483647	2147483647	0
Ready AreaSerially	0	0	0	0	0	0
Snap Quiet Point	1	1	1	1	1	0
Hold Retravl Locks	0	0	0	0	0	0
Coarse Buf Lockng	0	0	0	0	0	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- Buffer Count** This field displays the actual buffer count used by database processes. Updating this field has no effect on the process since you cannot dynamically change the number of buffers.
- APF Enabled** This field indicates whether asynchronous pre-fetch operations are enabled. The default value 1 indicates that asynchronous pre-fetch operations are enabled and the value "0" indicates that they are disabled. You can override the default value with the RDMSBIND_APF_ENABLED logical name.
- APF Depth** This field displays the asynchronous pre-fetch depth setting. You can override the default value of 5 buffers with the RDMSBIND_APF_DEPTH logical name.
- DAPF Enabled** This field indicates whether detected asynchronous pre-fetch operations are enabled. The default value 1 indicates that detected asynchronous pre-fetch operations are enabled and the value "0" indicates that they are disabled. You can override the default value with the RDMSBIND_DAPF_ENABLED logical name.

- DAPF Depth Count This field displays the detected asynchronous pre-fetch depth setting.
- DAPF Start Count This field displays the detected asynchronous pre-fetch start count.
- ABW Enabled This field indicates whether asynchronous batch write operations are enabled. The default value 1 indicates that asynchronous batch write operations are enabled and the value "0" indicates that they are disabled. You can override the default value with the RDM\$BIND_ABW_ENABLED logical name.
- ABW Clean BufCount This field displays the asynchronous batch write clean buffer count setting. You can override the default value of 5 buffers with the RDM\$BIND_CLEAN_BUF_CNT logical name.
- ABW Batch Max This field displays the asynchronous batch write maximum batch size setting.
- Lock Timeout Intvl This field displays the number of seconds for a process to wait during a lock conflict before timing out. The default value is 2147483647 seconds. You can override the default value with the RDM\$BIND_LOCK_TIMEOUT_INTERVAL logical name.
- Ready AreaSerially This field indicates whether Oracle Rdb grants lock requests for logical and physical areas in the order that the lock requests were made. The default value "0" indicates that lock requests are not granted serially. You can override the default value with the RDM\$BIND_READY_AREA_SERIALLY logical name.
- Snap Quiet Point This field indicates whether the snapshot locks acquire a quiet-point lock. The default value 1 indicates that a quiet-point lock is acquired and the value "0" indicates that a quiet-point lock is not required. You can override the default value with the RDM\$BIND_SNAP_QUIET_POINT logical name.
- Hold Retrvl Locks This field indicates whether hold retrieval locks are enabled. The default value "0" indicates that hold retrieval locks are disabled and the value 1 indicates that they are enabled. You can override the default value with the RDM\$BIND_HRL_ENABLED logical name.
- Coarse Buf Lockng This field indicates whether coarse buffer locking is enabled. The default value "0" indicates that coarse buffer locking is disabled and the value 1 indicates that it is enabled. You can override the default value with the RDM\$BIND_CBL_ENABLED logical name.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
INFORMATION



Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.



Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Per-Process Journal Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_BUFFERS logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

SCREEN EXAMPLE

The following is an example of the “Per-Process Journal Dashboard” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor   8-FEB-1998 07:00:42.52
Rate: 1.00 Second        Per-Process Journal Dashboard   Elapsed: 00:01:38.37
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1   Mode: Online
```

```
-----
```

Database.....	Current...	Previous..	Lowest....	Highest...	Original..	Chng
Attribute.Name....	Value.....	Value.....	Value.....	Value.....	Value.....	Cnt.
For Process: 38211CB8:1 - RICK7						
Min AIJ IO Bytes	0	0	0	0	0	0
Max AIJ IO Bytes	65024	65024	65024	65024	65024	0
AIJ Stall Interval	50	50	50	50	50	0
Root Stall Intervl	50	50	50	50	50	0
Switch Global Ckpt	1	1	1	1	1	0
Check Control Recs	1	1	1	1	1	0
Ckpt Blocks	256	256	256	256	256	0
Ckpt Time Interval	600	600	600	600	600	0
Ckpt Tx Interval	2147483647	2147483647	2147483647	2147483647	2147483647	0
CTJ Tx Interval	256	256	256	256	256	0
RUJ Alloc Blocks	127	127	127	127	127	0
RUJ Extend Blocks	127	127	127	127	127	0

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

- Min AIJ IO Bytes** This field displays the minimum after-image journal group commit I/O buffer size, in bytes. The default value is 4096 bytes. You can override the default value with the RDM\$BIND_AIJ_IO_MIN logical name.
- Max AIJ IO Bytes** This field displays the maximum after-image journal group commit I/O buffer size, in bytes. The default value is 65024 blocks. You can override the default value with the RDM\$BIND_AIJ_IO_MAX logical name.
- AIJ Stall Interval** This field displays the after-image journal group commit stall time, in milliseconds. The stall time permits a larger number of transactions in the group commit operation. You can override the default value with the RDM\$BIND_AIJ_STALL logical name.
- Root Stall Intervl** This field displays the TSNBLK group commit stall time, in milliseconds. You can override the default value with the RDM\$BIND_COMMIT_STALL logical name.
- Switch Global Ckpt** This field indicates whether to perform a global checkpoint after an after-image journal switch-over has occurred. The default value 1 indicates that a global checkpoint will be performed and the value “0” indicates that a global checkpoint will not be per-

formed. You can override the default value with the RDM\$BIND_AIJ_SWITCH_GLOBAL_CKPT logical name.

- Check Control Recs** This field indicates whether to check for control records during after-image journal cache formatting. The default value 1 indicates that Oracle Rdb will check for control records and the value "0" indicates that Oracle Rdb will not check for control records. You can override the default value with the RDM\$BIND_AIJ_CHECK_CONTROL_RECS logical name.
- Ckpt Blocks** This field indicates the number of after-image journal blocks after which a checkpoint will occur. The default value is 0 blocks. You can override the default value with the RDM\$BIND_CKPT_BLOCKS logical name.
- Ckpt Time Interval** This field indicates the amount of time, in seconds, after which a checkpoint will occur. The default value is 0. You can override the default value with the RDM\$BIND_CKPT_TIME logical name.
- Ckpt Tx Interval** This field indicates the number of committed transactions after which a checkpoint will occur. By default, if you have fast commit processing enabled, a process checkpoints when the after-image journal block size limit is reached or the time interval limit is exceeded, whichever occurs first. You can specify the number of transactions as the checkpoint trigger by defining the RDM\$BIND_CKPT_TRANS_INTERVAL logical name.
- CTJ TX Interval** This field indicates the number of transactions to be allocated as a single batch when the "Commit To Journal optimization" feature is enabled. You can specify the number of transactions to be allocated by defining the RDM\$BIND_TSN_INTERVAL logical name.
- RUJ Alloc Blocks** This field indicates the number of blocks the .RUJ file is created with. The default value is 127 blocks. You can override the default value with the RDM\$BIND_RUJ_ALLOC_BLKCNT logical name.
- RUJ Extend Blocks** This field indicates the number of blocks the .RUJ file is extended by. The default value is 127 blocks. You can override the default value with the RDM\$BIND_RUJ_EXTEND_BLKCNT logical name.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

- Options** This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.
- Update** Allows the updating of selected dashboard items.

SCREEN
 INFORMATION



Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.



Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
 OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Per-Process Row Cache Dashboard Screen

This screen displays the actual database parameter and attribute settings being used by the processes attached to the database. Optionally, qualified users are allowed to dynamically update certain database parameters and attributes on a single node at run-time. These changes can be either global to all processes on the node, or local to specific processes. The net effect of these changes can be examined at run time without having to restart database processes. These updates are non-persistent; the changes are lost when the database is closed on that node.

For example, when the RDM\$BIND_BUFFERS logical name is defined at the process level, there is no method available from the operating system for the database administrator (DBA) to examine its setting at runtime for a specific process. The same is true for any database logical name and the majority of interesting database attributes that affect runtime performance of the database. However, the Database Dashboard provides this information dynamically.

The Database Dashboard facility allows you to drive the database faster or slower, and immediately see the impact of increasing or decreasing certain database settings.

You can use the “Update” on-screen menu option, by typing the letter “U”, to change the value of a database attribute. Before you can update database attributes, you need to start your RMU Show Statistic utility session with the OPTIONS=UPDATE qualifier and you need OpenVMS WORLD, BYPASS, and SYSNAM privileges.



Caution: You should use the “Update” on-screen menu option carefully. The RMU Show Statistic utility does not perform error checking on the updated values, other than ensuring the values are within specified range criteria. The utility does not validate the effect of a parameter setting on the operation of the database.

SCREEN LOCATION

This screen resides in the “Database Dashboard” menu.

The following is an example of the “Per-Process Row Cache Dashboard” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 07:00:49.09
Rate: 1.00 Second        Per-Process Row Cache Dashboard  Elapsed: 00:01:44.94
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
```

```
-----
Database..... Current... Previous.. Lowest.... Highest... Original.. Chng
Attribute.Name.... Value..... Value..... Value..... Value..... Value..... Cnt.
For Process: 38211CB8:1 - RICK7
Insert Enabled          1          1          1          1          1          0
RCRL Count              0          0          0          0          0          0
Latch Spin Count       1024       1024       1024       1024       1024       0
```

```
-----
Config Exit Help Menu Options Set_rate Update Write !
```

SCREEN FIELDS

Insert Enabled This field indicates whether or not rows can be inserted into the row cache. The default value 1 indicates that rows can be inserted into the cache and the value “0” indicates that rows cannot be inserted into the cache. You can override the default value with the RDMSBIND_RCACHE_INSERT_ENABLED logical name.

RCRL Count This field displays the number of reserved row cache slots. You can override the default value of 20 slots with the RDMSBIND_RCACHE_RCRL_COUNT logical name.

Latch Spin Count This field indicates the number of times a process trying to acquire a latch should spin before hibernating.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

Update Allows the updating of selected dashboard items.

SCREEN
 INFORMATION



Note that updates made to any attributes are not stored in the database rootfile. The purpose of updating attributes is to test and measure the effects of changes on the database, so that you can later make persistent changes to appropriate database attributes using interactive SQL.



Database attributes are updated on the current node only.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
 OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Use active notification of change This configuration option indicates that the changes to the database attributes and parameters should be actively and immediately broadcast to all users on the current node.

Use passive notification of change This configuration option indicates that the changes to the database attributes and parameters should be passively and lazily broadcast to all users on the current node.

Notify users of previous changes Active notification means that all processes attached to the database on the node will be immediately notified that a database attribute setting has been modified and that they should reset their individual parameter values. Passive notification means that the user processes will be notified passively at intervals determined by the database software. Passive notification is the default, as it is non-intrusive to system performance.

The advantage of active notification is that each database user on the node is immediately notified of the change and that the new parameter settings are immediately used. The disadvantage is that system performance may be temporarily impacted while the users respond to the notification.

The advantage of passive notification is that there is no impact on the system. The disadvantage is that you have no control over when the individual processes will be notified that a particular database parameter setting has changed.

When using passive notification, a broadcast notification can be performed using the “Notify users of previous changes” option of the “Config” sub-menu.

Chapter

22

Online Analysis & Information Screens

RMU Show Statistic Online Analysis & Information Screens

The “Online Analysis & Information” menu screens display online analysis information related to potential statistic areas of interest.

Buffer Analysis Screen

This screen provides information about local and global buffer performance. You can use the information in this screen to examine the effectiveness of local and global buffers.

SCREEN LOCATION

This screen resides in the "Online Analysis & Info." menu.

SCREEN EXAMPLE

The following is an example of the "Buffer Analysis" screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second          Buffer Analysis                Elapsed: 00:00:13.42
Page: 1 of 1      DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

nchronous pre-fetch is disabled

This message indicates that the Asynchronous Pre-Fetch (APF) feature is not available for your database. The APF feature reduces the amount of time that a process waits for page retrievals from the storage area by anticipating which pages will need to be available next. These anticipated pages can then be retrieved before the application actually requests the pages.

When the APF feature is available, Oracle Rdb examines each process' operations and attempts to predict the process' future access patterns; the APF feature identifies sequential scans. When Oracle Rdb detects a sequential scan being performed, it fetches the pages that it anticipates the application will subsequently request. This pre-fetching of pages is asynchronous.

The APF feature is available by default. You can also use the **RDM\$BIND_APF_ENABLED** logical to enable or disable the APF feature. The value "0" will disable the feature and the value "1" will enable the feature.

Oracle recommends that you use the SQL syntax to enable or disable the APF feature.

ected asynchronous pre-fetch is disabled

This message indicates that the Detected Asynchronous Pre-Fetch (DAPF) feature is not available for your database. The DAPF feature reduces the amount of time that a process waits for page retrievals from the storage area by anticipating which pages will need to be available next. These anticipated pages are then subsequently retrieved before the application actually requests the pages.

When the DAPF feature is available, Oracle Rdb examines the operations of each process and attempts to predict the future access patterns of the application. The DAPF feature is able to detect more than just the identification of sequential scans. It applies artificial intelligence heuristics to predict access patterns, even within line-item-type clusters. This pre-fetching of pages is asynchronous.

The DAPF feature is available by default. You can also use the **RDMSBIND_DAPF_ENABLED** logical to enable or disable the DAPF feature. The value "0" will disable the feature and the value "1" will enable the feature.

Oracle recommends that you use the SQL syntax to enable or disable the DAPF feature.

nchronous batch-write is disabled

This message indicates that the Asynchronous Batch-Write (ABW) feature is not available for your database. The ABW feature reduces the number of stalls experienced by database processes while waiting for write I/O operations to the storage areas to complete.

Earlier versions of the Oracle Rdb product performed page-write operations in batches. Whenever a data page needed to be written back to the storage area, a set of other modified data pages would also be gathered together and written to disk in parallel. However, the process still had to wait until all of the disk writes completed. This is essentially a synchronous write I/O operation, with the result that the process could not perform additional application work until all write operations completed. Also, the impact of these write I/O operations on the disk device and disk controller is often quite intrusive, requiring additional resources to manage the sudden increase in I/O activity.

The goal of the ABW feature is to increase database performance by making it possible for a certain number of buffers in the allocate set of each process to have disk write I/O operations in progress at any time without requiring the process to wait for

any of those write I/O operations to complete. For each individual process, performance is optimal when a constant and sustainable series of disk write I/O operations are being performed. This allows the disk controller to manage I/O queues and disk-head movement in a much more controlled manner, thereby improving overall system performance for all database users.

The ABW feature is very important when using the Fast Commit feature, because checkpoints cause the entire set or portion of modified pages for a process to be written asynchronously back to the database storage areas.

The Fast Commit feature improves database throughput by delaying buffer write I/O operations as long as possible. This improves the re-use of hot pages and eliminates unnecessary write I/O operations. However, one of the major performance problems of the Fast Commit feature is being blasted for one of these modified pages by another user, and especially being blasted by the database recovery process (DBR). Your process essentially stops everything, synchronously writes RUJ, sometimes synchronously writes the SPAM and finally writes the affected data page to disk and releases the page lock. You can avoid this by pre-writing the pages asynchronously, including the RUJ journal, using the ABW feature. However, proper settings are necessary. Oracle recommends that you review the Guide to Database Performance and Tuning for additional information.



Note that it is difficult to recommend an ideal setting because each application is different. In some cases, it is sufficient to maintain your buffer pool 50% at empty, but sometimes 75% may be a better value; you need to decide. A further complication is that the DBR process prevents your process from doing application work due to the database freeze. Therefore, DBR requesting pages sometimes results in your process receiving deadlock on freeze on pages being requested by the DBR process. Even though these deadlocks are internal they still wait the interval specified by the OpenVMS SYSGEN DEADLOCK_WAIT parameter. The DEADLOCK_WAIT SYSGEN parameter defaults to 10 seconds; Oracle highly recommends that you change this value to one or two seconds at most.

The ABW feature is available by default. You can also use the **RDMSE_BIND_ABW_ENABLED** logical to enable or disable the ABW feature. The value "0" will disable the feature and the value "1" will enable the feature.

Oracle recommends that you use the SQL syntax to enable or disable ABW feature.

'AS <n.m>% page hit below <x.y>% threshold

This message indicates that the allocate set effectiveness is less than the corresponding threshold. This message applies to both local and global buffers; when using local buffers, the buffer pool is the allocate set.

The percentage is computed as the total number of pages found in the local buffer or global buffer allocate set, divided by the total number of page requests. In general, a higher percentage implies better buffering and, therefore, fewer stalls and better throughput. If every page request found the page in the buffer pool, the buffer pool effectiveness would be 100%.

However, an allocate set effectiveness of 100% is clearly impossible, since your process attaches to the database with an empty buffer pool. Therefore, at least one page request will result in the page not being found in the buffer pool.

Also, note that when using global buffers, this analysis is measuring the allocate set size instead of the overall global buffer pool effectiveness. It is quite possible that the global buffer pool effectiveness is adequate while the allocate set effectiveness is not. Typically, reducing the number of attached users or increasing the number of buffers improves the allocate set effectiveness.

The default allocate set effectiveness threshold is 75%. You can override the default value with the **RDMSBIND_STATS_LB_PAGE_HIT_RATIO** logical name, the **LB_PAGE_HIT_RATIO** configuration variable or by selecting the LB/AS page hit threshold option from the “Config” on-screen menu.

<n.m>% pool hit below <n.m>% threshold

This message indicates that your global buffer pool effectiveness is less than the corresponding threshold.

The percentage is computed as the total number of pages found in the global buffer pool, divided by the total number of page requests. In general, a higher percentage implies better buffering and, therefore, fewer stalls and better throughput. If every page request found the page in the buffer pool, the buffer pool effectiveness would be 100%.

However, a global buffer pool effectiveness of 100% is clearly impossible, since the monitor opens the database with an empty global buffer. Therefore, at least one page request will result in the page not being found in the buffer pool.

The default local buffer pool effectiveness threshold is 85%. You can override the default value with the **RDMSBIND_STATS_GB_POOL_HIT_RATIO** logical name, the **GB_POOL_HIT_RATIO** configuration variable or by selecting the “GB pool hit threshold” option from the “Config” on-screen menu.

<n.m>% IO-saved below <n.m>% threshold

This message indicates that the percentage of page requests that saved a disk I/O is less than the corresponding threshold. This message is very useful in determining whether local buffers are more efficient than global buffers.

The saved I/O percentage is computed as the total number of pages requested divided by the total number of pages found. In general, a higher percentage implies a better benefit obtained from using global buffers. As the percentage approaches 0%, you may want to consider using local buffers. Conversely, as the percentage approaches 100%, you may want to continue using global buffers.

Adjusting the size of the global buffer affects this value. The larger the global buffer, the higher likelihood of being able to save read I/O operations.

The default local buffer pool effectiveness threshold is 85%. You can override the default value with the **RDMSBIND_STATS_GB_IO_SAVED_RATIO** logical name, the **GB_IO_SAVED_RATIO** configuration variable or by selecting the “GB IO-saved threshold” option from the “Config” on-screen menu.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
 OPTIONS

You can use the Config menu option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

- LB/AS page hit threshold This configuration option allows you to change the default threshold, which is 75%. You can override the default value with the `RDMSBIND_STATS_LB_PAGE_HIT_RATIO` logical name.
- GB pool hit threshold This configuration option allows you to change the default threshold, which is 85%. You can override the default value with the `RDMSBIND_STATS_GB_POOL_HIT_RATIO` logical name.
- GB IO-saved threshold This configuration option allows you to change the default threshold, which is 85%. You can override the default value with the `RDMSBIND_STATS_GB_IO_SAVED_RATIO` logical name.

Transaction Analysis Screen

This screen provides information about transaction performance. Primarily, this screen examines average transaction duration.

SCREEN LOCATION

This screen resides in the “Online Analysis & Info.” menu.

SCREEN EXAMPLE

The following is an example of the “Transaction Analysis” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second        Transaction Analysis           Elapsed: 00:00:13.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

m>% verb rollback above <n.m>% threshold

This message indicates that the SQL verb rollback percentage has exceeded the corresponding threshold.

The percentage is computed by dividing the number of verb failures by the total number of verbs. A SQL verb is an atomic statement or action. For example, a record insert is a verb, as is a record deletion.

Excessive verb failures are usually an indication of a failed constraint, such as uniqueness criteria, or an invalid DDL statement.



Note that in the case of cursors and scans, reaching the end-of-stream always results in a verb failure.

Note that SQL performs its own internal queries to identify metadata, such as relation or index names.

Also, within a compound statement each individual statement is atomic and Oracle Rdb performs a verb-success operation after processing each one. To avoid this overhead, you can use the SQL BEGIN atomic statement to treat the entire block as a single verb.

Oracle Rdb rarely issues a verb-failure unless there is an exception of some kind, such as a constraint failure.

The default verb rollback threshold is 25 percent. You can override the default value with the **RDMSBIND_STATS_VERB_SUCCESS_RATIO** logical name, or by defining the **VERB_SUCCESS_RATIO** configuration variable, or by selecting the “Verb success threshold” option from the “Config” on-screen menu.

>th %ile transaction duration: <n.m> seconds

This message indicates that the nth percentile aggregate transaction duration has exceeded the corresponding threshold. The indicated number of seconds represents the nth percentile time, not the specified threshold. Note that this message is for both read-write and read-only transactions combined; there are separate messages for just read-write and read-only transaction durations.

The duration of each transaction measures the first SQL SET TRANSACTION statement through the completion of the transaction with a COMMIT or ROLLBACK statement. As each transaction completes, its duration increases the cumulative value maintained by the database.

This message is very useful for eliminating the infrequent long-duration transaction from the determination of the typical transaction duration.

The default transaction duration percentile is 95%. You can override the default value with the **RDMSBIND_STATS_MAX_TX_PERCENTILE** logical name, or by defining the **MAX_TX_PERCENTILE** configuration variable, or by selecting the “Transaction duration percentile” option from the “Config” on-screen menu.

The default transaction duration threshold is 15 seconds. You can override the default value with the **RDMSBIND_STATS_MAX_TX_DURATION** logical name, or by defining the **MAX_TX_DURATION** configuration variable, or by selecting the “Transaction duration threshold” option from the “Config” on-screen menu.

>th %ile read/write transaction duration: <n.m> seconds⁵

⁵ Oracle Rdb 7.0.2.0

This message indicates that the n^{th} percentile read/write transaction duration has exceeded the corresponding threshold. The indicated number of seconds represents the n^{th} percentile time, not the specified threshold. Note that this message is for read-write transactions only; there are separate messages for aggregate and read-only transaction durations.

The duration of each transaction measures the first SQL SET TRANSACTION statement through the completion of the transaction with a COMMIT or ROLLBACK statement. As each transaction completes, its duration increases the cumulative value maintained by the database.

This message is very useful for eliminating the infrequent long-duration transaction from the determination of the typical transaction duration.

The default transaction duration percentile is 95%. You can override the default value with the **RDM\$BIND_STATS_MAX_RW_TX_PERCENTILE** logical name, or by defining the **MAX_RW_TX_PERCENTILE** configuration variable, or by selecting the “Read/Write Transaction duration percentile” option from the “Config” on-screen menu.

The default transaction duration threshold is 15 seconds. You can override the default value with the **RDM\$BIND_STATS_MAX_RW_TX_DURATION** logical name, or by defining the **MAX_RW_TX_DURATION** configuration variable, or by selecting the “Read/Write Transaction duration threshold” option from the “Config” on-screen menu.

n^{th} %ile read-only transaction duration: <n.m> seconds⁶

This message indicates that the n^{th} percentile read-only transaction duration has exceeded the corresponding threshold. The indicated number of seconds represents the n^{th} percentile time, not the specified threshold. Note that this message is for read-only transactions only; there are separate messages for aggregate and read/write transaction durations.

The duration of each transaction measures the first SQL SET TRANSACTION statement through the completion of the transaction with a COMMIT or ROLLBACK statement. As each transaction completes, its duration increases the cumulative value maintained by the database.

This message is very useful for eliminating the infrequent long-duration transaction from the determination of the typical transaction duration.

⁶ Oracle Rdb 7.0.2.0

The default transaction duration percentile is 95%. You can override the default value with the **RDM\$BIND_STATS_MAX_RO_TX_PERCENTILE** logical name, or by defining the **MAX_RO_TX_PERCENTILE** configuration variable, or by selecting the “Read-Only Transaction duration percentile” option from the “Config” on-screen menu.

The default transaction duration threshold is 15 seconds. You can override the default value with the **RDM\$BIND_STATS_MAX_RO_TX_DURATION** logical name, or by defining the **MAX_RO_TX_DURATION** configuration variable, or by selecting the “Read-Only Transaction duration threshold” option from the “Config” on-screen menu.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Verb success threshold This configuration option allows you to change the default threshold, which is 25%. You can override the default value with the **RDM\$BIND_STATS_VERB_SUCCESS_RATIO** logical name.

Transaction duration percentile This configuration option allows you to change the default threshold, which is 95%. You can override the default value with the **RDM\$BIND_STATS_MAX_TX_PERCENTILE** logical name.

Transaction duration threshold This configuration option allows you to change the default threshold, which is 15 seconds. You can override the default value with the **RDM\$BIND_STATS_MAX_TX_DURATION** logical name.

Read-Only Transaction duration percentile

This configuration option allows you to change the default threshold, which is 95%. You can override the default value with the `RDM$BIND_STATS_MAX_RO_TX_PERCENTILE` logical name.

d-Only Transaction duration threshold

This configuration option allows you to change the default threshold, which is 15 seconds. You can override the default value with the `RDM$BIND_STATS_MAX_RO_TX_DURATION` logical name.

AIJ Analysis Screen

This screen provides information about AIJ journal performance. You can use the information in this screen to ensure that AIJ journals are on their own device. You can also examine the AIJ request block to I/O ratio, blocks per I/O ratio, and the formatting of the AIJ request block.

SCREEN LOCATION

This screen resides in the "Online Analysis & Info." menu.

SCREEN EXAMPLE

The following is an example of the "AIJ Analysis" screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second          AIJ Analysis                Elapsed: 00:00:13.42
Page: 1 of 1      DISK$: [DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

AIJ journaling is disabled This message indicates that the AIJ journaling feature is not available. By default, a newly created database or a newly restored database does not use the AIJ journaling facility. When AIJ journaling is not available, the database is non-recoverable in the event of media failure, such as a disk crash.

Using the AIJ journaling facility without the Fast Commit feature usually incurs about a 10% reduction in performance. However, when combined with the Fast Commit feature, a performance *improvement* of up to 50% is often possible. The Commit To Journal optimization feature is able to improve performance even more. These types of features rely on the after-image journal for recovery in the event of process failure.

abase should be configured for at least <n> journals

This message indicates that the database uses either a single extensible AIJ journal, or two fixed-size AIJ journals. Both of these combinations provide less than optimal performance characteristics.

Extensible AIJ journals are excellent when you cannot predict transaction work-loads, or you are short on disk space and cannot afford to pre-allocate multiple fixed-size AIJ journals. Extensible AIJ journals will extend automatically as they become full. However, all transaction activity suspends during the AIJ extension operation, on *all* database nodes.

The AIJ extension operation performs at a maximum rate of 2,000 blocks per second. For example, at 10,000 block extension takes a minimum of five seconds to complete, even under optimal circumstances. This pause can be quite noticeable on OLTP-type applications where the suspension of transactional activity directly impacts throughput.

Fixed-size AIJ journals eliminate the extension problem by providing multiple ready-to-use AIJ journals. When one AIJ journal fills up, you simply switch to another available AIJ journal. However, fixed-size journals must be available in order to be accessible. This means that you must backup the old journals faster than the application can fill them. This is not often possible with just two fixed-size journals. Therefore, Oracle recommends that you always have a minimum of three fixed-size journals and, preferably, at least five. Oracle also always recommends that you use the AIJ Backup Server to automate the AIJ backup operation.

The value <n> is based on the number of reserved AIJ slots.

east <n> journals should be reserved

This message indicates that the number of available AIJ journal slots is less than two. AIJ slots are pre-allocated by the DBA, and determine the maximum number of AIJ journals that can be created for a database.

With zero or one AIJ slots available, your database is at risk of suspending during an AIJ switch-over operation. Oracle recommends that your database always have at least two available AIJ journal slots and, preferably, at least five.

Log server is manual This message indicates that the database Monitor does not start the AIJ Log Server (ALS) unless specifically directed by the DBA. There is no threshold associated with this message.

Oracle recommends that the ALS always be invoked automatically.

When the database Monitor does not invoke the ALS process automatically, application processes need to acquire a local lock for their own updated buffers, as well as the updated buffers of other processes, before doing a group-commit of the updated

buffers to disk. When high update activity is occurring on the database, this lock contention can cause a decrease in performance and system throughput.

Also, when using circular AIJ journals, the ALS process is able to automatically create emergency AIJ journals in the event of an AIJ switch-over suspension.

Log server is automatic This message indicates that the database Monitor starts the AIJ Log Server (ALS) automatically at database open time. There is no threshold associated with this message.

Oracle recommends that the ALS process always be invoked automatically.

Journal overwrite is enabled This message indicates that the AIJ journals can be re-used without having first been backed up. This attribute is only useful for those databases that can be re-created from scratch, as they are non-recoverable.

Oracle recommends that the Journal Overwrite feature always be disabled. Use of the Emergency AIJ Journal facility is a preferable method of ensuring continued application processing in the event of suspended AIJ switch-over operations.

throughput <n.m> excessive for allocation <n> blocks (minimum <m>)

This message indicates that the AIJ throughput will cause the current AIJ journal to extend or switch-over in a time period that is less than the corresponding threshold.

AIJ Request Blocks (ARB) are pieces of information to be written to the AIJ journal. Each ARB is approximately 2k in size but often contains less information than that. In addition, ARBs contain different types of data; for instance, the most common ARB type is data but ARBs can also contain commit information, rollback and checkpoint information, and a host of other interesting types.

Each process initializes an ARB and submits it to be formatted by the group commit process. If you are not using the AIJ Log Server (ALS), one of the application processes becomes the group commit process. If you are using the ALS, then the ALS is always the group commit process. Using the ALS is desirable for several reasons. The first reason is that formatting the ARBs is CPU-intensive and allowing application processes to do this means they are NOT doing application work. The second reason is that the ALS formats ARBs and nothing else; it uses a highly optimized algorithm, using a double-buffering scheme to perform recoverable asynchronous I/O to the AIJ journal.

Note: Oracle Rdb is the only production database that allows asynchronous I/O to the AIJ journal.

The ALS algorithm always attempts to format more ARBs while the previously formatted ARBs are being written to disk. A foreground format occurs when ARBs are formatted while no AIJ journal write I/O operations are pending. A background for-

mat occurs when ARBs are formatted while AIJ journal write I/O operations are pending. Ideally, it is desirable to format all ARBs in background since this results in the best utilization of the ALS process. Remember that the AIJ journal is a sequential file. It is extremely desirable to keep the data streaming to the disk so that disk-head movement is optimal. Requiring the ALS to stop and format ARBs then issue an I/O means that the disk head is not properly positioned for optimum efficiency, with the result of increased latency.

Suppose 8.9% of your ARBs are formatted in the background. This means that 91.1% of the time, you are formatting ARBs while no AIJ journal disk I/O is pending. This results in longer AIJ journal I/O duration, which typically impacts transaction commit operation durations.

What are the causes of this? Typically, this means that your transactions commit intermittently (most frequent reason), or they update relatively little information. Using the Dashboard facility of the RMU Show Statistic utility is the easiest method available for experimenting with various database attribute adjustments. Try adjusting the “Max IO Blocks” dashboard entry lower. Possibly using the Fast Commit feature will help improve this event. The interesting use of this message is in determining whether the expense of using the ALS process is worth the benefit. If you have almost no background ARB formats, you might be better off not using the ALS process.

The default AIJ seconds to extend threshold is 60 seconds. You can override the default value with the `RDMSBIND_STATS_AIJ_SEC_TO_EXTEND` logical name, or by defining the `AIJ_SEC_TO_EXTEND` configuration variable, or by selecting the “Seconds to AIJ extend/switch” option from the “Config” on-screen menu.

AIJ Backup Server is manual This message indicates that the database Monitor will not automatically invoke the AIJ Backup Server (ABS) on an AIJ journal switch-over operation.

Failure to perform a timely backup of the AIJ journals may cause the AIJ journal switch-over operation to suspend. During the suspension, all database transactional activity suspends until an available AIJ journal can be found.

Oracle recommends the automatic invocation of the ABS process when using circular AIJ journals.

cannot be invoked because only 1 journal is active

This message indicates that the database Monitor will invoke the AIJ Backup Server (ABS) automatically. However, the database Monitor is unable to invoke the ABS process because extensible AIJ journals are being used.

The database Monitor only invokes the ABS process following an AIJ journal switch-over operation. Since extensible AIJ journals are unable to switch-over, the database Monitor will not invoke the ABS.

Oracle recommends that you use fixed-size (circular) AIJ journals.

↳ cannot be invoked because a journal has been overwritten

This message indicates that the database Monitor will invoke the AIJ Backup Server (ABS) automatically. However, the database Monitor is unable to invoke the ABS process because one or more of the fixed-size AIJ journals have been over-written. Over-writing the previous contents of the AIJ journal is destructive, and the recovery information cannot be re-obtained.

Oracle recommends that you immediately backup your database as it cannot be recovered in the event of media failure.

↳ cannot be invoked because hard data loss has occurred

This message indicates that the database Monitor invokes the AIJ Backup Server (ABS) automatically. However, the database Monitor is unable to invoke the ABS process because one or more of the fixed-size AIJ journals have experienced hard data loss.

Hard data loss occurs when a process is unable to write some AIJ modifications to any AIJ journal and the information must be discarded. This typically occurs when the disk device on which the AIJ journals reside is not accessible in some manner. The database is always shutdown following hard data loss to preserve integrity. However, the database can be restarted once new AIJ journals have been made available to the system.

Oracle recommends that you immediately backup your database as it cannot be recovered in the event of media failure.

↳ cannot be invoked because backup filename has not been specified

This message indicates that the database Monitor invokes the AIJ Backup Server (ABS) automatically. However, the database Monitor is unable to invoke the ABS process because none of the AIJ journals have a valid backup filename specification.

You can define the AIJ journal backup filename specifications either defined globally for all AIJ journals, or on a per-journal basis.

<AIJ_name> device <device_name> same as storage area

This message indicates that the named AIJ journal resides on a disk device that one or more storage areas share.

Oracle highly recommends that AIJ journals reside on their own disk device. This eliminates a double point-of-failure. For example, if the AIJ journal were to reside on the same device as a storage area, then losing that disk device would jeopardize both the AIJ journal and the storage area; you would be unable to recover the storage area because you also lost the source of the recovery information.

ARB pool exhausted <n> times This message indicates that the AIJ Request Block (ARB) pool has been fully consumed the indicated number of times.

By default, there are 300 ARBs per database per node. All applications users on the particular node share these ARBs. When the application processes consume all of the ARBs, all transaction processing stops until the group commit process formats some ARBs and makes them available to the application processes.

Running out of available ARBs is an extremely serious situation; it should be always avoided, if possible. The most frequent cause of this event is that the number of application processes greatly exceed the number of available ARBs. Another cause is transactions that update a large amount of data at once, such as bulk loads or mass updates.

In Oracle Rdb 7.0.1.0, the RMU Show Statistic utility displays the number of available ARBs on the “AIJ Journal Information” screen. Oracle Rdb7 allows you to increase the maximum ARB value up to 8,192 using the “ARB Count” dashboard entry, located in the “AIJ Dashboard.” Oracle Rdb7 also allows you to increase the maximum ARB value up to 8,192 using the **RDMSBIND_AIJ_ARB_COUNT** logical. There is no configuration variable to control the number of ARBs.

m>% ARBs per I/O below <n.m>% threshold

This message indicates that the average number of AIJ Request Blocks (ARB) formatted for each AIJ journal I/O operation is less than the corresponding threshold.

This message is computed by dividing the total number of ARBs submitted by the total number of AIJ journal write I/O operations.

As a process submits each ARB, the group commit process formats them into a single I/O buffer and ultimately writes the completed buffer to the AIJ journal. Ideally, there should be many ARBs per AIJ journal I/O, since this allows more transactional modifications to be written more efficiently.

Conversely, an extremely large number of ARBs per AIJ journal I/O may be an indication of small transactional modifications performed by many processes simultaneously. This usually results in sub-optimal ARB utilization, typically resulting in the ARB pool being frequently exhausted.

An ARB can contain approximately 2k in transaction modification information. The AIJ journal buffer size contains a maximum of 63k in transaction modification infor-

mation. Therefore, an optimal number of ARBs per AIJ journal I/O operation is in the 30 to 32 range. However, it is extremely difficult to achieve this goal in real-life. Therefore, a more reasonable number of ARBs is in the 20-25 range.

The default ARBs per I/O threshold is 2 ARBs. You can override the default value with the **RDMSBIND_STATS_AIJ_ARBS_PER_IO** logical name, or by defining the **AIJ_ARBS_PER_IO** configuration variable, or by selecting the “ARBs per AIJ I/O” option from the “Config” on-screen menu.

m>% blocks written per I/O below <n.m>% threshold

This message indicates that the average number of blocks written per AIJ journal write I/O operation is less than the corresponding threshold.

This message is computed by dividing the total number of AIJ journal blocks written by the total number of AIJ journal write I/O operations.

There is a maximum number of 127 blocks written to the AIJ journal per write I/O operation. Ideally, you want the average number of blocks to be as close to that value as possible.

The AIJ Log Server (ALS) is extremely efficient at issuing as many write I/O operations to the AIJ journal as possible. By doing this, the ALS process is able to keep the data streaming to the AIJ disk device. However, this sometimes sacrifices performance because more I/O is being written than is really necessary.

If the averaging number of blocks per AIJ journal write I/O operation is small, you can increase the minimum group commit buffer size using the **RDMSBIND_AIJ_IO_MIN** logical. There is no configuration variable for this value.

The default blocks per I/O threshold is 2 blocks. You can override the default value with the **RDMSBIND_STATS_AIJ_BLK_PER_IO** logical name, or by defining the **AIJ_BLK_PER_IO** configuration variable, or by selecting the “Blocks per AIJ I/O” option from the “Config” on-screen menu.

m>% background-formatted ARBs below <n.m>% threshold

This message indicates that the average number of AIJ Request Blocks (ARB) formatted in background by the AIJ Log Server (ALS) is less than the corresponding threshold.

The percentage is computed by dividing the total number of ARBs formatted in background by the total number of ARBs formatted (foreground and background).

Using the ALS is desirable for several reasons. The primary reason is that formatting the ARBs is CPU-intensive and allowing application processes to do this means they are NOT doing application work. The secondary reason is that the ALS formats

ARBs and nothing else; it uses a highly optimized algorithm, using a double-buffering scheme to perform recoverable asynchronous I/O to the AIJ journal.

The ALS algorithm always attempts to format more ARBs while the previously formatted ARBs are being written to disk. A foreground format results when an ARB is formatted while no disk I/O is pending. A background format results when an ARB is formatted while disk I/O is pending. Ideally, all ARBs should be formatted in background since this results in the best utilization of the ALS process. Remember that the AIJ journal is a sequential file. It is extremely desirable to keep the data streaming to the disk so that disk-head movement is optimized. Having to stop and format ARBs then issue an I/O means that the disk head is not properly positioned for optimum efficiency, which the result of increased latency.



Note that it is impossible to format 100% of the ARBs in the background, because the first AIJ journal write I/O operation is always a foreground format.

The default background ARB formatting threshold is 50%. You can override the default value with the **RDMSBIND_STATS_AIJ_BKGRD_ARB_RATIO** logical name, or by defining the **AIJ_BKGRD_ARB_RATIO** configuration variable, or by selecting the “Blocks per AIJ I/O” option from the “Config” on-screen menu.

IJ cache overflowed <n> times This message indicates that the AIJ cache has overflowed the specified number of times. There is no corresponding threshold for this message.

The AIJ cache is the buffer into which submitted ARBs are formatted. It can contain up to 63k of AIJ modification information. As ARBs are formatted into the cache, a number of events determine when a write I/O operation can be initiated to the AIJ journal. In general, however, the algorithm is optimized towards keeping data streaming to the AIJ journal in a steady manner. Thus, it is seldom the case that the AIJ cache becomes full before issuing a write I/O operation.

Therefore, it is a significant event when the AIJ cache overflows. This is typically indicative of excessive amounts of AIJ modifications being submitted at one time. It is possible with this type of behaviour to overwhelm the ALS into issuing a write I/O operation while an I/O is still pending. This causes the queue lengths on the disk device to become excessively large, causing a degradation of I/O throughput, which in turn causes the AIJ cache to overflow more often. A catch-22 situation that should be avoided.

cache overflow above <n.m>% threshold

This message indicates that the average number of AIJ cache overflows has exceeded the corresponding threshold.

The percentage is computed by dividing the total number of AIJ cache overflows by the total number of cache format operations.

This message is useful for eliminating the occasional AIJ cache overflow occurrence from the continual cache overflow problem.

The default AIJ cache overflow threshold is 10%. You can override the default value with the **RDMSBIND_STATS_AIJ_CACHE_OVF_RATIO** logical name, or by defining the **AIJ_CACHE_OVF_RATIO** configuration variable, or by selecting the “Cache overflow threshold” option from the “Config” on-screen menu.

Emergency AIJ creation not enabled

This message indicates that the Emergency AIJ facility is not enabled.

If the AIJ switch-over operation cannot complete because there are no available AIJ journals, the database enters the “AIJ suspended” state. During this state, the DBA can add new AIJ journals or perform database backups, but all other AIJ-related activities are temporarily suspended until such time as an AIJ journal becomes available.

During the AIJ suspended period, any DBR invocation causes the database to be shutdown. This is required because the DBR always writes either a commit or rollback record to the AIJ journal.



Note that even a DBR invoked for a read-only transaction causes the database to be shutdown.


The work-around to this problem is to always ensure that adequate AIJ journals are available for the AIJ switch-over operation. A new logical has been added to the product: **RDMSBIND_ALS_CREATE_AIJ**. This logical indicates whether or not the ALS server is to create an emergency AIJ journal if the AIJ switch-over operation ever enters the suspended state. The default value “0” indicates that the ALS should not create an AIJ journal. The value “1” indicates that the ALS should attempt to create an AIJ journal. The logical must reside in the LNM\$SYSTEM_TABLE logical name table.

The database recovery process DBR is also allowed to create emergency AIJ journals if the logical is enabled.

When the **RDMSBIND_ALS_CREATE_AIJ** logical is set to the value “1”, the ALS will attempt to create an emergency AIJ journal using the previous AIJ journal as a template. This means that the emergency AIJ journal will be created in the same directory, and with the same allocation, as the journal being switched from. If adequate disk space, or any other error occurs, the database will simply enter the “AIJ suspended” state and the DBA must resolve the situation.


WARNING: The emergency AIJ journal is not a temporary AIJ journal. DO NOT delete it using any means other than through the standard database syntax (SQL or RMU). Manually deleting the emergency AIJ journal (i.e. from DCL) will cause your database to be shutdown.

The location of the emergency AIJ journal (device/directory) can also be specified using the **RDMSBIND_AIJ_EMERGENCY_DIR** logical. This logical, if defined in the LNM\$SYSTEM_TABLE logical name table, should only specify the device & directory where the emergency AIJ journal is to be created.


 Note that, if defined, the **RDMSBIND_AIJ_EMERGENCY_DIR** logical applies to all databases on the current node. Furthermore, the logical must not contain any non-system concealed logical definitions.

The ALS will notify the DBA via the operator notification facility that an emergency AIJ journal has been created. Furthermore, the RMU Dump Header utility will identify any AIJ journal created by the ALS server process.

The RMU Show Statistic utility will highlight any identified emergency AIJ journal.

 Note that an emergency AIJ journal is a normal AIJ journal in all respects. It is simply created by the ALS (or DBR) process in order to avoid the AIJ switch-over suspension state. DBR invocations due to application process failure during the AIJ journal creation do not cause database shutdown.

When created, the name and file specification of the emergency AIJ journal is "EMERGENCY_XXX" where XXX is a series of 16 characters used to create a unique name.

 Note that there is no way to remove the emergency status of an emergency AIJ journal. However, an emergency AIJ journal is a normal AIJ journal in all other respects.

<n> journals are inaccessible This message indicates that the specified number of AIJ journals are inaccessible. There is no threshold associated with this message.

An AIJ journal becomes inaccessible when some event causes the AIJ journal to be unreadable or unwritable. The most frequently occurring cause is the accidental deletion of an AIJ journal.

The loss of *any* AIJ journal is cause for concern. The database is non-recoverable when it detects the loss of an AIJ journal.

Oracle recommends that you immediately backup your database as it cannot be recovered in the event of media failure.

network <n.m> BPS above <n.m> BPS threshold⁷

⁷ Oracle Rdb 7.0.2.0

This message indicates that the current Hot Standby network transmission rate exceeds the corresponding threshold. This value is expressed in terms of bits per second.

The value is computed using the total hot standby network packet size converted into bits per second.

This message is very useful for determining when the current AIJ throughput exceeds your network capacity. Exceeding your network capacity may cause your master database throughput to degrade in order to allow the standby database sufficient time to apply the database modifications.

The default AIJ network bandwidth threshold is 60 seconds. You can override the default value with the **RDMSBIND_STATS_AIJ_NETWORK_BPS** logical name, or by defining the **AIJ_NETWORK_BPS** configuration variable, or by selecting the "Network bandwidth threshold" option from the "Config" on-screen menu.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

C O N F I G U R A T I O N
O P T I O N S

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Seconds to AIJ extend/switch This configuration option is used to change the default threshold, which is 60 seconds. You can override the default value with the RDM\$BIND_STATS_AIJ_SEC_TO_EXTEND logical name.

ARBs per AIJ I/O This configuration option is used to change the default threshold, which is 2. You can override the default value with the RDM\$BIND_STATS_AIJ_ARBS_PER_IO logical name.

Blocks per AIJ I/O This configuration option is used to change the default threshold, which is 2. You can override the default value with the RDM\$BIND_STATS_AIJ_BLK_PER_IO logical name.

Background ARB threshold This configuration option is used to change the default threshold, which is 50%. You can override the default value with the RDM\$BIND_STATS_AIJ_BKGRD_ARB_RATIO logical name.

Cache Overflow threshold This configuration option is used to change the default threshold, which is 10%. You can override the default value with the RDM\$BIND_STATS_CACHE_OVF_RATIO logical name.

Network bandwidth threshold This configuration option is used to change the default threshold, which is 60 seconds. You can override the default value with the RDM\$BIND_STATS_AIJ_NETWORK_BPS logical name.

RUJ Analysis Screen

This screen provides information about RUJ journal performance.

SCREEN LOCATION

This screen resides in the "Online Analysis & Info." menu.

SCREEN EXAMPLE

The following is an example of the "RUJ Analysis" screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second        RUJ Analysis                Elapsed: 00:00:13.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1  Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

m>% synchronous RUJ I/O above <n.m>% threshold

This message indicates that the average number of synchronous write I/O operations to the RUJ journal exceeds the corresponding threshold.

The percentage is computed as the number of synchronous write I/O operations to the RUJ journal divided by the total number of write I/O operations, both synchronous and asynchronous.

Ideally, most write I/O operations to the RUJ should be asynchronous. Most synchronous write I/O operations to the RUJ journal occur in response to a request for a particular modified page by another process. The modifications on that page must first be hardened to the RUJ journal before the page can be written to the storage area and released to the requesting process. This synchronous I/O operation causes the process to stall for a relatively long time, approximately 20ms on most disk devices.

Unfortunately, there is no method available to indicate that you wish to ignore page requests from other processes.

Asynchronous write I/O operations to the RUJ journal are preferable in that the process is able to continue performing application work while the I/O operations are still pending. Batch writes are the frequent cause of asynchronous write I/O operations to the RUJ journal. These batch write operations typically occur at transaction commit, rollback and checkpoint intervals, but they can also occur in response to a batch-write request.

Ideally, you want to always pre-write modified pages back to the database before they are requested by other users. Adjusting the asynchronous batch write facility parameters is a good method to accomplish this. In particular, you should examine the clean buffer count and batch size attributes. You can use the **RDMSBIND_BATCH_MAX** logical, or the equivalent SQL ABW syntax, to specify the number of buffers a process will write asynchronously as part of the ABW operation. The database uses this value whenever the number of unmodified (clean) buffers for the process is less than the value specified by the **RDMSBIND_CLEAN_BUF_CNT** logical.

The value of the **RDMSBIND_CLEAN_BUF_CNT** logical adjusts automatically at runtime so that it does not exceed more than half the number of buffers available to the process.

It is important that the value of the **RDMSBIND_BATCH_MAX** logical be neither too high nor too low. In general, the logical value should probably not exceed the number of storage areas or disk devices, with the result that a sudden burst of write I/O operation requests should not overwhelm the disk controllers. If the logical value is too low, the process will not benefit from any disk device parallelism. For example, if you are using five disk devices but have the **RDMSBIND_BATCH_MAX** logical set to the value "2", then an ABW operation will write to only two disk devices in parallel.

The default RUJ journal synchronous I/O threshold is 10%. You can override the default value with the **RDMSBIND_STATS_RUJ_SYNC_IO_RATIO** logical name, or by defining the **RUJ_SYNC_IO_RATIO** configuration variable, or by selecting the "Synchronous RUJ I/O threshold" option from the "Config" on-screen menu.

m>% RUJ extends above <n.m>% threshold⁸

This message indicates that the average number of file extension operations to the RUJ journal exceeded the corresponding threshold.

The percentage is computed as the number of file extension operations to the RUJ journal divided by the total number of write I/O operations, both synchronous and asynchronous.

Ideally, there should be only one extend of the RUJ journal per database process attach. An RUJ journal is initially created with only header information; this always re-

⁸ Oracle Rdb 7.0.2.0

sults in at least one file extend operation. Beyond the initial extension operation, each subsequent RUJ journal file extension greatly impacts the transaction throughput. Oracle recommends setting the **RDMSBIND_RUJ_ALLOC_BLCNT** and **RDMSBIND_RUJ_EXTEND_BLCNT** logicals to a reasonable value, such that no RUJ journal extends occur at runtime.

The default RUJ journal extend threshold is 2%. You can override the default value with the **RDMSBIND_STATS_RUJ_EXTEND_RATIO** logical name, or by defining the **RUJ_EXTEND_RATIO** configuration variable, or by selecting the “RUJ extend thresh- old” option from the “Config” on-screen menu.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
 OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Asynchronous RUJ I/O threshold This configuration option is used to change the default threshold, which is 10%. You can override the default value with the **RDMSBIND_STATS_RUJ_SYNC_IO_RATIO** logical name.

Recovery Analysis Screen

This screen provides information about process failure and database recovery performance.

SCREEN LOCATION

This screen resides in the “Online Analysis & Info.” menu.

SCREEN EXAMPLE

The following is an example of the “Recovery Analysis” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second        Recovery Analysis          Elapsed: 00:00:13.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

m>% process failure rate above <n.m>% threshold

This message indicates that the average number of database recovery processes exceeds the corresponding threshold. The percentage is computed by dividing the total number of database recovery operations by the total number of database attaches.

When an application or server process terminates prematurely, the database Monitor freezes all database activity and invokes a database recovery process (DBR). The DBR process is responsible for ensuring that all transactional changes performed by the failed process either commit or rollback. While the DBR process is performing this recovery operation, the database Monitor suspends all other transactional activity on the entire database, including on other nodes in the cluster.



Note that the DBR process will recover a process using read-only transactions, even though there may not actually be anything to recover.

Ideally, there should be no application process failures, but it is often impossible to prevent external failures such as network disconnects or power outages. In particular, Oracle recommends that you disallow the Ctrl/Y STOP action on OpenVMS; this is a major source of process recovery operations.

The default process failure threshold is 15%. You can override the default value with the **RDMSBIND_STATS_DBR_RATIO** logical name, or by defining the **DBR_RATIO** configuration variable, or by selecting the “DBR invocation threshold” option from the “Config” on-screen menu.

m>% process recovery duration above <n.m>% threshold⁹

This message indicates that the average duration of the database recovery processes exceeds the corresponding threshold. The percentage is computed by dividing the total durations of all database recovery operations by the total number of database recovery operations.

This message is effectively indicating the average time the database is unavailable as a result of each process failure. When an application or server process terminates prematurely, the database Monitor freezes all database activity and invokes a database recovery process (DBR). The DBR process is responsible for ensuring that all transactional changes performed by the failed process either commit or rollback. While this recovery operation is being performed, the database Monitor suspends all other transactional activity on the entire database, including on other nodes in the cluster.



Note that the DBR process performs recovery for all failed processes, including those processes using read-only transactions, even though there may not actually be anything to recover.

Ideally, there should be no application process failures, but it is often impossible to prevent some types of external events such as network disconnects or power outages.

Database attributes such as the Fast Commit feature affect the process recovery duration. Also, features that require the Fast Commit feature often increase the process recovery duration significantly. Ideally, the process recovery duration should be as short as possible, since it affects all other processes accessing the database.

The default process failure threshold is 2 seconds. You can override the default value with the **RDMSBIND_STATS_DBR_DURATION** logical name, or by defining the **DBR_DURATION** configuration variable, or by selecting the “DBR recovery duration” option from the “Config” on-screen menu.

database backup has not been performed since <date> _____

⁹ Oracle Rdb 7.0.2.0

This message indicates that the database has either *never* been backed up since it was created, or that the last backup database occurred in the past more days than the corresponding threshold.

If your database has never been backed up, Oracle highly recommends that you do so immediately. Also, never, ever use the OpenVMS backup utility to backup your database; it does not work correctly, especially when using global buffers or record caches.

Ideally, you want to backup your database every day. However, this is often impossible in today's 7*24 environment. Therefore, Oracle recommends that you backup your database a minimum of once a week. You always need to customize a database backup strategy for your specific operating environment. The backup strategy should take into consideration the system workload, usage schedule, importance of data, and hardware environment of the database. There are some general guidelines that apply to all databases. Incorporating these guidelines into your strategy will ensure more reliable and more cost effective backups.

The following rules are essential to a proper database backup strategy. Always use after-image journals. Always maintain a log of your backup schedule. Occasionally store backups in a separate location, preferably off-site. Do not use the OpenVMS Backup Utility on your database; Oracle even recommends that you set the do-not-backup attribute on database files. Spread your database across several disks. Always consider your existing hardware configuration and system workload. Determine the volatility and value of your data. Verify the integrity of your database regularly. Verify the integrity of your backup media after each backup. Failure to perform these simple recommendations may result in your database being non-recoverable.

The default database backup threshold is 6 days. You can override the default value with the **RDMSBIND_STATS_FULL_BACKUP_INTRVL** logical name, or by defining the **FULL_BACKUP_INTRVL** configuration variable, or by selecting the "Full database backup threshold" option from the "Config" on-screen menu.

Database is single-file¹⁰ This message indicates your database is a single-file database. There are no thresholds associated with this message.

Oracle highly recommends that you convert your single-file database into a multi-file database.

ON-SCREEN MENU Options This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

¹⁰ Oracle Rdb 7.0.2.0

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

DBR invocation threshold This configuration option is used to change the default threshold, which is 15%. You can override the default value with the RDM\$BIND_STATS_DBR_RATIO logical name.

Full database backup threshold This configuration option is used to change the default threshold, which is 6 days. You can override the default value with the RDM\$BIND_STATS_FULL_BACKUP_INTRVL logical name.

Record Analysis Screen

This screen provides information about record storage and retrieval performance, including SPAM page retrieval performance.

SCREEN LOCATION

This screen resides in the “Online Analysis & Info.” menu.

SCREEN EXAMPLE

The following is an example of the “Record Analysis” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second          Record Analysis              Elapsed: 00:00:13.42
Page: 1 of 1      DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

m>% page discard rate above <n.m>% threshold (avg <n.m> I/Os)

This message indicates that the average number of pages discarded has exceeded the corresponding threshold, resulting in the average number of extra read I/O operations.

The percentage is computed using the total number of pages discarded divided by the total number of pages checked. The average number of read I/O operations incurred as a result of the discarded pages is computed using the total number of discarded pages that required a read I/O operation divided by the total number of pages checked.

Performance degradation most often results when storing new records in a database. Sometimes this is due to the way that Oracle Rdb manages free space on data pages. There are times when Oracle Rdb will read many pages in a storage area before it finds a page with sufficient space on it to store a record.

The most simple way to verify if any of the various pages checked scenarios described are occurring is to use RMU Show Statistic utility. On the “Record Statistics” screen you should examine the following two statistics: “record stored” and “pages checked”. If you see that the “pages checked” number is significantly higher than the “record stored” number then it is very likely that one of the situations described in the following sections is being encountered.

The default pages checked threshold is 10%. You can override the default value with the **RDMSBIND_STATS_PAGES_CHECKED_RATIO** logical name, or by defining the **PAGES_CHECKED_RATIO** configuration variable, or by selecting the “Pages checked threshold” option from the “Config” on-screen menu.

m>% SPAM page fetch rate above <n.m>% total fetched threshold¹¹

This message indicates that the average number of SPAM page fetches exceeds the corresponding threshold for the total number of SPAM page fetches.

The percentage is computed using the total number of SPAM pages fetched to store a record divided by the total number of SPAM page fetches for any reason. In most real-world applications, new records stored exceed the number of record modifications or deletions.

The default SPAM fetched threshold is 80%. You can override the default value with the **RDMSBIND_STATS_SPAM_PAG_FET_RATIO** logical name, or by defining the **SPAM_PAG_FET_RATIO** configuration variable, or by selecting the “SPAM pages fetched threshold” option from the “Config” on-screen menu.

m>% SPAM page fetch rate above <n.m>% record stored threshold¹²

This message indicates that the average number of SPAM page fetches exceeds the corresponding threshold for stored records.

The percentage is computed using the total number of SPAM pages fetched to store a record divided by the total number records stored.

When inserting a record into a full large uniform format area, a large number of SPAM reads is normal. Updates that cause fragmentation have the same effect when inserting the new fragment. This occurs when the logical area is scanned for free space. The SPAM pages make this much faster than reading all the data pages would be. If the logical area is full then every SPAM page must be read to check that there is any free space. The logical area is then expanded by one clump, which contains

¹¹ Oracle Rdb 7.0.2.0

¹² Oracle Rdb 7.0.2.0

enough pages to fill one buffer. The next few inserts will be fast, until this new clump is full, when the entire area is scanned again, just in case another process may have deleted something. So the average number of SPAM reads per insert should be approximately the number of SPAM pages in the area divided by the number of records that will fit in a clump.

However, this will not divide evenly. All but one record insert into a clump will access just one SPAM page; the other record insert will read all SPAM pages. This fits in with only some transactions being affected by excessive SPAM page fetches.

The default SPAM record stored threshold is 20%. You can override the default value with the **RDM\$BIND_STATS_SPAM_REC_STO_RATIO** logical name, or by defining the **SPAM_REC_STO_RATIO** configuration variable, or by selecting the “SPAM records stored threshold” option from the “Config” on-screen menu.

area_name> has most pages checked

This message indicates that the identified storage area contains the most pages checked. There is no threshold associated with this message; it is for informational purposes only.



Note that this message does not take into account other activity that might be occurring on the identified storage area; it simple indicates that this storage area has the most. It is often quite possible that storage areas experiencing a lot of pages checked may contain tables that should be moved to another storage area, either in their entirety or via partitioning to multiple storage areas.

area_name> has <n.m>% synchronous pages checked

This message indicates that the identified storage area contains the largest percentage of pages checked. There is no threshold associated with this message; it is for informational purposes only.

The percentage is determined using the total number of pages checked divided by the total number of synchronous read I/O operations. The synchronous read I/O operations are used because asynchronous read I/O operations do not cause the process to stall and, therefore, are considered light-weight. However, synchronous read I/O operations cause the process to stall for approximately 20ms, during which time no other work is being performed. This makes the pages checked a disruptive operation to high throughput.



Note that this message does not take into account other activity that might be occurring on the identified storage area; it simple indicates that this storage area has the highest percentage. It is often quite possible that storage areas experiencing a lot of pages checked may contain tables that should be moved to another storage area, ei-

ther in their entirety or via partitioning to multiple storage areas.

m>% fragmented records stored above <n.m>% threshold

This message indicates that the average number of fragmented stored exceeds the corresponding threshold.

The percentage is computed using the total number of fragmented records stored divided by the total number of record stored.

Ideally, this value should be as close to 0% as possible, indicating the lack of record fragmentation.

There are several reasons why there may not be adequate space on a page to store a new record, thus causing the record to be fragmented. Thresholds can be set incorrectly on a mixed format storage area or on logical area thresholds defined for a table or index. Some or all of the free space is locked for use by other database users. The length used to represent the size of the data records does not always represent the actual length of the data records; this is always the case for duplicate nodes on indexes that allow duplicates. The SPAM page entry is actually incorrect

The default fragmented records stored threshold is 20%. You can override the default value with the **RDMSBIND_STATS_RECS_STORED_RATIO** logical name, or by defining the **RECS_STORED_RATIO** configuration variable, or by selecting the “Records stored threshold” option from the “Config” on-screen menu.

m>% fragmented records fetched above <n.m>% threshold

This message indicates that the average number of fragmented record retrieved exceeds the corresponding threshold.

The percentage is computed using the total number of fragmented records retrieved divided by the total number of record retrieved.

Ideally, this value should be as close to 0% as possible, indicating the lack of record fragmentation.

The default fragmented records fetched threshold is 20%. You can override the default value with the **RDMSBIND_STATS_RECS_FETCHED_RATIO** logical name, or by defining the **RECS_FETCHED_RATIO** configuration variable, or by selecting the “Records fetched threshold” option from the “Config” on-screen menu.

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “C”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Pages checked threshold This configuration option is used to change the default threshold, which is 10%. You can override the default value with the RDMSBIND_STATS_PAGES_CHECKED_RATIO logical name.

Records stored threshold This configuration option is used to change the default threshold, which is 20%. You can override the default value with the RDMSBIND_STATS_RECS_STORED_RATIO logical name.

Records fetched threshold This configuration option is used to change the default threshold, which is 20%. You can override the default value with the RDMSBIND_STATS_RECS_FETCHED_RATIO logical name.

Area Analysis Screen

This screen provides information about storage area performance.

The “Area Analysis” screen is unique in that it does not have a configuration sub-menu. There are no thresholds associated with any messages in this screen.

SCREEN LOCATION

This screen resides in the “Online Analysis & Info.” menu.

SCREEN EXAMPLE

The following is an example of the “Area Analysis” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second          Area Analysis                Elapsed: 00:00:13.42
Page: 1 of 1      DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

ice <device_name> exceeds average storage area count <n>

This message indicates that the identified device contains more than the average number of storage areas. This may be an indication that the disk device contains too many storage areas and may be performing sub-optimally.

Ideally, storage areas should reside on their own disk devices. However, this may not always be possible. In those cases where spreading the storage areas over multiple disk devices is not possible, it is best to mix heavily utilized storage areas with minimal-use storage areas. Also, mixing update-intensive storage areas with read-only storage areas may be helpful.

rea_name> extended <n> times total <n> times

This message indicates that the identified storage area has extended the indicated number of times since the database Monitor opened the database. The identified storage area has also extended the total number of times since the database was created.

There are several reasons why storage areas extend. The most obvious reason is that the storage area was completely full and needed to extend to allow the continued storing of new records. However, another frequent cause of storage area extension is improperly sized mixed-format storage areas used for hash indexes.

For the hash index to operate properly, sufficient space must be allocated in the storage area to ensure that all records distribute evenly and randomly across all pages of the storage area. When records begin filling pages in a non-uniform manner, performance of the hash index will degrade. The reason for this degradation is that the hashing algorithm formulates the target DBKEY from the record search key. If the target page is full, then an adjacent page will be retrieved and scanned using the SPAM thresholds you established for the storage area. Of course, using the new page will cause cascading problems when *that* page is subsequently selected as a target page.

Ultimately, the storage area will be extended even though the hash index may not be completely full.

rea_name> async read I/O stalls <n.m> exceeded average <n.m>

This message indicates that the average asynchronous I/O read operation duration for the identified storage area exceeded the average of all storage areas.

This message is meant to provide an indication of which storage areas reside on slow disk devices, or over-utilized disk devices. It may be necessary or prudent to move these storage areas to another disk device, if possible.

rea_name> async write I/O stalls <n.m> exceeded average <n.m>

This message indicates that the average asynchronous I/O write operation duration for the identified storage area exceeded the average of all storage areas.

This message is meant to provide an indication of which storage areas reside on slow disk devices, or over-utilized disk devices. It may be necessary or prudent to move these storage areas to another disk device, if possible.

rea_name> sync read I/O stalls <n.m> exceeded average <n.m>

This message indicates that the average synchronous I/O read operation duration for the identified storage area exceeded the average of all storage areas.

This message is meant to provide an indication of which storage areas reside on slow disk devices, or over-utilized disk devices. It may be necessary or prudent to move

these storage areas to another disk device, if possible.

rea_name> sync write I/O stalls <n.m> exceeded average <n.m>

This message indicates that the average synchronous I/O write operation duration for the identified storage area exceeded the average of all storage areas.

This message is meant to provide an indication of which storage areas reside on slow disk devices, or over-utilized disk devices. It may be necessary or prudent to move these storage areas to another disk device, if possible.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

The “Area Analysis” screen is unique amongst the online analysis screens, in that it does not have a configuration sub-menu. There are no thresholds associated with any messages in this screen.

Locking Analysis Screen

This screen provides information about locking performance.

SCREEN LOCATION

This screen resides in the “Online Analysis & Info.” menu.

SCREEN EXAMPLE

The following is an example of the “Locking Analysis” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second        Locking Analysis           Elapsed: 00:00:13.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

cess <n.m> excessive stall on lock <x> for <object>

This message indicates that the identified process has stalled on a lock for the identified object for a period of time exceeding the corresponding threshold.

To reduce lock contention with other processes, Oracle recommends you keep your transactions as short in duration as possible, while providing as much data sharing as possible. If you need to retrieve a look-up record to display information, you should specify an appropriate sharing level that allows other transactions concurrent access to the same information.

Oracle recommends you use the Adjustable Locking Granularity (ALG) feature. The ALG feature allows Rdb to maintain the minimum number of record locks while ensuring maximum concurrent access to those records.

Reducing the number of locks is sometimes more important than reducing the transaction duration. For example, page-level locking might be appropriate for storage areas containing clustered record. Often times, a 5:1 reduction in lock requests is possi-

ble in some situations. However, page-level locking is not always appropriate for all applications.

The default lock stall threshold is 2 seconds. You can override the default value with the RDMS\$BIND_STATS_MAX_LOCK_STALL logical name, the MAX_LOCK_STALL configuration variable, or the “Lock stall threshold” option from the “Config” on-screen menu.

Process <n.m> excessive timeouts <n> on <object>

This message indicates that the identified process has experienced excessive lock time-outs, the most recent of which was on the identified object. There is no threshold associated with this message. This message is displayed when the process’ number of lock time-outs exceeds the average of all processes’ number of lock time-outs.

Process <n.m> excessive deadlocks <n> on <object>

This message indicates that the identified process has experienced excessive lock deadlocks, the most recent of which was on the identified object. There is no threshold associated with this message. This message is displayed when the process’ number of lock deadlocks exceeds the average of all processes’ number of lock deadlocks.

ON-SCREEN MENU
 OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
 INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Lock stall threshold This configuration option is used to change the default threshold, which is 2 seconds. You can override the default value with the RDM\$BIND_STATS_MAX_LOCK_STALL logical name.

Index Analysis Screen

This screen provides information about B-tree and hash index performance.

SCREEN LOCATION

This screen resides in the “Online Analysis & Info.” menu.

SCREEN EXAMPLE

The following is an example of the “Index Analysis” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second        Index Analysis                Elapsed: 00:00:13.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

m>% duplicate btree fetch above <n.m>% threshold

This message indicates that the average number of btree index duplicate node fetches exceeds the corresponding threshold.

The default btree index duplicate fetch threshold is 15%. You can override the default value with the RDMSBIND_STATS_BTR_FETCH_DUP_RATIO logical name, the BTR_FETCH_DUP_RATIO configuration variable, or the “B-tree duplicate fetch threshold” option from the “Config” on-screen menu.

m>% duplicate btree store above <n.m>% threshold

This message indicates that the average number of btree index duplicate node stores exceeds the corresponding threshold.

The default btree index duplicate store threshold is 15%. You can override the default value with the RDMSBIND_STATS_BTR_STORE_DUP_RATIO logical name, the

BTR_STORE_DUP_RATIO configuration variable, or the “B-tree duplicate store threshold” option from the “Config” on-screen menu.

m>% leaf-node btree fetch below <n.m>% threshold

This message indicates that the average number of btree index leaf-node retrievals is less than the corresponding threshold.

This may be an indication that the btree index contains too many levels. You may want to adjust the fan-out factors during index creation to use a larger btree that contains fewer levels. This may help improve overall btree index throughput and utilization.

The default lock stall threshold is 25%. You can override the default value with the **RDM\$BIND_STATS_BTR_LEF_FETCH_RATIO** logical name, the **BTR_LEF_FETCH_RATIO** configuration variable, or the “B-tree leaf node store threshold” option from the “Config” on-screen menu.

m>% duplicate hash index fetch above <n.m>% threshold¹³

This message indicates that the average number of hash index duplicate node fetches exceeds the corresponding threshold.

The default hash index duplicate fetch threshold is 15%. You can override the default value with the **RDM\$BIND_STATS_HASH_FETCH_DUP_RATIO** logical name, the **HASH_FETCH_DUP_RATIO** configuration variable, or the “Hash duplicate fetch threshold” option from the “Config” on-screen menu.

m>% duplicate hash index store above <n.m>% threshold¹⁴

This message indicates that the average number of hash index duplicate node stores exceeds the corresponding threshold.

The default hash index duplicate store threshold is 15%. You can override the default value with the **RDM\$BIND_STATS_HASH_STORE_DUP_RATIO** logical name, the **HASH_STORE_DUP_RATIO** configuration variable, or the “Hash duplicate store threshold” option from the “Config” on-screen menu.

¹³ Oracle Rdb 7.0.2.0

¹⁴ Oracle Rdb 7.0.2.0

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

B-tree duplicate fetch threshold This configuration option is used to change the default threshold, which is 15%. You can override the default value with the RDMSBIND_STATS_BTR_FETCH_DUP_RATIO logical name.

B-tree duplicate store threshold This configuration option is used to change the default threshold, which is 15%. You can override the default value with the RDMSBIND_STATS_BTR_STORE_DUP_RATIO logical name.

B-tree leaf node fetch threshold This configuration option is used to change the default threshold, which is 25%. You can override the default value with the RDMSBIND_STATS_BTR_LEF_FETCH_RATIO logical name.

Hash duplicate fetch threshold This configuration option is used to change the default threshold, which is 25%. You can override the default value with the RDMSBIND_STATS_HASH_FETCH_DUP_RATIO logical name.

Hash duplicate store threshold This configuration option is used to change the default threshold, which is 25%. You can override the default value with the `RDM$BIND_STATS_HASH_STORE_DUP_RATIO` logical name.

Row Cache Analysis Screen

This screen provides information about row cache effectiveness and performance.

SCREEN LOCATION

This screen resides in the “Online Analysis & Info.” menu.

SCREEN EXAMPLE

The following is an example of the “Row Cache Analysis” screen:

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second        Row Cache Analysis      Elapsed: 00:00:13.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1      Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

Row cache is not enabled This message indicates that the row caching feature is not available on your database. There is no threshold associated with this message.

The following requirements must be met in order for the row caching feature to be enabled:





At least 1 storage area must be enabled for record caching


Note that logical area row caches do not correlate into the determination whether or not row caching is enabled.

Row cache is not allowed This message indicates that the row caching feature is not allowed on your database. There is no threshold associated with this message.

The following requirements must be met in order for the row caching feature to be allowed:


 Row Caching must be enabled (see above)


 At least one row cache must be active


 The database node count must be “1”

 The Fast Commit feature must be enabled

The Fast Commit feature requires the following:

 AII journaling must be enabled

 Fast commit must be enabled

 All snapshot files must be allowed

<rcache_name> Hash Queues: This message indicates that the hash queues for the identified record cache exceed the corresponding threshold.

The default lock stall threshold is 25%. You can override the default value with the **RDMSBIND_STATS_MAX_HASH_QUE_LEN** logical name, the **MAX_HASH_QUE_LEN** configuration variable, or the “Hash table queue length threshold” option from the “Config” on-screen menu.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the “Cluster Statistic Collection” facility.

Zoom Screen This screen does not have a zoom screen.

C O N F I G U R A T I O N
O P T I O N S

You can use the Config menu option, by typing the letter “**C**”, to display the configuration sub-menu. The configuration sub-menu provides the following options:

Hash table queue length threshold This configuration option allows you to change the default threshold, which is 2 rows. You can override the default value with the RDM\$BIND_STATS_MAX_HASH_QUE_LEN logical name.

Process Analysis Screen¹⁵

This screen provides information about process effectiveness and performance.

This screen is only available on OpenVMS.

This screen resides in the “Online Analysis & Info.” menu.

SCREEN LOCATION

The following is an example of the “Process Analysis” screen:

SCREEN EXAMPLE

```
Node: MYNODE (1/1/1)      Oracle Rdb X7.0-00 Perf. Monitor  8-FEB-1998 10:35:25.29
Rate: 1.00 Second        Process Analysis              Elapsed: 00:00:13.42
Page: 1 of 1            DISK$:[DB_HOMEDIR.WORK.STATS]MF_PERSONNEL.RDB;1    Mode: Online
-----
```

Nothing to report with current configuration

```
-----
Config Exit Help Menu Set_rate Write !
```

SCREEN FIELDS

Process n ENQLM <x.y> ratio below <x.y>% threshold

This message indicates that the process ENQCNT to ENQLM ratio is below the corresponding threshold.

The default ENQLM threshold is 25%. You can override the default value with the RDM\$BIND_STATS_PROC_ENQLM_RATIO logical name, the PROC_ENQLM_RATIO configuration variable, or the “ENQLM threshold” option from the configuration submenu.

Process n ASTLM <x.y> ratio below <x.y>% threshold

¹⁵ Available in Oracle Rdb 7.0.2.

This message indicates that the process ASTCNT to ASTLM ratio is below the corresponding threshold.

The default ASTLM threshold is 25%. You can override the default value with the RDM\$BIND_STATS_PROC_ASTLM_RATIO logical name, the PROC_ASTLM_RATIO configuration variable, or the “ASTLM threshold” option from the configuration sub-menu.

Process n BIOLM <x.y> ratio below <x.y>% threshold

This message indicates that the process BIOCNT to BIOLM ratio is below the corresponding threshold.

The default BIOLM threshold is 25%. You can override the default value with the RDM\$BIND_STATS_PROC_BIOLM_RATIO logical name, the PROC_BIOLM_RATIO configuration variable, or the “BIOLM threshold” option from the configuration sub-menu.

Process n DIOLM <x.y> ratio below <x.y>% threshold

This message indicates that the process DIOCNT to DIOLM ratio is below the corresponding threshold.

The default DIOLM threshold is 25%. You can override the default value with the RDM\$BIND_STATS_PROC_DIOLM_RATIO logical name, the PROC_DIOLM_RATIO configuration variable, or the “DIOLM threshold” option from the configuration sub-menu.

Process n PGFLQUOTA <x.y> ratio below <x.y>% threshold

This message indicates that the process PAGFILCNT to PGFLQUOTA ratio is below the corresponding threshold.

The default PGFLQUOTA threshold is 25%. You can override the default value with the RDM\$BIND_STATS_PROC_PGFLLM_RATIO logical name, the PROC_PGFLLM_RATIO configuration variable, or the “PGFLQUOTA threshold” option from the configuration submenu.

ON-SCREEN MENU
OPTIONS

This section discusses the on-screen menu options available to this screen. Remember that the Exit, Help (?), Menu, Set_rate, Write and (!) on-screen menu options are available on all screens unless otherwise noted.

Options This on-screen menu option displays a menu containing options for writing all of the screens to a text file named STATISTICS.RPT. This option should be used with care because it generates a very large output file.

SCREEN
INFORMATION

This section discusses screen-specific issues.

Page Navigation This screen does not contain multiple pages of information and, therefore, does not contain any special navigation keystrokes.

Binary File Support This screen is not recorded in the binary output file produced using the OUTPUT qualifier. Consequently, this screen is not available when you replay a binary file using the INPUT qualifier.

Cluster Statistic Collection Support This screen is not integrated into the "Cluster Statistic Collection" facility.

Zoom Screen This screen does not have a zoom screen.

CONFIGURATION
OPTIONS

You can use the Config menu option, by typing the letter "C", to display the configuration sub-menu. The configuration sub-menu provides the following options:

Hash table queue length threshold This configuration option allows you to change the default threshold, which is 2 rows. You can override the default value with the RDMSBIND_STATS_MAX_HASH_QUEUE_LEN logical name.

Index

/INPUT, 8, 11, 14, 18, 22, 25, 29, 31, 36, 39, 43, 46, 50, 53, 57, 59, 63, 66, 68, 71, 74, 77, 80, 83, 86, 89, 92, 97, 100, 103, 106, 109, 112, 116, 120, 124, 128, 133, 137, 141, 144, 148, 152, 156, 159, 164, 169, 173, 180, 185, 198, 201, 206, 211, 214, 216, 220, 224, 227

/OPTIONS, 114, 118, 122, 126, 130, 135, 139, 142, 146, 150, 154, 157, 161, 166, 171

/OUTPUT, 8, 11, 14, 18, 22, 25, 29, 31, 116, 120, 124, 128, 133, 137, 141, 144, 148, 152, 156, 159, 164, 169, 173, 180, 185, 198, 201, 206, 211, 214, 216, 220, 224, 227

/SCREEN, 7, 11, 14, 18, 22, 25, 28, 30, 35, 39, 42, 45, 49, 52, 56, 59, 62, 65, 68, 71, 74, 77, 80, 83, 86, 89, 92, 96, 99, 102, 105, 108, 111, 116, 120, 124, 127, 133, 136, 140, 144, 148, 152, 155, 159, 163, 168, 172, 180, 197, 201, 205, 210, 214, 216, 220, 223, 227

/TIME, 127, 158, 168

"Lock stall threshold, 216

ACTIVE, 76, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111

Adjustable Locking Granularity, 215

AIJ Backup Server, 143, 188, 190, 191

AIJ cache, 194

AIJ Log Server, 143, 151, 188, 189, 193

AIJ Request Block, 192

AIJ Request Blocks, 189, 192, 193

AIJ_ARBS_PER_IO, 193

AIJ_BKGRD_ARB_RATIO, 194

AIJ_BLKES_PER_IO, 193

AIJ_CACHE_OVF_RATIO, 195

AIJ_NETWORK_BPS, 197

AIJ_SEC_TO_EXTEND, 190

AIJDB, 67, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111

AIJFB, 70, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111

ARBs per AIJ I/O, 193

Asynchronous Batch-Write, 177

Asynchronous Pre-Fetch, 176

Blocks per AIJ I/O, 193, 194

BTR_FETCH_DUP_RATIO, 218

BTR_LEF_FETCH_RATIO, 219

BTR_STORE_DUP_RATIO, 219

B-tree duplicate fetch threshold, 218

B-tree duplicate store threshold, 219

B-tree leaf node store threshold, 219

Cache overflow threshold, 195

CLIENT, 85, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111

Commit To Journal optimization, 127, 168, 187

CPT, 79, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111

database recovery process, 203, 204

DBR invocation threshold, 204
 DBR recovery duration, 204
DBR_DURATION, 204
DBR_RATIO, 204
 Detected Asynchronous Pre-Fetch, 177
 Emergency AIJ, 195
 Emergency AIJ facility, 151, 189, 195
 Emergency AIJ Journal, 189
 Fast Commit, 178, 187, 190, 204, 223
 FILID, 58, 95, 98, 101, 104, 107, 110
 Full database backup threshold, 205
FULL_BACKUP_INTRVL, 205
 GB IO-saved threshold, 180
 GB pool hit threshold, 179
GB_IO_SAVED_RATIO, 180
GB_POOL_HIT_RATIO, 179
 Hash duplicate fetch threshold, 219
 Hash duplicate store threshold, 219, 225, 226
 Hash table queue length threshold, 223
HASH_FETCH_DUP_RATIO, 219
HASH_STORE_DUP_RATIO, 219, 225, 226
 Hot Standby, 130, 131, 133, 197
 Journal Overwrite, 189
 LB/AS page hit threshold, 179
LB_PAGE_HIT_RATIO, 179
 Log Replication Server, 131, 132, 143
MAX_HASH_QUE_LEN, 223
MAX_LOCK_STALL, 216
MAX_RO_TX_DURATION, 185
MAX_RO_TX_PERCENTILE, 185
MAX_RW_TX_DURATION, 184
MAX_TX_DURATION, 183
MAX_TX_PERCENTILE, 183
 Network bandwidth threshold, 197
 Online Analysis facility, 175, 176, 182, 187, 199, 203, 207, 212, 215, 218, 222, 225
 Pages checked threshold, 208
PAGES_CHECKED_RATIO, 208
 RCACHE, 82, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111, 135, 136, 172
 RDM\$BIND_ABS_OVERWRITE_ALLOWED, 147
 RDM\$BIND_ABS_OVERWRITE_IMMEDIATE, 147
 RDM\$BIND_ABS_PRIORITY, 143
 RDM\$BIND_ABS_QUIET_POINT, 146, 147
 RDM\$BIND_ABW_ENABLED, 116, 163, 178
RDM\$BIND_AIJ_ARB_COUNT, 192
 RDM\$BIND_AIJ_CHECK_CONTROL_RECS, 124, 151, 168
RDM\$BIND_AIJ_EMERGENCY_DIR, 196
 RDM\$BIND_AIJ_IO_MAX, 123, 167
 RDM\$BIND_AIJ_IO_MIN, 122, 123, 167, 193
 RDM\$BIND_AIJ_SHUFFLE_DISABLED, 124
 RDM\$BIND_AIJ_STALL, 123, 167
 RDM\$BIND_AIJ_SWITCH_GLOBAL_CKPT, 123, 150, 151, 168
 RDM\$BIND_ALS_CREATE_AIJ, 151, 195
 RDM\$BIND_ALS_PRIORITY, 143
 RDM\$BIND_APF_DEPTH, 115

RDM\$BIND_APF_ENABLED, 115, 162, 176
RDM\$BIND_BATCH_MAX, 200
RDM\$BIND_BUFFERS, 114, 161, 166, 171
RDM\$BIND_CBL_ENABLED, 119, 163
RDM\$BIND_CKPT_BLOCKS, 126, 127, 168
RDM\$BIND_CKPT_TIME, 127, 158, 168
RDM\$BIND_CKPT_TRANS_INTERVAL, 127, 168
RDM\$BIND_CLEAN_BUF_CNT, 116, 163, 200
RDM\$BIND_COMMIT_STALL, 123, 167
RDM\$BIND_DAPF_ENABLED, 115, 162, 177
RDM\$BIND_DBR_BUFFER_CNT, 154, 155
RDM\$BIND_DBR_PRIORITY, 143
RDM\$BIND_HOT_ABS_SUSPEND, 133
RDM\$BIND_HOT_CHECKPOINT, 132
RDM\$BIND_HOT_DATA_SYNC_MODE, 132
RDM\$BIND_HOT_NETWORK_TIMEOUT, 130, 131
RDM\$BIND_HRL_ENABLED, 119, 163
RDM\$BIND_LCS_CONNECT_TIMEOUT, 131
RDM\$BIND_LCS_PRIORITY, 143
RDM\$BIND_LOCK_TIMEOUT_INTERVAL, 118, 119, 163
RDM\$BIND_LRS_GAP_TIMEOUT, 131, 132
RDM\$BIND_LRS_GOVERNOR_ENABLED, 133
RDM\$BIND_LRS_PRIORITY, 143
RDM\$BIND_MAX_DBR_COUNT, 143
RDM\$BIND_RCACHE_INSERT_ENABLED, 135, 136, 172
RDM\$BIND_RCACHE_RCRL_COUNT, 136, 172
RDM\$BIND_RCS_CHECKPOINT, 158
RDM\$BIND_RCS_CKPT_BUFFER_CNT, 157, 158
RDM\$BIND_RCS_MAX_COLD, 159
RDM\$BIND_RCS_MIN_COLD, 159
RDM\$BIND_RCS_PRIORITY, 144
RDM\$BIND_RCS_SWEEP_INTERVAL, 158
RDM\$BIND_READY_AREA_SERIALLY, 119, 163
RDM\$BIND_RUJ_ALLOC_BLKCNT, 139, 140, 142, 168, 201
RDM\$BIND_RUJ_EXTEND_BLKCNT, 140, 168, 201
RDM\$BIND_SNAP_QUIET_POINT, 119, 163
RDM\$BIND_STATS_AJ_ARBS_PER_IO, 193
RDM\$BIND_STATS_AJ_BKGRD_ARB_RATIO, 194
RDM\$BIND_STATS_AJ_BLKS_PER_IO, 193
RDM\$BIND_STATS_AJ_CACHE_OVF_RATIO, 195
RDM\$BIND_STATS_AJ_NETWORK_BPS, 197
RDM\$BIND_STATS_AJ_SEC_TO_EXTEND, 190
RDM\$BIND_STATS_BTR_FETCH_DUP_RATIO, 218
RDM\$BIND_STATS_BTR_LEF_FETCH_RATIO, 219
RDM\$BIND_STATS_BTR_STORE_DUP_RATIO, 218
RDM\$BIND_STATS_DBR_DURATION, 204
RDM\$BIND_STATS_DBR_RATIO, 204
RDM\$BIND_STATS_FULL_BACKUP_INTRVL, 205
RDM\$BIND_STATS_GB_IO_SAVED_RATIO, 180, 181
RDM\$BIND_STATS_GB_POOL_HIT_RATIO, 179, 181
RDM\$BIND_STATS_HASH_FETCH_DUP_RATIO, 219
RDM\$BIND_STATS_HASH_STORE_DUP_RATIO, 219, 225, 226
RDM\$BIND_STATS_LB_PAGE_HIT_RATIO, 179, 181
RDM\$BIND_STATS_MAX_HASH_QUE_LEN, 223, 224, 227

RDM\$BIND_STATS_MAX_LOCK_STALL, 216
RDM\$BIND_STATS_MAX_RO_TX_DURATION, 185
RDM\$BIND_STATS_MAX_RO_TX_PERCENTILE, 185
RDM\$BIND_STATS_MAX_RW_TX_DURATION, 184
RDM\$BIND_STATS_MAX_RW_TX_PERCENTILE, 184
RDM\$BIND_STATS_MAX_TX_DURATION, 183
RDM\$BIND_STATS_MAX_TX_PERCENTILE, 183
RDM\$BIND_STATS_PAGES_CHECKED_RATIO, 208
RDM\$BIND_STATS_RECS_FETCHED_RATIO, 210
RDM\$BIND_STATS_RECS_STORED_RATIO, 210
RDM\$BIND_STATS_RUJ_EXTEND_RATIO, 201
RDM\$BIND_STATS_RUJ_SYNC_IO_RATIO, 200
RDM\$BIND_STATS_SPAM_PAG_FET_RATIO, 208
RDM\$BIND_STATS_SPAM_REC_STO_RATIO, 209
RDM\$BIND_STATS_VERB_SUCCESS_RATIO, 183
RDM\$BIND_TSN_INTERVAL, 127, 168
Read/Write Transaction duration percentile, 184
Read/Write Transaction duration threshold, 184
Read-Only Transaction duration percentile, 185
Read-Only Transaction duration threshold, 185
Records fetched threshold, 210
Records stored threshold, 210
RECS_FETCHED_RATIO, 210
RECS_STORED_RATIO, 210
RMU Dump Header, 196
Rootfile, 54, 94, 116, 120, 124, 128, 133, 137, 140, 144, 148, 152, 155, 159, 164, 169, 173
Row Cache Server, 30, 144
RTUPB, 73, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111
RUJ extend threshold, 201
RUJ_EXTEND_RATIO, 201
RUJ_SYNC_IO_RATIO, 200
Seconds to AIJ extend/switch, 190
SEQBLK, 61, 95, 98, 99, 101, 104, 107, 110
single-file database, 205
Snapshot File, 223
SPAM pages fetched threshold, 208
SPAM records stored threshold, 209
SPAM_PAG_FET_RATIO, 208
SPAM_REC_STO_RATIO, 209
Synchronous RUJ I/O threshold, 200
Transaction duration percentile, 183
Transaction duration threshold, 183
TSNBLK, 61, 64, 95, 96, 98, 99, 101, 102, 104, 105, 107, 108, 110, 111, 123, 167
VERB_SUCCESS_RATIO, 183
VLM, 5, 7