



Rdb Controller For Oracle Rdb (5.2a/5.3)

User's Guide

May 2003

ALI Database Consultants/Empirical Software

1151 Williams Dr.
Aiken, SC 29803 USA
Toll Free: (866) 257-8970
(803) 648-5931
Fax: (803) 641-0345
www.aliconsultants.com

DBAnalyzer for Oracle Rdb User's Guide for VAX/VMS

Version 1.0, February 1992
Version 2.0, May 1992
Version 2.0/2.2, October 1992
Version 2.5, August 1993
Version 2.6, October 1993
Version 3.0, September 1994

DBAnalyzer for Oracle Rdb User's Guide for Alpha AXP

Version 2.7, March 1994
Version 3.0, September 1994

DBAnalyzer for Oracle Rdb User's Guide (Combined Version)

Version 3.7, June 1995
Version 3.8, July 1995
Version 3.9, October 1996
Version 5.0, June 1997
Version 5.2, November 1998
Version 5.2a, October 2000
Version 5.2a/5.3, September 2001
Version 5.2a/5.3, June 2002
Version 5.2a/5.3, August 2002
Version 5.2a/5.3, September 2002
Version 5.2a/5.3, May 2003

Copyright © 1992, 1993, 1994 Information Systems Group, Inc.
Copyright © 1994-1996 The Database Solutions Company of Virginia
Copyright © 1997-1999 Empirical Software, Inc.

DBTune for Oracle Rdb User's Guide for VAX/VMS

Version 1.0, May 1992
Version 2.0/2.2, October 1992
Version 2.5, August 1993
Version 2.6, October 1993
Version 3.0 (FT2), July 1994
Version 3.0, September 1994

DBTune for Oracle Rdb User's Guide for Alpha AXP

Version 2.7, March 1994
Version 3.0, September 1994

DBTune for Oracle Rdb User's Guide (Combined Version)

Version 3.5, February 1995
Version 3.7, June 1995
Version 3.8, July 1995
Version 3.9, September 1996
Version 4.0, June 1997
Version 5.0, June 1997
Version 5.1, March 1998
Version 5.2a, October 2000
Version 5.2a/5.3, September 2001
Version 5.2a/5.3, June 2002
Version 5.2a/5.3, August 2002
Version 5.2a/5.3, September 2002
Version 5.2a/5.3, May 2003

Copyright © 1992, 1993, 1994 Information Systems Group, Inc.

Copyright © 1994-1997 The Database Solutions Company of Virginia
Copyright © 1997-1999 Empirical Software, Inc.

DBXAct for Oracle Rdb User's Guide for VAX/VMS

Version 1.0, December 1993
Version 1.2, January 1994
Version 1.3, February 1994
Version 1.4, March 1994
Version 3.0 (FT2), October 1994
Version 3.5, January 1995

DBXAct for Oracle Rdb User's Guide for Alpha AXP

Version 1.0, December 1993
Version 1.2, January 1994
Version 1.3, February 1994
Version 1.4, March 1994
Version 3.0 (FT2), October 1994
Version 3.5, January 1995

DBXAct for Oracle Rdb User's Guide (Combined Version)

Version 3.7, June 1995
Version 3.8, July 1995
Version 3.9, March 1996
Version 5.0, March 1998
Version 5.1, April 1999

Version 5.1/5.2, September 2001
Version 5.1/5.2, August 2002
Version 5.2a/5.3, September 2002
Version 5.2a/5.3, May 2003

Copyright © 1993, 1994 Information Systems Group, Inc.
Copyright © 1994-1997 The Database Solutions Company of Virginia
Copyright © 1997-1999 Empirical Software, Inc.

Copyright © 1999-2002 ALI Database Consultants.

All rights reserved. Printed in the USA

Information in this *Rdb Controller for Oracle Rdb User's Guide* is subject to change without notice and does not represent a commitment on the part of the vendor. The software described in this *Rdb Controller for Oracle Rdb User's Guide* is furnished under a license agreement and may be used or copied only in accordance with the terms of the agreement.

DBAnalyzer, DBTune, DBXAct, and C/S Control Center are trademarks of ALI Database Consultants/Empirical Software.

DEC, Rdb, VAX, VMS, Alpha AXP, AXP, and OpenVMS are trademarks of Digital Equipment Corporation.

NOTE: This manual applies to Rdb Controller 5.2 for Vax and Alpha AXP running Rdb 7.0, and to Rdb Controller 5.3 for Rdb 7.1 on Alpha AXP.

Acknowledgments

The development team at ALI would like to thank the following consultants and clients who have taken time to review our products and make suggestions for improving our products. Rdb Controller for Oracle Rdb has many features that were suggested by these people. We regret that we may not have been able to implement all suggestions for this release, but we try to shape each release as closely as possible to clients' needs and believe that the success of our products depends on our ability to incorporate improvements that users request.

COMMONWEALTH ALUMINUM

Jean Dickens

ITSD

Richard Dean

NOVOPHARM

Mohammed Aziz

SAFILO

Jeff Wolfe

SEA CONTAINERS (UK)

Kevin Gordon

VNU

Latesha Williams

Contents

Introduction	1
Rdb Controller for Oracle Rdb	1
Chapter 1.....	5
What is DBAnalyzer?.....	5
Getting Started	8
Installing DBAnalyzer	10
Using DBAnalyzer	13
Online Use of DBAnalyzer	14
Batch Use of DBAnalyzer	15
DBAnalyzer Keystrokes	19
DBAnalyzer Report Generation Keystrokes	20
DBAnalyzer Process.....	21
DBAnalyzer Reports	31
Sample DBAnalyzer Windows	39
Sample DBAnalyzer Report	45
DBAnalyzer Help	52
Chapter 2.....	61
What is DBTune for Rdb?.....	61
Getting Started	65
Installing DBTune	70
Using DBTune.....	73
Online Use of DBTune.....	73

Batch Use of DBTune	74
DBTune Parameters	76
DBTune Keystrokes.....	77
DBTune Process	78
DBTune Reports.....	128
DBTune Help	143
Chapter 3.....	145
What is DBXAct?.....	145
Getting Started	149
System Requirements.....	149
Installing DBXAct for Rdb	151
Running DBXAct	153
Explanation of DBXAct Variables	156
Using DBXAct	157
Analyzing Data Gathered with DBXAct.....	158
Generating and Understanding Reports	161
Using C/S Control Center with DBXAct.....	165
DBXAct Logical Names	168
Answers to Commonly Asked Questions	169
Appendix A	171
Appendix B	173
Appendix C	175
<i>Database Terms</i>	<i>175</i>
<i>Network Terms</i>	<i>181</i>
Index.....	184

Introduction

What is Rdb Controller for Oracle Rdb?

Rdb Controller for Oracle Rdb

Rdb Controller™ provides a solution for many Oracle Rdb environments by tuning the entire database and monitoring the system to identify problem areas before they impact the user community. Rdb Controller not only properly tunes the storage areas, indexes, table distribution, and I/O distribution, it performs these actions in record time. Tuning efforts can be reduced drastically, and performance improvements typically range from 20 to 50 percent.

Additionally, to ensure your Oracle Rdb environment is constantly performing as expected, Rdb Controller for Rdb provides you with the ability to continuously monitor Oracle Rdb and identify problem areas before they impact your user community.

Minimizing Fragmentation

By using the DBTune feature in Rdb Controller for Rdb, users can automatically implement a database tuning solution to resolve fragmentation problems. DBTune partitions the database via storage areas, intelligently distributes storage areas across multiple disk drives, and appropriately sizes storage areas to handle growth within the database. This enables you to proactively tune the database, so your database will perform optimally for months.

Rdb Controller for Rdb first analyzes your database activity to determine the busiest storage areas and disk drives, and identifies the growth level for the storage areas. It also analyzes where to distribute tables and indexes, so that user access times are minimized. Then, Rdb Controller generates all the necessary tuning scripts to tune the entire database, or a limited number of specific storage items. These scripts can take days to generate manually, where Rdb Controller for Rdb can complete them in minutes and ensure all dependencies and steps are in place with error checking for the tuning process.

Rdb Controller for Rdb provides a complete tuning process implemented with a DCL driver that runs the generated SQL scripts to perform the tuning. These scripts can be executed to tune your entire database, or simply the most important areas, depending on the maintenance window you have available. Empirical's customers typically experience a 20 to 50 percent improvement in application performance.

Choose the Right Indexes

Rdb Controller for Rdb reviews indexes and database activity to identify conversion considerations. It provides recommendations to change the index type to either SORTED or HASHED based on likely access approaches. The decision should be made based on the actual retrieval and concurrency needs of the database users. Rdb Controller implements all necessary work to modify SORTED and HASHED indexes and structure them for optimal performance. It generates a SQL procedure that will change the database and tune it accordingly.

Minimizing Locking Occurrences

Rdb Controller for Rdb can assist with these advanced tuning requirements in several ways: by locating hashing opportunities and providing record counts, record sizes, and index information to calculate appropriate page size and allocation.

Rdb Controller for Rdb can be used to implement and properly tune for clustering and shadowing.

I/O Distribution

Rdb Controller for Rdb automatically balances I/O for individual storage areas, or the entire database, to ensure your users have the quickest access to data. It identifies which tables are the most active, sizes their storage areas accordingly, distributes them across separate disk drives, and follows the rules of thumb for I/O listed below.

- Do not place all your indexes on a single disk drive when not using large RAID units with much built-in caching.
- Separate the table and index if you are not using the **PLACEMENT VIA** option.
- Keep RDA and SNP files on separate disks.
- Keep your busiest storage areas on separate disks.

Features of the Products within Rdb Controller for Rdb

DBAnalyzer for Rdb

- Scans a database and performs a summary analysis.
- Provides four views of a database: a **MACRO** view, a **MICRO-TABLE** view, a **MICRO-INDEX** view, and a **MICRO-STORAGE AREA** view.
- Provides overall and individual statistics for databases.
- Produces organized output that provides you with a comprehensive view of a database and gives some measure of its tuning status.

DBTune for Rdb

- Creates a Performance Analysis Data file, using the logical and physical data gathered from the database. The Performance Analysis Data file contains volume, workload, and environment information, which you may customize to include additional transaction activity.
- Maximizes Rdb performance and maintenance without requiring excessive effort on your part.
- Separates larger tables and the associated indexes into their own tablespaces for ease of monitoring and load distribution to improve performance.
- Tunes the entire physical structure of the database, typically resulting in a 50 percent improvement in application performance.

DBXAct for Rdb

- Monitors performance of the database and generates baseline statistics and reports on Rdb activity, making it possible for you to clearly understand the hundreds of data points available for monitoring database activity.
- Produces detailed statistical reports that allow you to perform precise tuning and optimization of the Rdb database.
- Establishes benchmarks and presents the results of the tuning changes in order to determine if the tuning efforts have been successful.
- DBXAct produces a database activity file that can be used by DBTune to evaluate the tuning needs of the database and create SQL scripts for its physical redesign. The database activity file provides you with the information needed to successfully tune and optimize the Rdb database.

Chapter 1

DBAnalyzer for Rdb

What is DBAnalyzer?

The DBAnalyzer procedure scans a database pointed to by the logical `FRENDRDBIMPORT`. During this scan, the procedure performs a summary analysis, which is then displayed on your screen or written to an output file. If called in **BATCH** mode, the procedure will create the output file only.

DBAnalyzer provides four views of a database: a **MACRO** view, a **MICRO-TABLE** view, a **MICRO-INDEX** view, and a **MICRO-STORAGE AREA** view.

The **MACRO** view provides overall statistics for the database—the five largest tables, the five indexes with the most duplicates, etc. The purpose of these **MACRO** windows is to highlight likely hot spots in the database that may require tuning attention.

The **MICRO-TABLE** view provides statistics for individual tables/views in the database. You may scroll (alphabetically) through the tables/views or jump directly to a particular table by entering a search string (to which the table names are compared).

The **MICRO-INDEX** view provides statistics for individual indexes in the database. Here again, you may scroll through indexes (based on their associated tables) or jump directly to a specific table by entering a table name search string.

The **MICRO-STORAGE AREA** view provides statistics for each storage area in the database. As with the other MICRO views, you may scroll through each storage area or jump directly to a specific storage area by entering the storage area search string.

Output can also be obtained from this procedure and several options are provided. The default name for the output file is DBANALYZR.LST, but you can override this by typing in a different file name when prompted. See the DBAnalyzer Reports section.

Note DBAnalyzer 5.2 works with Rdb 7.1. DBAnalyzer 5.2 works with Rdb 7.0. You must use DBAnalyzer 3.9 if you are running an earlier version of Rdb.

Tune and Complexity Ratings

The purpose of DBAnalyzer is to produce organized output that provides you with a comprehensive view of a database and give some measure of its tuning status. Two indicators of a database's tuning status are the TUNE RATING and the COMPLEXITY RATING values. The TUNE RATING is a measure of the extent to which existing tuning options have been utilized in the database.

The TUNE RATING does not directly indicate the percentage increase in performance that can be achieved by tuning a database. Rather, the impact of tuning is more directly related to the complexity of the database. Therefore, the TUNE RATING in conjunction with the COMPLEXITY RATING is a more complete measure of a database's tuning status and the increase in performance that can be achieved by further tuning.

The INTEGRITY RATING of the database is a measure of how effectively the database has implemented constraints. The measurements that comprise this INTEGRITY RATING are:

1. Columnar Integrity Rating
2. Referential Integrity Rating
3. Referential Efficiency Rating

These individual components can be viewed on MACRO screen 14 or by displaying and/or printing the DBAnalyzer narrative. This narrative is available online by pressing the **Find** key, or in batch as part of the DBAnalyzer report. The narrative explains the meanings of the various graphs and ratings that are presented in DBAnalyzer.

Getting Started

This section provides the information needed for you to quickly install and run DBAnalyzer. It also includes system requirements necessary to run the application.

DEC VAX/OpenVMS

Recommended **minimum AUTHORIZE** settings for a DBAnalyzer user account. (Medium and large databases may need to increase these numbers.)

Username:	USER	Owner:	USER
Account:	USER	UIC:	[Group, Member]
CLI:	DCL	Tables:	DCLTABLES
Default:	<disk>:[dir]		
LGIMD:	LOGIN		
Login Flags:			
Primary Days:	Mon. Tues. Wed. Thurs. Fri.		
Secondary Days:	Sat. Sun.		
No access restrictions			
Expiration:	(none)	Pwdminimum:	0
Pwdlifetime:	(none)	Pwdchange:	
Last Login:		Login Fails:	0
Maxjobs:	0	Fillm:	512
Maxacctjobs:	0	Shrfillm:	100
Maxdetach:	0	BIOfm:	100
Prclm:	4	DIOfm:	100
Prio:	4	ASTlm:	113
Queprio:	0	TQElm:	10
CPU:	(none)	Enqlm:	8000
Authorized Privileges:	GROUP TMPMBX NETMBX	Bytln:	90000
Default Privileges:	GROUP MPMBX NETMBX	Pbytlm:	0
		Jtquota:	1024
		Wsdef:	1024
		Wsquo:	1024
		Wsexent:	1024
		Pqlfquo:	80000

Warning If these minimums are not in place when DBAnalyzer is executed, the analysis may fail!

DEC AXP/OpenVMS

Recommended **minimum AUTHORIZE** settings for a DBAnalyzer user account. (Medium and large databases may need to increase these numbers.)

Username:	USER	Owner:	USER
Account:	USER	UIC:	[Group, Member]
CLI:	DCL	Tables:	DCLTABLES
Default:	<disk>:[dir]		
LGIMD:	LOGIN		
Login Flags:			
Primary Days:	Mon. Tues. Wed. Thurs. Fri.		
Secondary Days:	Sat. Sun.		
No access restrictions			
Expiration:	(none)	Pwdminimum:	0
Pwdlifetime:	(none)	Pwdchange:	0
Last Login:			
Maxjobs:	0	Fillm:	512
Maxacctjobs:	0	Shrfillm:	0
Maxdetach:	0	BIOIm:	150
Prclm:	4	DIOIm:	150
Prio:	4	ASTIm:	250
Queprio:	0	TQElm:	10
CPU:	(none)	Enqlm:	8000
Authorized Privileges:	GROUP TMPMBX NETMBX		
Default Privileges:	GROUP MPMBX NETMBX		
		Bytlm:	90000
		Pbytlm:	0
		Jtquota:	1024
		Wsdef:	2000
		Wsquo:	4096
		Wsextent:	16384
		Pqlfquo:	80000

Warning If these minimums are not in place when DBAnalyzer is executed, the analysis may fail!

Installing DBAnalyzer

► **To install DBAnalyzer for Rdb from a tape drive:**

1. Back up your system disk (optional).
2. Log in under the SYSTEM account.
3. Put the DBAnalyzer distribution tape in the tape drive.
4. Type in the following command to invoke the VMS install facility to install DBAnalyzer on your system:

For VAX/VMS

```
$ @SYS$UPDATE:VMSINSTAL DBARDBVMS052 <<tape-drive>>:
```

For Alpha AXP

```
$ @SYS$UPDATE:VMSINSTAL DBARDBAXP052 <<tape-drive>>:
```

where <<tape-drive>> is the name of the device where the DBAnalyzer distribution tape has been mounted (e.g., MUA6:).

Note DBAnalyzer V5.2/5.3 should NOT be installed in the same directory with other versions of DBAnalyzer or any other product from ALI (i.e., DBTune).

5. After the VMS install has completed, place the following lines into the system startup command file

(SYS\$MANAGER:SYSTARTUP_VMS.COM) so that required logicals are set up when the system is rebooted:

```
$ DEFINE/SYSTEM/EXEC FRENDDBA$HOME <<disk>>: [dir]
```

```
$ DEFINE/SYSTEM/EXEC FRENDDBA$SCRATCH  
<<disk>>: [dir.scratch]
```

where <<disk>> and [dir] are the disk and directory to which DBAnalyzer was installed (e.g.,
`1DUA1: [DBARDBVMS52.SCRATCH]` or
`1DUA1: [DBARDBAXP52.SCRATCH]`).

6. Now, to obtain a license pak for DBAnalyzer, type in the following commands:

```
$ SET DEFAULT FRENDSDBA$HOME
```

```
$ EDIT DBANLZR.LICENSE
```

For each node (“machine”) on which you wish to run DBAnalyzer:

- Replace “your node name” with the node name of the machine on which you have installed DBAnalyzer. To get this information, type:

```
$ WRITE SYS$OUTPUT F$GETSYI (“nodename”)
```

- If the “operating system” value supplied with your license is not accurate for your system, replace it with the output generated from the following command:

```
$ WRITE SYS$OUTPUT F$GETSYI (“node_swtype”)
```

- Replace “your company name” with your company’s full name
- Exit and save the file

To obtain the appropriate registration ID for each machine entered, call ALI at (866) 257-8970 [or (803) 648-5931], or fax a copy of the altered DBANLZR.LICENSE file to (803) 641-0345. International clients may also obtain registration IDs/support through their local distributor’s office.

► **To install DBAnalyzer for Rdb from a CD-ROM:**

1. Mount the CD using a command like

```
$ MOUNT/OVER=ID <cd_device>:
```

2. Install the product with the command

```
$ @SYS$UPDATE:VMSINSTAL <product_name>  
<cd_device>: [INSTALL]
```

where <product_name> is the product you wish to install.

For example:

```
$ @sys$update:vmsinstal DBARDBVMS052 dka400: [INSTALL]
```

Using DBAnalyzer

DBAnalyzer may be executed either ONLINE or in BATCH. If executed ONLINE, DBAnalyzer is invoked by the command file DBA.COM. If executed in BATCH, DBAnalyzer is invoked by submitting the command file DBA_BATCH.COM. Before using DBAnalyzer, however, check with your system manager—VMS symbols may have been set up to facilitate use of this facility:

e.g., `DBA:==@FRIENDDBAHOME:DBA.COM`

There is one logical that is required to be assigned in order to execute DBAnalyzer in either ONLINE or BATCH mode:

`FRIENDRDBIMPORT`

DBAnalyzer scans the Rdb database pointed to by the logical FRIEND\$RDB\$IMPORT and expects this logical to be assigned prior to execution. You must also have privilege to access the database that you chose to analyze.

Online Use of DBAnalyzer

If executed ONLINE, the DBAnalyzer utility will check the assignment of FRENDSRDB\$IMPORT. If not assigned, you will be prompted for the location and name of the Rdb database to be analyzed. If the logical is already pointing to a database, however, you will be asked if you would like to point to a different database before continuing.

► **To invoke DBAnalyzer online:**

- If the symbol “DBA” has been created, type the following:

```
$ DBA
```

You will receive the following prompt:

The logical FRENDSRDB\$IMPORT, which must be assigned to the database that you wish to analyze, is currently unassigned. If you wish to continue, type in the disk, directory, and name of the database you wish to analyze or press Ctrl-Z to quit.

Example: DISK1 : [MYDATA.RDB] PERSONNEL.RDB

Specify a DATABASE: DISK5 : [ACCTNG.RDB] INVOICE.RDB

- If no symbol exists for DBAnalyzer, type the following:

```
$ @FRENDSDBA$HOME:DBA.COM
```

Batch Use of DBAnalyzer

If executed in BATCH, DBAnalyzer expects the FRENDSRDB\$IMPORT logical to have been assigned prior to execution—you will not be prompted. To this end, a command file has been provided to allow assignment of FRENDSRDB\$IMPORT. This command file—DBA_BATCH.COM—can be edited to select the database that will be scanned.

```

$!-----
$!   DBA_BATCH.COM
$!   - Command file to submit DBAnalyzer in BATCH mode...
$!
$!   Invoke this file with the command:
$!   $ @FRENDSDBA$HOME:DBA_BATCH
$!-----
$!
$!   This command procedure will submit itself to batch.
$!
$!   if p1 .eqs. "" .or. p1 .nes. "BATCH"
$!       then
$!
$!   Change the /name="" qualifier to specify a different name for the job.
$!
$!       cur_def = f$environment("DEFAULT")
$!       vfl = f$verify(0)
$!       set verify
$!       submit-
$!           /log-
$!           /noprint-
$!           /name="DBA in Batch"-
$!           /parameters=("BATCH","'cur_def'") -
$!           FRENDSDBA$HOME:DBA_BATCH.COM
$!       vfl = f$verify(vfl)
$!       exit
$!   endif
$!-----
$!   Change the following ASSIGN statement to point the database you wish to
$!   analyze and uncomment it by removing the "!" ...
$!
$!   ASSIGN "disk1:[directory]database_name" FRENDSRDB$IMPORT
$!
$!   Change the following SET PROC/NAME= to assign a different
$!   process name and uncomment it by removing the "!" ...
$!
$!   SET PROC/NAME="DBA in Batch"
$!
$!   Change the following SET DEFAULT to change the default
$!   directory where the report will be generated. Otherwise,
$!   output will be generated the user's current directory at the
$!   time the file was submitted.
$!
$!   SET DEFAULT 'p2'
$!-----
$!   @FRENDSDBA$HOME:DBA.COM

```

► **To invoke DBAnalyzer in batch:**

- Type the following:

```
$ @FREND$DBA$HOME:DBA_BATCH.COM
```

You can optionally set up a customized report parameter file that can be used to govern the format of the output report created when DBAnalyzer is executed in BATCH. You can specify a customized report parameter file for DBAnalyzer by assigning the logical FREND\$DBA\$PARAM_FILE and pointing it to a valid parameter file. This assignment can be made in the DBA_BATCH.COM file mentioned previously.

On the following page is an example of such a parameter file with valid entries and default values listed.

Example of an optional DBAnalyzer report parameter file

dba.report.name: MYFILE.RPT	Defaults [DBANALYZER.LST]
* * Selection Criteria *	
dba.select.full: YES or NO	[YES]
dba.select.brief: YES or NO	[NO]
dba.select.domains: YES or NO	[YES]
dba.select.narrative: YES or NO	[YES]
dba.select.storage_areas: YES or NO	[YES]
dba.select.tables: YES or NO	[YES]
dba.select.views: YES or NO	[YES]
* * Printer Criteria *	
dba.print_report: YES or NO	[NO]
dba.printer: SYS\$PRINT or blank	[]
dba.printer_options: {/FORM=LETTER16/HEADER}	[]
* * Domain Sort Order Criteria *	
dba.domain.sort.order: DOMAIN, COLUMN, TABLE	[DOMAIN]
* * Storage Input Criteria *	
dba.storage_source: ALL or FILE	[ALL]
* * Table Detail Criteria *	
dba.table_source: ALL or FILE	[ALL]
dba.table_option: FULL or BRIEF	[FULL]
dba.table_columns: YES or NO	[YES]
dba.table_indices: YES or NO	[YES]
dba.table_indices_option: FULL or BRIEF	[FULL]
* * View Detail Criteria *	
dba.view_source: ALL or FILE	[ALL]
dba.view_option: FULL or BRIEF	[FULL]

DBAnalyzer Keystrokes

The following keystrokes may be used when executing DBAnalyzer ONLINE using a VT220 (or higher) terminal interface:

<input type="text" value="Help"/>	or	<input type="text" value="PF2"/>	Receive help for the current DBAnalyzer context. Repeated execution will allow viewing of the entire help menu
<input type="text" value="Esc"/>	<input type="text" value="PF3"/>	<input type="text" value="PF4"/>	Exit DBAnalyzer or return to MACRO mode
<input type="text" value="Do"/>			Rescan the Database and recalculate analysis information
<input type="text" value="Select"/>			Toggle from MACRO mode to MICRO-TABLE mode to MICRO-INDEX mode to MICRO-STORAGE AREA mode
<input type="text" value="Next Screen"/>			Scroll to next window
<input type="text" value="Prev Screen"/>			Scroll to previous window
<input type="text" value="Find"/>			Display Narrative [MACRO window] Display View [MICRO table window]
<input type="text" value="F20"/>			Write report to an output file
<input type="text" value="Ctrl"/>		<input type="text" value="W"/>	Refresh screen display

DBAnalyzer Report Generation Keystrokes

The following keystrokes may be used when executing the DBAnalyzer report generation functions using a VT220 (or higher) terminal interface. (See also the DBAnalyzer Reports section on page 31.)

PF3	Exit back to SCAN mode
PF4	Backup one window (not implemented in this release)
Select	Highlight an option for execution
Return	Execute options
↑	Move up in selection window
↓	Move down in selection window
Next Screen	Page down in a selection window
Prev Screen	Page up in a selection window
Help	Help on Report Generation

DBAnalyzer Process

The following information is available from DBAnalyzer:

Rdb Statistics

Total number of Tables, Indexes, Storage Areas, etc., as well as database parameters such as Global Buffers, Buffer size, Number of Users, etc.

Complexity Rating

The **COMPLEXITY RATING** is a weighted measure of the size and complexity of a database. The **NARRATIVE** analysis provides a relative description of its significance.

Tune Rating

The **TUNE RATING** is a measure of the extent to which existing tuning options have been utilized in the database. A database with a **TUNE RATING** of between 80 to 100 is considered one that has taken advantage of available tuning options. A rating of zero suggests that no tuning (other than the provided defaults) has been performed.

The **TUNE RATING** does not directly indicate the percentage increase in performance that can be achieved by tuning a database. Rather, the impact of tuning is more directly related to the complexity of the database. Therefore, the **TUNE RATING** in conjunction with the **COMPLEXITY RATING** is a more complete measure of a database's tuning status and the increase in performance that can be achieved by further tuning.

The **TUNE RATING** indicates how well the physical design supports the current logical design. The rating is a ratio of storage area utilization (factored for allocation efficiency) and index design compared to the database complexity. The **NARRATIVE** analysis provides an explanation of the performance impact of the **TUNE RATING** for the given **COMPLEXITY**.

Integrity Rating

The **INTEGRITY RATING** of the database is a measure of how effectively the database has implemented constraints. The measurements that comprise this **INTEGRITY RATING** are:

1. Columnar Integrity Rating
2. Referential Integrity Rating
3. Referential Efficiency Rating

A database with an **INTEGRITY RATING** between 75 and 100 is considered to be one that has taken advantage of Rdb's constraint mechanisms to ensure the completeness and integrity of data. DBAnalyzer's **NARRATIVE** analysis provides a detailed description for a particular database.

Storage Area Allocation

The **STORAGE AREA ALLOCATION** consists of two graphs that show the percentage of blocks that have been created due to extensions. The graphs indicate the percentage for both the RDA and SNP files. These ratios are used to factor the storage area utilization portion of the **TUNE RATING**. That is, the higher the allocation percentages, the lower the tune rating.

Hashed Index Percentage

The **HASHED INDEX PERCENTAGE** is the ratio of hashed indexes to the total number of indexes in the database. If you have ten indexes of which two are hashed, then the **HASHED INDEX PERCENTAGE** would be 20 percent.

HASH-to-SORT/SORT-to-HASH Ratio

HASH-to-SORT/SORT-to-HASH RATIO is the ratio of indexes that have ADVISOR recommendations to be changed from either hashed to sorted or sorted to hashed. If you have ten indexes of which two sorted indexes are recommended to be hashed and one hashed index is recommended to be sorted, then the HASH-to-SORT/SORT-to-HASH RATIO would be 30 percent.

Macro View

The MACRO VIEW consists of 14 windows to provide statistics on significant items from various aspects of the database. Each of the 14 windows is described below. These windows are not intended to be a comprehensive list of all of the items that should be tuned. Rather, they are listed as ones that are likely to have significant tuning implications relative to other items in the database. Remember that overall tuning benefits available are relative. The expected tuning benefits are impacted by numerous factors. DBAnalyzer cannot predict exact improvements. You can, however, expect tuning benefits to have a direct relation to the COMPLEXITY RATING and an indirect relation to the TUNE RATING. That is, the higher the COMPLEXITY RATING and the lower the TUNE RATING, the more the database performance may be improved through tuning.

- Macro Window 1:
Tables with the Highest Record Counts

The five tables with the highest record counts (cardinality) are listed in descending order. These tables are likely candidates for special tuning attention.

- Macro Window 2:
Indexes with the Most Duplicates

The five indexes with the highest average number of duplicates are listed in descending order. Indexes with numerous duplicates can unnecessarily increase append, modification, and deletion processing time.

- Macro Window 3:
Tables with the Most Columns

The five tables with the most columns are listed in descending order. These tables may indicate the presence of redundant fields or fields that can be more efficiently stored in other tables. Rdb has clustering and shadowing capabilities that can potentially improve performance by greater normalization rather than reduced normalization. These capabilities often apply to one-to-many relationships.

- Macro Window 4:
Tables with the Largest Record Size

The five tables whose columns require the most bytes to store a 'full' record. Although Rdb stores each column as its own item, proper page sizing techniques make use of this effective record size.

- Macro Window 5:
SORTED Indexes Recommended to be HASHED

The first five SORTED indexes that DBAnalyzer encounters that it determines may be a candidate for Hashing. HASHED indexes may be used to decrease the number of I/Os required to retrieve a record and increase concurrency by reducing the number of required locks. Additional benefits may accrue when hashed indexes are used in conjunction with CLUSTERING and SHADOWING.

- Macro Window 6:
HASHED Indexes Recommended to be SORTED

The first five HASHED indexes that DBAnalyzer encounters that it determines may be a candidate for Sorting. SORTED indexes may be necessary to decrease the number of I/Os required for retrieving a range of records. An example of this type of retrieval is gathering all employees with a last name starting with 'SM'. Sorted indexes may reduce concurrency for records whose index values fall into the same index NODE. The effect of such occurrences depends on the values being accessed and the mode in which they are accessed.

- Macro Window 7:
Tables with the Most Indexes

The five tables with the most indexes are listed in descending order. The indexes for these tables should be reviewed to see if they are being used. Each index must be maintained whenever records are added, modified, or deleted. Indexes require disk space, and they use resources for their maintenance. Indexes that are not used should be removed.

- Macro Window 8:
Non-Indexed Tables with the Most Records

The five tables that do not have any indexes are listed in descending order by record count. Retrieval of data from these tables requires that all of the records in the table be retrieved. Performance may be improved if an index can be created that will allow the Rdb-Optimizer to select a smaller set of these tables' records.

- Macro Window 9:
Storage Areas with the Most Mapped Items

The storage areas, into which the most items are mapped, are listed in descending order by mapped item count. The mapped item can be a table, sorted index, or hashed index. These areas may not perform as well as other areas since multiple types of data must be stored in the same area. Thus, performance may not be as efficient as possible. The performance impact increases as the number of records increases for each of the items.

- Macro Window 10:
Storage Areas with the Most File Extensions

The five RDA files that have extended the most times are listed in descending order. File extensions occur when there is insufficient space to hold new data. File extensions may indicate insufficient pages, incorrect page sizes, or both.

- Macro Window 11:
RDA Files with the Most Extension Blocks

The five RDA files that have the most extension blocks are listed in descending order. File extensions occur when there is insufficient space to hold new data. File extensions may indicate insufficient pages, incorrect page sizes, or both.

- Macro Window 12:
SNP Files with the Most Extension Blocks

The five SNP files that have the most extension blocks are listed in descending order. File extensions occur when there is insufficient space to hold new data. SNP files are used to hold “snapshots” of pages that are locked, so that READ transactions can access these pages. Extended space indicates that the snapshot activity requires more space than that initially allocated.

- Macro Window 13:
Database Storage Area Extension Summary

The storage area extension summary displays the total number of RDA and SNP blocks that were initially allocated and how many are currently in use. Additionally, the total number of RDA extensions is available.

- Macro Window 14:
Integrity Rating Constituent Components

The integrity rating constituent components display a graphical and numeric representation of the three components of the database’s composite integrity rating.

Micro-Table View

The MICRO-TABLE VIEW is presented by pressing the **Select** key while in MACRO MODE. When you select **MICRO-TABLE MODE**, you will see the first non-system table in the database, in alphabetic order by table name. You may press the **Next Screen** key to progress through the database tables alphabetically. You may also type a table name or a portion of a table name and press **Return**. DBAnalyzer will locate the first table name that is greater than or equal to the entered letters. When DBAnalyzer locates the end of the list of tables, it returns to MACRO MODE.

MICRO Mode: Rdb Tables/Views	
Table: CUST_CONTACTS	
Cardinality (Record Count):	342
Number of Columns:	6
Number of Bytes (Record Size)	87
Storage Area: CUST_CONTACT_AREA	
Number of Indexes:	1
1 st Unique Index: CONTACT_TYPE	

The Micro-Table window presents critical tuning information for each non-system table in the database. The number of columns and their byte allocations are used to calculate proper page sizes, and the number of records indicates how many pages to allocate. Proper page size and allocation prevent fragmentation within the storage area and prevent fragmenting extents for the storage area file itself. The first storage area used by the table is listed. RDB\$SYSTEM is the default storage area. Those tables with the most records should have priority when deciding which tables to move into their own storage area.

The number of indexes is a quick check on possible database performance problems. A high number of indexes may be wasteful and inefficient, but no indexes may require extra I/O, especially for tables with a large number of records and frequent retrievals of a subset of the table records. The proper number of indexes is a trade-off in the processing time to maintain an index versus the savings from index-based retrieval.

VIEWS may be displayed in the MICRO-Table view by using the **Find** key. The display shows the columns and the selection that comprise the view.

Note The DBAnalyzer report shows all storage areas and indexes for a table. See the DBAnalyzer Reports section for more information.

Micro-Index View

The MICRO-INDEX VIEW is presented by pressing the **Select** key while in MICRO-TABLE MODE. When you select **MICRO-INDEX MODE**, you will see the first index for the first non-system table in the database, in alphabetic order by table and the table's index. You may press the **Next Screen** key to progress through the database indexes. You may also type a table name or a portion of a table name and press **Return**. DBAnalyzer will locate the first index for the first table name that is greater than or equal to the entered letters. When DBAnalyzer locates the end of the list of indexes it returns to MACRO MODE.

```

MICRO Mode:  Rdb Indexes

Table       :  CUST_CONTACTS
Index (001) :  CONTACT_TYPE
Index Type  :  SORTED, NON-UNIQUE
Avg Dups    :  27
Storage Area : CONTACT_TYPE_IDX_AREA
Index columns: 1
1st column  :  CUST_CONTACT_TYPE

```

The Micro-Index window presents information about a selected index. It shows the table to which the index belongs, whether it is Sorted or Hashed, the first storage area it is mapped into, the number of columns, and the first column in the index. The window allows a DBA to quickly see what indexes are available for record retrieval.

Those indexes with a high number of average duplicates are ones that should be reviewed. Unless specific techniques such as clustering and shadowing of records are being employed, indexes with a high number of average duplicates may be creating extra processing. Adding columns to the index may reduce the number of average duplicates, index processing and improve performance.

Note The DBAnalyzer report shows all storage areas and columns for an index. See the DBAnalyzer Reports section for more information.

Micro-Storage Area View

The MICRO-STORAGE AREA VIEW is presented by pressing the **Select** key while in MICRO-INDEX MODE. When you select **MICRO-STORAGE AREA MODE**, you will see the first storage area in the database, in alphabetical order by storage area name. You may press the **Next Screen** key to progress through the database storage areas. You may also type a storage area name or a portion of a storage area name and press **Return**. DBAnalyzer will locate the first storage area that is greater than or equal to the entered letters. When DBAnalyzer locates the end of the list of storage areas, it returns to MACRO Mode.

```

MICRO Mode:  Rdb Storage Areas

Storage Area: COMP_NAME_INDEX_AREA
Page Size:   4 blocks, Format: UNIFORM
Number of Extents:                               1
RDA Extensions (blocks):                          194
SNP Extensions (blocks):                           0
-----Stored Elements-----
          Tables:  0   Sorted:  1   Hashed:  0

```

The Micro-Storage Area window presents information on the initial and current pages for both the RDA and SNP files. Additionally, the number of each type of mapped item is displayed. Large storage areas may achieve better performance if only “related” items are clustered together. For example, clustering the HASH index and its table may achieve single I/O performance when the HASH key is used to retrieve the table.

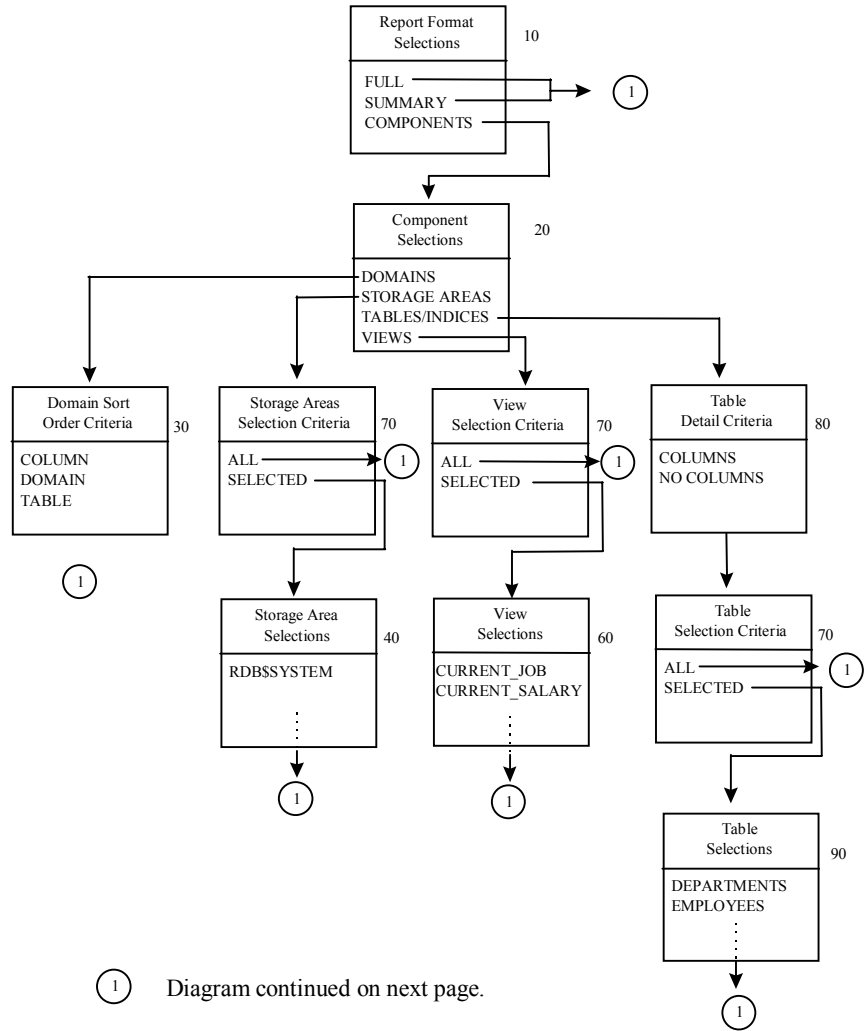
DBAnalyzer Reports

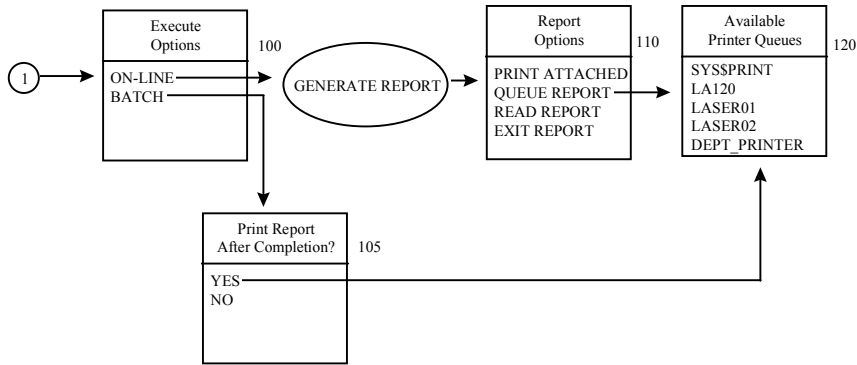
Report generation for DBAnalyzer has been modified in order to give you more flexibility in the type of report you generate.

DBAnalyzer provides a window environment for you to select the flavor of your report. **FULL** and **SUMMARY** report options are available, and with the window interface, you may specify a report to provide information on **DOMAINS**, **STORAGE AREAS**, **TABLES**, and **VIEWS**. These may be done exclusively of one another or combined together in one report.

In addition, you can generate variations on the **/Table**, **/Index**, and **/Storage Area** reports. The **STORAGE AREAS** report is tabular in orientation and reports **TABLES** and **VIEWS** independently of one another to more logically segregate your data needs. The **DOMAIN** report will provide you with a means of viewing your columns/tables in any one of three useful fashions.

Note Almost all portions of the report have 132 columns. Be advised that before printing, your printers should be set for compressed print. Before generating your reports, please review the next page and familiarize yourself with the map of windows that will guide your report generation.





DBAnalyzer REPORT DESCRIPTION	
Window Number	Description
10	Report Format Selections
20	Report Components
30	Domain Sort Order
40	Select Specific STORAGE AREAS
60	Select Specific VIEWS
70	Specify either ALL or SELECTED
80	Specify either COLUMNS or NO COLUMNS
90	Select Specific TABLES
100	Select Execution Options, ONLINE or BATCH
105	If in batch, print report when completed
110	Select Report Review Options
120	Select from Available Printer Queues

See the **DBAnalyzer Help** section on page 52 for more information on the various output options.

All of the above information is available via the DBAnalyzer report. This report can be accessed when the **F20** key is pressed or the **OUTPUT** menu option is selected. When the **OUTPUT** option is chosen, you will be prompted to enter a report file name. The default is **DBANALYZR.LST**. By default, the report file will be located in your current default directory.

After entering the filename, you will be presented with a window of options used to determine the format of the report.

Report Format Selections
FULL SUMMARY COMPONENTS

If you select either **FULL** or **SUMMARY**, you will be prompted to choose whether to run the report online or in batch. See **Execution Options** on page 37 for more information. If you select **COMPONENTS**, you will be presented with a window of items to be included on the report.

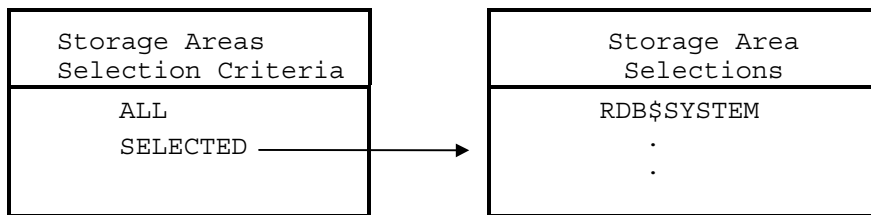
Component Selections
DOMAINS STORAGE AREAS TABLES/INDEXES VIEWS

Use the **UP** and **DOWN** arrow keys to move through the window(s). Press **Select** to highlight an option to be included on the report. More than one component can be selected. Press **Return** when the selection process is completed.

If you elect to include **DOMAINS** in the report listing, DBAnalyzer will present a window to determine the sort order for domains on the final report. Arrow between the three options and press **Return** for the one desired.

Domain Sort Order Criteria
COLUMN DOMAIN TABLE

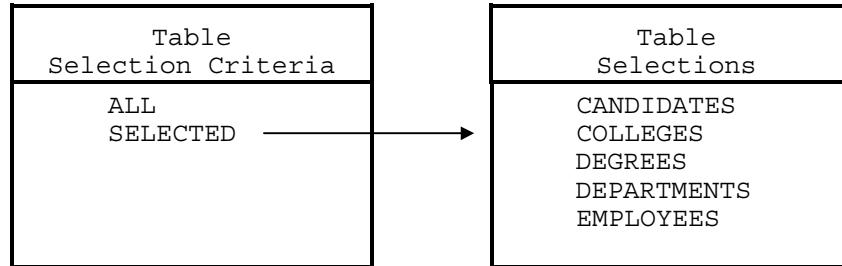
If you elect to include **STORAGE AREAS** in the report listing, then DBAnalyzer will present a window enabling you to choose whether to include **ALL** storage areas or only **SELECTED** areas. If you indicate **Selected Areas**, then a window will be presented to indicate which storage areas to include in the report. You can press **Select** for one or more storage areas on which to report. Pressing **Return** will complete the selection process.



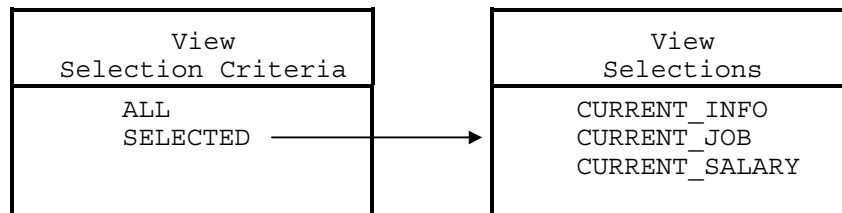
If you elected to include **TABLES/INDEXES** in the report, DBAnalyzer will present a window to select the level of detail for the **TABLE** report.

Table Detail Criteria
COLUMNS NO COLUMNS

You will then be asked to select the tables on which to report. You may specify either **ALL** tables or **SELECTED**. If the **SELECTED** option is chosen, a window listing all the tables in the database will be presented. You can then arrow through the tables, highlighting the ones to be reported with the **Select** key. Pressing **Return** will complete the selection process.

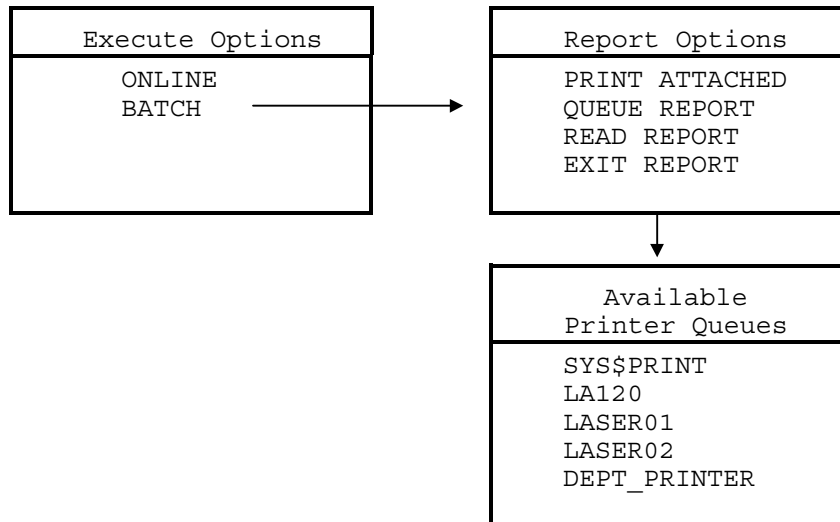


If you elect to include views in the report, you will be asked to select the views for reporting. You may specify either **ALL** views or **SELECTED**. If the **SELECTED** option is chosen, a window with all of the views will be presented. You can arrow through the views, highlighting the ones to be reported. Pressing **Return** will complete the selection process.



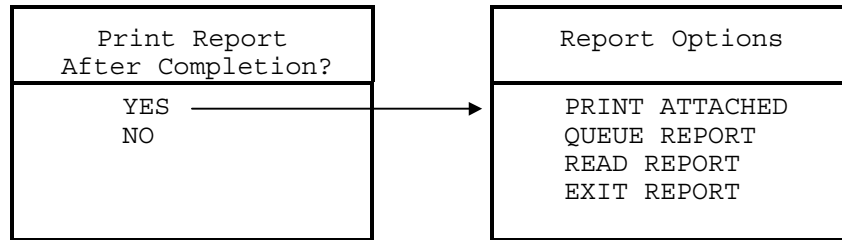
Execution Options:

Once the report specification is completed, you will be asked if the report is to be run **ONLINE** or in **BATCH**. A window will be used to present you with this option. If executed online, you will be presented with a window to select an available printer queue when the report completes.



Only one option can be selected at any one time from this menu. As long as you are selecting report options, this menu will continue to be presented.

If you elect to execute the report in batch, you will be asked whether or not to automatically print the report upon completion. If yes, you will be prompted for a queue to be used for the output. The queues are selected from among the print queues available on the system.



See the section on **BATCH Use of DBAnalyzer** on page 15 for information on creating the DBAnalyzer report in BATCH.

Sample DBAnalyzer Windows

MACRO MODE 1: (Tables with the Highest Record Counts)

Tables with the Highest Record Counts Macro Window 1	
SALARY_HISTORY	729
JOB_HISTORY	274
DEGREES	165
EMPLOYEES	100
DEPARTMENTS	26

MACRO MODE 2: (Indexes with the Most Duplicates)

Indexes with the Most Duplicates Macro Window 2	
DEG_COLLEGE_CODE	12
SH_EMPLOYEE_ID	7
JH_EMPLOYEE_ID	2
DEG_EMP_ID	1
EMP_LAST_NAME	1

MACRO MODE 3: (Tables with the Most Columns)

Tables with the Most Columns Macro Window 3	
EMPLOYEES	12
JOB_HISTORY	6
COLLEGES	5
DEGREES	5
DEPARTMENTS	5

MACRO MODE 4: (Tables with the Largest Record Size)

Tables with the Largest Record Size Macro Window 4	
CANDIDATES	280
EMPLOYEES	112
COLLEGES	56
DEPARTMENTS	47
JOB_HISTORY	34

MACRO MODE 5: (SORTED Indexes Recommended to be HASHED)

SORTED Indexes Recommended to be HASHED Macro Window 5	
Index / Table	
DEG_EMP_ID	/ DEGREES
EMP_EMPLOYEE_ID	/ EMPLOYEES
JH_EMPLOYEE_ID	/ JOB_HISTORY
SH_EMPLOYEE_ID	/ SALARY_HISTORY

MACRO MODE 6: (HASHED Indexes Recommended to be SORTED)

HASHED Indexes Recommended to be SORTED Macro Window 6	
*** None were found ***	

MACRO MODE 7: (Tables with the Most Indexes)

Tables with the Most Indexes	
Macro Window 7	
INVOICE_INFORMATION	5
EXPENSES	3
INVOICE_LINE_ITEMS	3
MEAL_EXPENSES	3
TIME_WEEK_HEADER	3

MACRO MODE 8: (Non-Indexed Tables with the Most Records)

Non-Indexed Tables with the Most Records	
Macro Window 8	
*** None were found ***	

MACRO MODE 9: (Storage Areas with the Most Mapped Items)

Storage Areas with the Most Mapped Items	
Macro Window 9	
CONTRACT_INDEX_C1	4
PRODUCTS_INDEX_C2	4
CONTRACT_ITEMS_C3	3
CONTRACT_AREA	2
CONTRACT_ITEMS_C5	2

MACRO MODE 10: (Storage Areas with the Most File Extensions)

Storage Areas with the Most File Extensions	
Macro Window 10	
TIME_TABLE	8
RDB\$SYSTEM	4
TIME_INDEX	2
CONTRACT_INDEX_C1	1
CONTRACT_INDEX_C2	1

MACRO MODE 11: (RDA Files with the Most Extension Blocks)

RDA Files with the Most Extension Blocks	
Macro Window 11	
RDB\$SYSTEM	3,608
TR_TABLE	3,396
PRODUCT_INDEX_P5	672
TIME_INDEX	597
PAYROLL_INDEX	576





MACRO MODE 12: (SNP Files with the Most Extension Blocks)

SNP Files with the Most Extension Blocks	
Macro Window 12	
CONTRACT_INDEX_C6	384
INVOICE_INDEX_i1	384
PAYROLL_INDEX_R5	384
TIME_TABLE	384
TIME_ITEM_TABLE	384

MACRO MODE 13: (Database Storage Area Extension Summary)

Database Storage Area Extension Summary	
Macro Window 13	
Total RDA Extensions:	59
Initial RDA Alloc (blocks)	5,635
Current RDA Alloc (blocks)	29,389
Initial SNP Alloc (blocks)	656
Current SNP Alloc (blocks)	4,952

MACRO MODE 14: (Database Integrity Ratings)

	DB Integrity (200.00)
	Referential Rating (116.67)
	Efficiency Rating (66.67)
	Column Integrity (70.00)

MICRO-TABLE MODE

```

MICRO Mode:  Rdb Tables/Views

Table:  DEGREES
Cardinality (Record Count):      165
Number of Columns:                5
Number of Bytes (Record Size)    29
Storage Area:  RDB$SYSTEM
Number of Indexes:                2
No Unique Indexes Exist for This Table

```

MICRO-INDEX MODE

```

MICRO Mode:  Rdb Indexes

Table:  DEGREES
Index (002):DEG_COLLEGE_CODE
Index Type :  SORTED, NON-UNIQUE
Avg Dups   :  12
Storage Area:  RDB$SYSTEM
Index columns:  2
1st column:  COLLEGE_CODE

```

MICRO-STORAGE AREA MODE

```

MICRO Mode:  Rdb Storage Areas

Storage Area:  COMP_NAME_INDEX_AREA
Page Size:    4 blocks,  Format: UNIFORM
Number of Extents :  1
RDA Extensions (blocks):  194
SNP Extensions (blocks):  0
-----Stored Elements-----
Tables:  0  Sorted:  1  Hashed:  0

```

Sample DBAnalyzer Report

FULL FORMAT

```

*****
*
* FRENDSDBA$SCRATCH:DBANALYZR.LST
*
* Output generated by Empirical Software's DBAnalyzer V5.2
* 11/08/98 12:57
*
* Rdb Name: PERSONNEL
*
*****

          === Rdb Analysis Summary ===

Complexity Rating      :          6  Buffer Size      :          6
Tune Rating           :          0  Global Buffers : Disabled
Integrity Rating      :      285.71 # of Global Buffers : 250
Number of Users       :          50 Global Buffers/User : 5
Number of Nodes       :          16 After Image Journal : Disabled
Open Mode             : Automatic Journal Fast Commit : Disabled
Number of Buffers     :          20 Adjustable Lock Gran : Enabled
Number of Recovery Buffer:          20 Lock Timeout Interval: 0

Number of Tables      :          10 Number of records : 1,333
Number of Indices     :          7  Avg Recs per Table : 133
Number of Domains     :          28 Number of Triggers : 3
Number of Table Columns :          51 RDA Ext. Block % : 100%
Number of Storage Areas :          0  SNP Ext. Block % : 50%
Number of Views       :          3  Hashed Index % : 0%
Number of Constraints :          20 % Indices to Review : 71%

-----
Tables with the Highest Record Counts      Indices with the Most Duplicates
-----
SALARY_HISTORY          729  DEG_COLLEGE_CODE          12
JOB_HISTORY             274  SH_EMPLOYEE_ID            7
DEGREES                 165  JH_EMPLOYEE_ID            2
EMPLOYEES               100  DEG_EMP_ID                 1
DEPARTMENTS             26   EMP_LAST_NAME              1

-----
Tables with the Most Columns              Tables with the Largest Record Size
-----
EMPLOYEES                12  CANDIDATES                  280
JOB_HISTORY               6   EMPLOYEES                   112
COLLEGES                 5   COLLEGES                    56
DEGREES                  5   DEPARTMENTS                  47
DEPARTMENTS              5   JOB_HISTORY                   34

SORTED Indices Recommend to be HASHED   HASHED Indices Recommend to be SORTED Index /
Table...                               *** None were found ***
DEG_EMP_ID / DEGREES
DEG_COLLEGE_CODE / DEGREES
EMP_EMPLOYEE_ID / EMPLOYEES
JH_EMPLOYEE_ID / JOB_HISTORY
SH_EMPLOYEE_ID / SALARY_HISTORY

(Continued)

```

```

Empirical Software's DBAnalyzer
Rdb Name: PERSONNEL
Page 2

Rdb Analysis (continued)...
=====
Tables with the Most Indices          Non-Indexed Tables with the Most Records
-----
DEGREES                2  DEPARTMENTS                26
EMPLOYEES              2  JOBS                       15
COLLEGES               1  CANDIDATES                 3
JOB_HISTORY            1  RESUMES                    3
SALARY_HISTORY         1  WORK_STATUS                 3

Storage Areas with Most Mapped Items  Storage Areas with Most File Extensions
-----
RDB$SYSTEM              17  RDB$SYSTEM                  4

RDA files with Most Extension Blocks  SNP files with Most Extension Blocks
-----
RDB$SYSTEM              1,736  RDB$SYSTEM                  198

Storage Area Extension Summary
-----
Total RDA Extensions   :           4
Init RDA Alloc (blks)  :           0
Curr RDA Alloc (blks)  :          1,736
Init SNP Alloc (blks)  :           200
Curr SNP Alloc (blks)  :           398

=====

Database Complexity
-----

The complexity rating is a weighted measure of the database
design and its stored records. A rating of 6 indicates a
relatively small database. Tuning requirements are simple.
The largest factor in this rating is the domain count.
It accounts for 17% of the complexity rating. The next largest
component of this rating is the column count, which is also 17%.
The complexity will increase as records and Rdb items (e.g.,
tables, columns) are added.

Database Tune Rating
-----

The tune rating is a composite measure of the physical storage
design as it applies to the logical structure of this database.
The complexity rating of 6 and the tune rating of 0 indicate
that more tuning could be done, but for a database of this size
the physical storage strategy is probably sufficient.
The tune rating measures significant factors that affect the
physical storage design. It does not consider every factor, but
does objectively measure factors critical to successful Rdb tuning.
Remember that overall performance is a factor of many things,
including system load, system tuning, and application design, in
addition to the physical storage strategy.

(Continued)

```

```
Empirical Software's DBAnalyzer                               Page 3
Rdb Name: PERSONNEL

Rdb Analysis (continued)...
=====

    DBTune, a complementary Empirical Software product, can be used to improve
the tune rating for a database. DBTune uses volume, environment,
and activity data to tune storage areas, indices, and database
parameters. SQL scripts to implement these tuning changes are
automatically generated as well as reports that give advice
on tuning options and SYSUAF/SYSGEN parameter settings.
@COURIER6IND12 =
Database Storage Area Allocation
-----

    The 'Storage Area Allocation' graphs indicate the percentage of
allocated space that has been extended for both RDA and SNP files.
Of the total RDA pages, a large percentage (100%) have been
extended as these storage areas have been loaded. RDA pages are
extended when the data requirements for tables and/or indices exceed
the existing allocation of pages. The RDA areas have been extended
4 times.

    Of the total SNP pages, a large percentage (50%) have been
created as these storage areas have been utilized. SNP files are
used to enable READ-only transactions to access data concurrently
while WRITE transactions are active. SNP pages are only used if
SNAPSHOTS ARE ENABLED.
@COURIER6IND12 =
Database Index Analysis
-----

    The index analysis graphs show that 0% of the database indices are HASHED.
Thus, 100% of the indices are SORTED. DBAnalyzer reviewed the indices
and found that five of the SORTED indices are candidates to be HASHED
and none of the HASHED indices are candidates to be SORTED. Review
MACRO windows 5 and 6 to see which indices have been selected.

    DBAnalyzer looks for certain key words within the index columns. It
assumes certain types of queries will be made based on these key words.
The person responsible for maintaining the database should review
actual usage to determine whether to modify the index.
    NOTE:  HASHED indices facilitate exact match queries.
           They incur narrow locks for updates.
           SORTED indices facilitate sequential retrievals.
           They incur broader lock contention for updates
           than HASHED indices.
@COURIER6IND12 =
MACRO and MICRO windows
-----

    DBAnalyzer is a tool to increase the productivity of the
individual who manages Rdb databases. The overview provides
database-wide objective measurements. The MACRO windows
highlight significant items that affect the database, its
applications, and its users. The MICRO windows provide
additional details to assist in making decisions to maintain
and improve the usefulness of an Rdb database.

(Continued)
```

```
Empirical Software's DBAnalyzer
Page 4
Rdb Name: PERSONNEL

Rdb Analysis (continued)...

=====
==

Database Integrity Ratings
-----

The Database Integrity Rating is an overall measure of the utilization
of database constraints. A rating of 285.71 indicates that few additional
constraints could be implemented to improve database integrity. As
more tables and indices are added to the database, this rating may
drop unless a corresponding number of constraints are added as well.
Following are ratings that focus on more specific areas of database
integrity.

The Columnar Integrity Rating considers column-level constraints such
as CHECK, NOT NULL, and FOREIGN KEY constraints that enforce data validation
in and between tables. A rating of 82.35 indicates that few additional
constraints could be implemented to improve the columnar integrity of
the database. An example of such a column constraint is:
CHECK(STATUS IN ('ACTIVE','INACTIVE') OR STATUS IS NULL).

The Referential Integrity Rating considers UNIQUE, PRIMARY KEY,
and FOREIGN KEY constraints that enforce uniqueness or referential
validation in and between tables. A rating of 116.67 indicates
that few additional foreign key constraints could be implemented to
improve the referential integrity of the database.

The Referential Efficiency Rating measures the extent to which existing
indices mirror UNIQUE and PRIMARY KEY constraints. If such a constraint
constraint more efficiently by taking advantage of the underlying
index. A rating of 33.33 indicates that additional unique
indices could be implemented to improve the referential efficiency
of the database.

(Continued)
```



```

Empirical Software's DBAnalyzer                               Page 5
Rdb Name: PERSONNEL

Storage Area Information:
=====
Area: RDB$SYSTEM ..... (PgSz: 2, Fmt:UNI, Tbls:10, Sort:7, Hash:0,
                          Exts/blks--RDA:4/1736, SNP:1/198)

      Stored  Elmnt Cardin- Column  Byte  Index  Average
      Elements Type  ality  Count  Count Type  Unique  Dups  Comments
      -----
CANDIDATES  TABLE      3      4    280
COLLEGES    TABLE     15      5     56
COLL_COLLEGE_CODE INDEX      1      4 SORTED YES      N/A  TBL: COLLEGES
DEGREES     TABLE    165      5     29

.
.
.
@COURIER6IND12 =

Empirical Software's DBAnalyzer                               Page 6
Rdb Name: PERSONNEL

Individual Table Statistics...
=====
Table : JOB_HISTORY ... (Card:274, Col:6, Byt:34, Idx:1, ..... Areas:RDB$SYSTEM)
>> Index: JH_EMPLOYEE_ID (SORTED, Dup:2, Col:1, Byt:5, Node:default, Areas:RDB$SYSTEM)

      Columns for Table:
      -----
      JOB_HISTORY      Column Domain      Domain/Datatype Description
001 EMPLOYEE_ID      ID_DOM      CHAR(5)
002 JOB_CODE      JOB_CODE_DOM      CHAR(4)
003 JOB_START      DATE_DOM      DATE VMS
004 JOB_END      DATE_DOM      DATE VMS
005 DEPARTMENT_CODE DEPARTMENT_CODE_DOM      CHAR(4)
006 SUPERVISOR_ID      ID_DOM      CHAR(5)

      Columns for Index:
      -----
      JH_EMPLOYEE_ID      Column Domain      Domain/Datatype Description
001 EMPLOYEE_ID      ID_DOM      CHAR(5)

.
.
.

(Continued)
    
```

```

Empirical Software's DBAnalyzer                               Page 7
Rdb Name: PERSONNEL

Individual View Statistics ...
=====
View : CURRENT_INFO
>> information related to work status codes

      Column Name                Domain Name          Data Type Description
-----
LAST_NAME                       LAST_NAME_DOM       CHAR(14)
FIRST_NAME                      FIRST_NAME_DOM      CHAR(10)
ID                               ID_DOM              CHAR(5)
DEPARTMENT                      DEPARTMENT_NAME_DOM CHAR(30)
JOB                             JOB_TITLE_DOM       CHAR(20)
JSTART                          DATE_DOM            DATE VMS
SSTART                          DATE_DOM            DATE VMS
SALARY                          SALARY_DOM          INTEGER(2)

SELECT
  CJ.LAST_NAME,
  CJ.FIRST_NAME,
  CJ.EMPLOYEE_ID,
  D.DEPARTMENT_NAME,
  J.JOB_TITLE,
  CJ.JOB_START,
  CS.SALARY_START,
  CS.SALARY_AMOUNT
FROM CURRENT_JOB CJ,
DEPARTMENTS D,
JOBS J,
CURRENT_SALARY CS
WHERE CJ.DEPARTMENT_CODE = D.DEPARTMENT_CODE
AND CJ.JOB_CODE = J.JOB_CODE
AND CJ.EMPLOYEE_ID = CS.EMPLOYEE_ID

                                     (Continued)
.
.
.

```

```

Empirical Software's DBAnalyzer                               Page 8
Rdb Name: PERSONNEL

Domain Summary ...
=====
  Domain Name      Column Name having Domain  Table Name  DataType Description
-----
ADDRESS_DATA_1_DOM  ADDRESS_DATA_1             EMPLOYEES   CHAR(25)
>> standard definition for street addresses
ADDRESS_DATA_2_DOM  ADDRESS_DATA_2             EMPLOYEES   CHAR(20)
>> standard definition for apartments, suites,
BUDGET_DOM          BUDGET_ACTUAL              DEPARTMENTS INTEGER(0)
>> standard definition for departmental budget
                   BUDGET_PROJECTED         DEPARTMENTS

* Report generation parameters:
* dba.report_name = DBANALYZR.LST
* dba.select.full = YES
* dba.select.brief = NO
* dba.select.domains = YES
* dba.select.narrative = YES
* dba.select.storage_areas = YES
* dba.select.tables = YES
* dba.select.views = YES
*
* dba.print_report = NO
* dba.printer =
* dba.printer_options =
*
* dba.domain.sort.order = DOMAIN
*
* dba.storage_source = ALL
*
* dba.table_source = ALL
* dba.table_option =
* dba.table_columns = YES
* dba.table_indices = YES
* dba.table_indices_option = FULL
* dba.table_storage_areas =
*
* dba.view_source = ALL

```

DBAnalyzer Help

SET_REPORT_NAME

Changing the report output name from the default DBANALYZR.LST offers you an opportunity to give the report a more meaningful name that is related to the contents. If multiple users are running DBAnalyzer and are generating reports on or about the same time, it is advisable to change the report name.

REPORT_PARAMETERS

Report Format Selections
FULL SUMMARY COMPONENTS

There are two basic formats available for the report(s) generated by DBAnalyzer. The **FULL** option will generate a comprehensive report about all aspects of the database.

The **SUMMARY** report only reports the information available in the MACRO windows of DBAnalyzer, plus a narrative explaining the complexity and tune rating of the database.

To tailor the report, you should select **COMPONENTS**. This will allow you to report on selected components of the database at various levels of detail. In this manner, you can report on selected storage areas or selected tables, as opposed to the entire database at one time.

REPORT_COMPONENTS

Component Selections
DOMAINS STORAGE AREAS TABLES/INDEXES VIEWS

You can select from four options when formatting a report by components. Multiple selections are available from this menu. To report on a component, position the bar over an item using the arrow keys, and press the **Select** key. When the selection process is complete, press **Return**.

DOMAIN_SORT_ORDER

Domain Sort Order Criteria
COLUMN DOMAIN TABLE

The **Domain Sort Order** determines the sorted order in which the domains and fields in the database are presented. In **DOMAIN** order, the fields will be sorted by the **DOMAIN** name. Thus, all fields sharing a common domain will be printed together. In **TABLE** order, all the fields in a given table will be printed together, yielding a table layout. In **FIELD** order, tables and domains are ordered by field name. This can be useful when looking for incorrect definitions of a field in different tables.

STORAGE_AREA_SELECTION

Storage Areas Selection Criteria
ALL SELECTED

You can specify either **ALL** or **SELECTED** storage areas. If **ALL** is selected, then all storage areas in the database will be reported. If you opt for **SELECTED**, a window will be presented which will allow you to specify the area(s) to be reported.

SELECTED_STORAGE_AREAS

Storage Areas Selection Criteria
ALL SELECTED

Storage Area Selections
RDB\$SYSTEM . . .

This window presents you with a selection of all storage areas in the database. You can select one or more storage areas by arrowing up and down through the window and highlighting the areas to be included by pressing the **Select** key. Pressing **Return** will include the highlighted areas in the final report.

VIEW_SELECTIONS

View Selection Criteria
ALL SELECTED

You can specify either **ALL** or **SELECTED** views. If **ALL** is selected, then all views in the database will be reported. If you opt for **SELECTED**, a window will be presented which will allow you to specify the view(s) to be reported.

SELECTED_VIEWS

Storage Areas Selection Criteria	
ALL SELECTED	
	View Selections
	CURRENT_INFO CURRENT_JOB CURRENT_SALARY

This window presents you with a selection of all views in the database. You can select one or more views by arrowing up and down through the window and highlighting the ones to be included by pressing the **Select** key. Pressing **Return** will include the highlighted views in the final report.

TABLE_DETAIL_OPTIONS

Table Detail Criteria
COLUMNS NO COLUMNS

You can select from several levels of detail concerning the tables in the database. The **COLUMNS** option will give a listing of the columns in the table, columns in each index in the table, and such summary information as record size, column count, and cardinality.

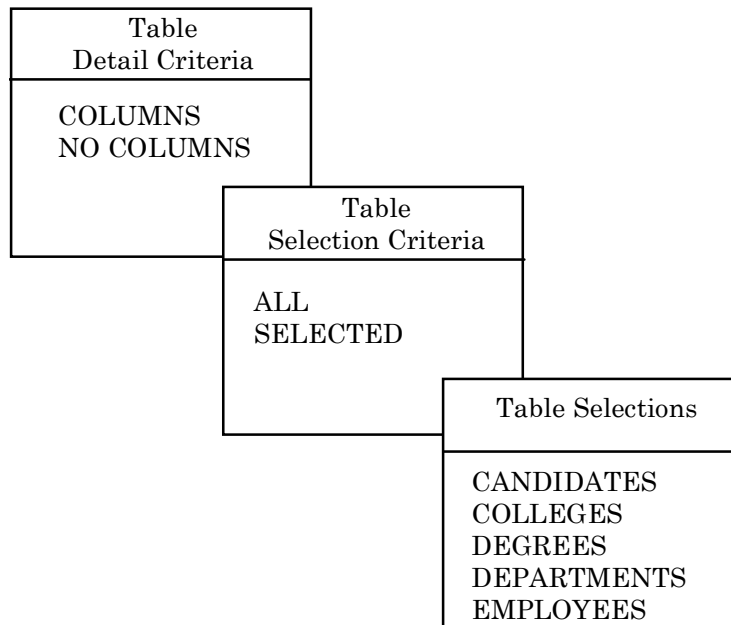
The **NO COLUMNS** option presents only summary information for tables/indexes and omits the column listings.

TABLE_SELECTIONS

Table Detail Criteria	
COLUMNS NO COLUMNS	
	Table Selection Criteria
	ALL SELECTED

You can specify either **ALL** or **SELECTED** tables. If **ALL** is selected, then all tables in the database will be reported. If you opt for **SELECTED**, a window will be presented which will allow you to specify the table(s) to be reported.

SELECTED_TABLES



This window presents you with a selection of all tables in the database. You can select one or more tables by arrowing up and down through the window and highlighting the ones to be included by pressing the **Select** key. Pressing **Return** will include the highlighted tables in the final report.

EXECUTION_OPTIONS

Execute Options
ONLINE BATCH

After specifying the format of the report, you may elect to run the report in **BATCH** or **ONLINE**. Running the report in **BATCH** frees up your screen to return to scanning the database or to specify a new report. For large databases and reports that include **DOMAINS** (including the **FULL** report), it is recommended that the report be run in batch.

PRINT_REPORT

Execute Options
ON-LINE BATCH
Print Report After Completion?
YES NO

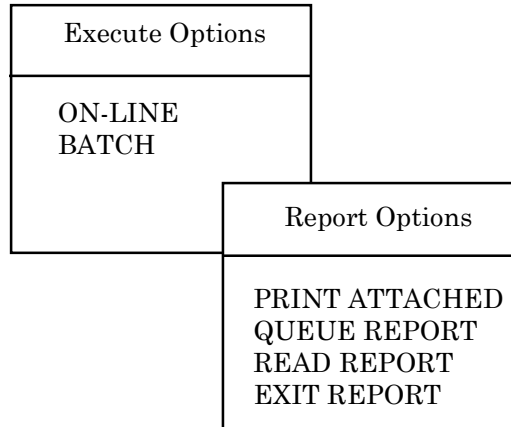
If you elect to run the report in **BATCH**, you may elect to print the report on completion. If you elect to print the report, you will be given a list of available printer queues to specify for the output.

AVAILABLE_QUEUES

Execute Options		
ON-LINE BATCH		
	Print Report After Completion?	
	YES NO	
		Available Printer Queues
		SYS\$PRINT LA120 LASER01 LASER02 DEPT_PRINTER

This is a list of all available printer queues on your system. Be aware that the report(s) generated by DBAnalyzer require the printer be set to 132 column mode. Select your printer accordingly.

REPORT_OUTPUT_OPTIONS



PRINT ATTACHED will print the report to your attached printer port device.

QUEUE REPORT will allow you to print the report to one of the available printer queues on the system.

READ REPORT will put you into the **OpenVMS EDT** editor to peruse the output. To read the report in 132 column mode, after entering the file, press <<PF1>>W. An **EDTINI** file is included with DBAnalyzer that defines the following keys:

<<PF1>>W	Set screen 132
<<PF1>>N	Set screen 80
<<PF1>>R	Shift screen right
<<PF1>>L	Shift screen left
<<PF1>>Q	Quit out of editor

EXIT REPORT returns you to scan mode.

Chapter 2

DBTune for Rdb

What is DBTune for Rdb?

DBTune for Rdb is intended to maximize your Rdb performance without requiring excessive effort. To gather information necessary for the tuning process, it considers both:

1. The existing Rdb design (both physical and logical), and
2. Performance Analysis Data.

Thus, both the database design and its actual usage are incorporated in the database optimization procedure, typically resulting in a 30 to 50 percent improvement in performance (direct I/O, CPU usage, and elapsed time).

Using the logical and physical data gathered from the database, DBTune creates a Performance Analysis Data (PAD) file. This PAD file contains volume, workload, and environment information. You may customize this information with additional transaction activity. Additionally, the PAD file can be supplemented with a dynamic analysis of observed transaction data, which you can create manually or automatically using a third-party tool. The tuning process then combines the Performance Analysis data, the logical and physical Rdb design, and the transaction analysis results to create an optimized design.

DBTune achieves this transformation of the Rdb database without affecting the programs and queries that access it. Rather, the Rdb database is converted to a more efficient physical structure. The transformation takes advantage of the high-performance capabilities of Rdb, including the use of multi-file structures. Each storage area is tuned so that the appropriate page size and allocation are assigned according to the analysis inputs provided. Both sorted and hashed indexes are stored so as to maximize throughput and minimize lock contention. Additionally, database parameters are tuned and storage areas are distributed across available disk drives to take advantage of system capacity for the particular physical design.

To assist with Rdb tuning plans, DBTune provides an objective measure of a database's complexity and tune rating. Empirical data indicate that databases with a complexity rating of at least 15 to 25 can achieve significant performance improvement through increase in complexity and the increase in database usage (e.g., more users, reports, etc.).

Rdb Version 7.0

Note DBTune 5.2 works with Rdb 7.1 .DBTune 5.2 works with Rdb 7.0 . You must use DBTune 4.0 if you are running an earlier version of Rdb.

Temporary Tables

DBTune 5.2/5.3 correctly tunes both global and local temporary tables. Temporary tables, a new feature in Rdb 7, provide the convenience of storing and manipulating short-term data in a table, rather than using a flat file or repeatedly creating and then dropping a table. Temporary tables can be used to store the output of a query or other intermediate results. Global temporary tables allow data to be shared between different modules in a single SQL session. Local temporary tables do not allow data to be shared. Metadata for both global and local temporary tables is stored in the database and persists beyond a SQL session. Data in temporary tables does not persist beyond the SQL session.

Since the data in the temporary table does not exist after the SQL session, temporary tables will always be empty after tuning. DBTune 5.2/5.3 correctly tunes and preserves any global or local temporary tables present in an Rdb 7 database.

Row Cache

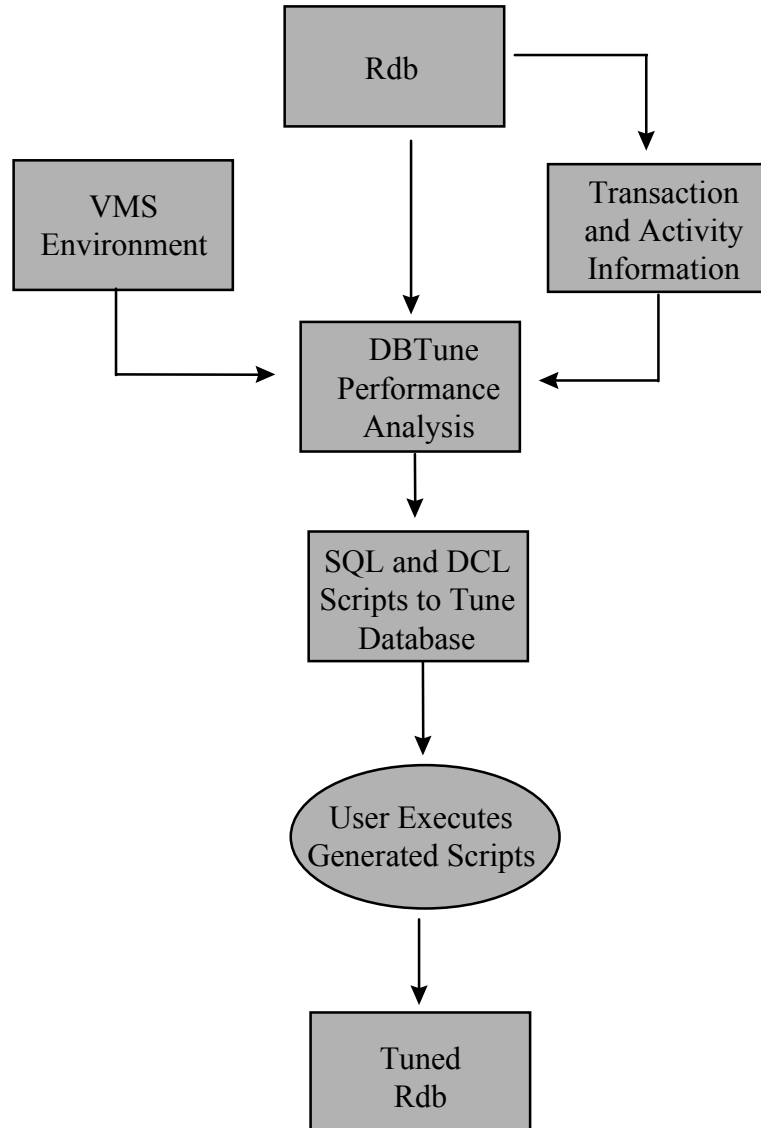
Row cache allows the most frequently accessed rows of a logical or physical database area to remain in memory even after the associated pages are flushed back to the disk. Row cache can be used on single node VAX and AXP systems and can significantly improve performance if memory is ample and most database accesses involve a limited number of table or index rows.

At this time, Rdb does not support read/write row caching for production databases.

Sorted Ranked Indexes

Rdb 7 supports a ranked B-tree structure for sorted indexes. Sorted ranked indexes allow better optimization of queries, especially queries involving range retrievals. The sorted ranked index type can also reduce lock contention and disk I/O. DBTune 5.2/5.3 preserves any sorted ranked indexes while running an Rdb 7.0 database.

**DBTune Considers
Dynamic Workload, Volume, and Environment Inputs
During Its Performance Analysis**



Getting Started

This section provides the information needed for you to quickly install and run DBTune. It also includes system requirements necessary to run the application.

DEC VAX/OpenVMS

Recommended **minimum AUTHORIZE** settings for a DBTune user account. (Medium and large databases may need to increase these numbers.)

Username:	USER	Owner:	USER
Account:	USER	UIC:	[Group, Member]
CLI:	DCL	Tables:	DCLTABLES
Default:	<disk>.[dir]		
LGIMD:	LOGIN		
Login Flags:			
Primary Days:	Mon. Tues. Wed. Thurs. Fri.		
Secondary Days:	Sat. Sun.		
No access restrictions			
Expiration:	(none)	Pwdminimum:	0
Pwdlifetime:	(none)	Pwdchange:	
Last Login:		Login Fails:	0
Maxjobs:	0	Fillm:	512
Maxacctjobs:	0	Shrfillm:	0
Maxdetach:	0	BIOfm:	100
Prclm:	4	DIOfm:	100
Prio:	4	ASTfm:	113
Queprio:	0	TQElm:	10
CPU:	(none)	Enqlm:	8000
Authorized Privileges:	GROUP TMPMBX NETMBX	BytIm:	120000
Default Privileges:	GROUP MPMBX NETMBX	PbytIm:	0
		Jtquota:	1024
		Wsdef:	1024
		Wsquo:	5120
		Wsextent:	5120
		Pqlfquo:	160000

Warning If these minimums are not in place when DBTune is executed, the tuning process may fail!

Note See the **REVIEW_AND_GUIDE.REPORT** file created by DBTune for suggested **AUTHORIZE** settings for individual database user accounts.

Below are **MINIMUM** settings recommended for several **SYSGEN** parameters. If changes are made to any of the **SYSGEN** parameters listed below, your system will need to be **REBOOTED** to make the changes effective. It is recommended that **AUTOGEN** be used to make any required changes.

SYSTEM PARAM	MINIMUM SETTING	DESIRED SETTING
CHANNELCNT *	512	2047
VIRTUALPAGECNT *	160000	240000
LOCKIDTBL **	2048	10240
LOCKIDTBL_MAX **	2048	61440
RESHASHTBL **	512	2560
PROCSECTCNT	64	64
GBLPAGES *	100000	200000
GBLSECTIONS	600	600
GBLPAGFIL ***	50000	100000
CTLPAGES ****	100	200

(*) The values for these parameters may need to be increased if using Rdb **Global Buffers**. Please see the **REVIEW_AND_GUIDE.REPORT** generated by DBTune for more details on a particular database.

(**) The setting for **LOCKIDTBL** must be four (4) times the setting for **RESHASHTBL**. If you change one setting, you should change the other as well. The setting for **LOCKIDTBL_MAX** must be equal to or greater than the setting for **LOCKIDTBL**.

(***) Total system **PAGE FILE** space must be larger than the setting for **GBLPAGFIL**. Thus, if **GBLPAGFIL** is increased, ensure that adequate **PAGEFILE.SYS** space exists. You can view the current **PAGE FILE** sizes on a system with the DCL command: `$ SHOW MEMORY/FILES`.

(****) The setting for **CTLPAGES** is recommended but not required. However, if logicals are being used to specify database storage areas, the suggested setting for **CTLPAGES** should be adhered to.

DEC AXP/OpenVMS

Recommended **minimum AUTHORIZE** settings for a DBTune user account.
(Medium and large databases may need to increase these numbers.)

Username:	USER	Owner:	USER
Account:	USER	UIC:	[Group, Member]
CLI:	DCL	Tables:	DCLTABLES
Default:	<disk>:[dir]		
LGIMD:	LOGIN		
Login Flags:			
Primary Days:	Mon. Tues. Wed. Thurs. Fri.		
Secondary Days:	Sat. Sun.		
No access restrictions			
Expiration:	(none)	Pwdminimum:	0
Pwdlifetime:	(none)	Pwdchange:	
Last Login:		Login Fails:	0
Maxjobs:	0	Fillm:	512
Maxacctjobs:	0	Shrfillm:	0
Maxdetach:	0	BIOlm:	150
Prclm:	4	DIOLm:	150
Prio:	4	ASTlm:	250
Queprio:	0	TQEIm:	10
CPU:	(none)	Enqlm:	8000
Authorized	GROUP TMPMBX	Bytlm:	120000
Privileges:	NETMBX	Pbytlm:	0
Default	GROUP MPMBX	Jtquota:	1024
Privileges:	NETMBX	Wsdef:	2000
		Wsquo:	4096
		Wsexent:	16384
		Pqlfquo:	160000

Warning If these minimums are not in place when DBTune is executed, the tuning process may fail!

Note See the **REVIEW_AND_GUIDE.REPORT** file created by DBTune for suggested **AUTHORIZE** settings for individual database user accounts.

Below are **MINIMUM** settings recommended for several **SYSGEN** parameters. If changes are made to any of the **SYSGEN** parameters listed below, your system will need to be **REBOOTED** to make the changes effective. It is recommended that **AUTOGEN** be used to make any required changes.

SYSTEM PARAM	MINIMUM SETTING	DESIRED SETTING
CHANNELCNT *	512	2047
VIRTUALPAGECNT *	160000	240000
LOCKIDTBL **	2048	10240
LOCKIDTBL_MAX **	2048	61440
RESHASHTBL **	512	2560
PROCSECTCNT	64	64
GBLPAGES *	100000	200000
GBLSECTIONS	600	600
GBLPAGFIL ***	50000	100000
CTLPAGES ****	100	200

(*) The values for these parameters may need to be increased if using Rdb Global Buffers. Please see the **REVIEW_AND_GUIDE.REPORT** generated by DBTune for more details on a particular database.

(**) The setting for **LOCKIDTBL** must be four (4) times the setting for **RESHASHTBL**. If you change one setting, you should change the other as well. The setting for **LOCKIDTBL_MAX** must be equal to or greater than the setting for **LOCKIDTBL**.

(***) Total system **PAGE FILE** space must be larger than the setting for **GBLPAGFIL**. Thus, if **GBLPAGFIL** is increased, ensure that adequate **PAGEFILE.SYS** space exists. You can view the current **PAGE FILE** sizes on a system with the DCL command: `$ SHOW MEMORY/FILES`.

(****) The setting for **CTLPAGES** is recommended but not required. However, if logicals are being used to specify database storage areas, the suggested setting for **CTLPAGES** should be adhered to.

Installing DBTune

► To install DBTune for Rdb from a tape drive:

1. Back up your system disk (optional).
2. Log in under the SYSTEM account.
3. Put the DBTune distribution tape in the tape drive.
4. Type in the following DCL command to invoke the VMS install facility to install DBTune on your system:

For VAX/VMS

```
$ @SYS$UPDATE:VMSINSTAL DBTRDBVMS52 <<tape-drive>>:
```

For Alpha AXP

```
$ @SYS$UPDATE:VMSINSTAL DBTRDBAXP52 <<tape-drive>>:
```

where <<tape-drive>> is the name of the device where the DBTune distribution tape has been mounted (e.g., MUA6:).

Note DBTune V5.2/5.3 should NOT be installed in the same directory with other versions of DBTune or any other product from ALI (i.e., DBAnalyzer).

5. After the VMS install has completed, place the following lines into the system startup command file (SYS\$MANAGER:SYSTARTUP_VMS.COM) so that required logicals are set up when the system is rebooted:

```
$ DEFINE/SYSTEM/EXEC FRENDDBTUNE$HOME <<disk>>: [dir]
$ DEFINE/SYSTEM/EXEC FRENDDBTUNE$SCRATCH
  <<disk>>: [dir.SCRATCH]
```

where <<disk>> and [dir] are the disk and directory to which DBTune was installed (e.g., \$1\$DUA1:[DBTRDBVMS52] or \$1\$DUA1:[DBTRDBAXP52]).

- Now, to obtain a license pak for DBTune, type in the following commands:

```
$ SET DEFAULT FRENDDBTUNE$HOME
$ EDIT DBTUNE.LICENSE
```

For each node (“machine”) on which you wish to run DBAnalyzer:

- Replace “your node name” with the node name of the machine on which you have installed DBAnalyzer. To get this information, type:

```
$ WRITE SYS$OUTPUT F$GETSYI (“nodename”)
```

- If the “operating system” value supplied with your license is not accurate for your system, replace it with the output generated from the following command:

```
$ WRITE SYS$OUTPUT F$GETSYI (“node_swtype”)
```

- Replace “your company name” with your company’s full name.
- Exit and save the file.

To obtain the appropriate registration ID for each machine entered, call ALI at (866) 257-8970 [or (803) 648-5931], or fax a copy of the altered DBTUNE.LICENSE file to (803) 641-0345. International clients may also obtain registration IDs or support through their local distributor’s office.

► To install DBTune for Rdb from a CD-ROM:

- Mount the CD using a command like

```
$ MOUNT/OVER=ID <cd_device>:
```

2. Install the product with the command

```
$ @SYS$UPDATE:VMSINSTAL <product_name>  
<cd_device>: [INSTALL]
```

where <product_name> is the product you wish to install. For example:

```
$ @sys$update:vmsinstal DBTRDBVMS052 dka400: [INSTALL]
```


Using DBTune

DBTune requires that you have either the VMS privilege **SYSPRV** or the appropriate database and **RMU** privileges. If you do not have **SYSPRV** privilege, DBTune requires the following **RMU** privileges:

```
"RMU$BACKUP," "RMU$UNLOAD," "RMU$LOAD," "RMU$DUMP,"  
"RMU$OPEN," and "RMU$ANALYZE"
```

Note DBTune 5.3 requires that the database to be tuned be Rdb 7.1 .DBTune 5.2 requires that the database to be tuned be Rdb 7.0 .

Online Use of DBTune

► To run DBTune online:

```
$ @FREND$DBTUNE$HOME:DBTUNE.COM
```

Note Setting up a VMS symbol can make this easier.

```
$ DBTUNE ::= "@FREND$DBTUNE$HOME:DBTUNE.COM"  
allowing you to execute DBTune by typing:  
$ DBTUNE
```

Batch Use of DBTune

► To run DBTune in BATCH:

```
$ @FRIEND$DBTUNE$HOME:DBTUNE.COM
```

If executed in **BATCH**, DBTune expects the **FRIEND\$RDB\$IMPORT** logical to have been assigned prior to execution— you will not be prompted. To this end, a command file has been provided to allow assignment of the **FRIEND\$RDB\$IMPORT** logical. This command file— **DBTUNE_BATCH.COM**—can be found in the directory **FRIEND\$DBTUNE\$HOME**. After editing **DBTUNE_BATCH.COM**, DBTune can be invoked in batch with the command:

```
@$ @FRIEND$DBTUNE$HOME:DBTUNE_BATCH.COM
```

Following is a copy of the unedited batch command procedure:

```

$!-----
$!   DBTUNE_BATCH.COM
$!   - Command file to submit DBTune in BATCH mode...
$!
$!   Invoke this file with the command:
$!   $ @FRIEND$DBTUNE$HOME:DBTUNE_BATCH
$!-----
$!
$!   This command procedure will submit itself to batch.
$!
$!   if pl .eqs. "" .or. pl .nes. "BATCH"
$!   then
$!
$!   Change the /name="" qualifier to specify a different name for
$!   the job.
$!
$!       cur_def = f$environment("DEFAULT")
$!       vfl = f$verify(0)
$!       set verify
$!       submit-
$!           /log='cur_def'-
$!           /noprint-
$!           /name="DBTune in Batch"-
$!           /parameters=("BATCH","'cur_def'") -
$!           FRIEND$DBTUNE$HOME:DBTUNE_BATCH.COM
$!       vfl = f$verify(vfl)
$!       exit
$!   endif
$!-----
$!   Change the following ASSIGN statement to point the database
$!   you wish to analyze and uncomment it by removing the "!" ...
$!
$!   ASSIGN "disk1:[directory]database_name" FRIEND$RDB$IMPORT
$!
$!   Change the following SET PROC/NAME= to assign a different
$!   process name and uncomment it by removing the "!" ...
$!
$!   SET PROC/NAME="DBTune in Batch"
$!
$!   Change the following SET DEFAULT to change the default
$!   directory where the report(s) will be generated. Otherwise,
$!   output will be generated the user's current directory at the
$!   time the file was submitted.
$!
$!   SET DEFAULT 'p2'
$!-----
$! @FRIEND$DBTUNE$HOME:DBTUNE.COM

```

DBTune Parameters

Prior to executing DBTune, it is recommended that you review the default parameter settings provided with the tool. DBTune parameters allow you to control the Rdb transformation procedure that is generated. All of the parameters are pre-set to handle typical scenarios. You are encouraged to tailor these parameters to match your particular environment. This customization can be accomplished by editing the parameter file prior to invoking DBTune. DBTune then parses the parameter file and ensures that valid values have been selected. If a parameter value is found to be invalid, its default value is used.

The parameter file that DBTune will parse is pointed to by the logical `FREND$DBTUNE$PARAMS`. If this logical is not assigned prior to execution, DBTune will use the default parameter file that is provided (`FREND$DBTUNE$HOME:DBTUNE_DEFAULT.PARAMS`). To customize the parameter settings, you can do one of two things:

1. either edit the default parameter file directly

or
2. copy the default parameter file to a new file with a different name and assign the logical `FREND$DBTUNE$PARAMS` to this new file.

If you are using DBTune for multiple databases, it is recommended to create a parameter file for each database. The parameter file contains environment information, such as available disk drives, along with the space that can be used for each database. Thus, using separate parameter files for each database will facilitate repeated use of DBTune with minimal setup time.

Note It may be beneficial to carry this idea a step further and set up a subdirectory for each database that will hold all of the DBTune input and output files. This arrangement provides a great deal of easy-to-access documentation for each database.

Descriptions of each of DBTune's parameters and their possible values can be found later in this manual on page 80 under the section titled **DBTune Process - Step 2: Load Parameters**.

DBTune Keystrokes

The following keystrokes may be used when executing DBTune online using a VT220 (or higher) terminal interface:

[Help], [PF2]	Receive help for the current DBTune context or option.
[Esc], [PF3], [PF4]	Exit DBTune; if you are in a selection window to change a DBTune parameter, then function keys will simply return you to the main menu.
[Do]	Continue with the DBTune tuning process.
[Select]	Select a DBTune parameter to change for the current session.
[Ctrl] [W]	Refresh screen display.

After the database has been initially scanned and the screen displayed, a menu is presented on the bottom two lines of the display:

- **[DO]**-Create Tuning Scripts
- **[SELECT]**-Edit Parameters
- **[HELP]**-HELP
- **[ESC]**-Exit

If you are on a VT200+ terminal, these options can be executed with the designated function key. If you are on a PC or VT200- terminal, you can use the arrow keys to move around between the options. To select a function by means of the arrow keys, highlight the designated option and press **Return**.

If you wish to change a parameter setting, choosing the **Select** menu option will present a list of DBTune parameters. You can arrow between selections and press **Return** on the parameter to be changed. A window will then be presented in which to enter the new value. Online help is available at either the menu or in the change window. After completing the editing process, pressing **ESC** or **PF3** will return you to the main menu.

DBTune Process

The DBTune transformation process is automated to provide a simple method to achieve excellent Rdb performance. The process includes nine steps. The first eight steps occur automatically via the DBTune program and produce a command procedure that you manually execute during the ninth step to perform the Rdb transformation. The process can be executed online or in batch. The final transformation step can be invoked immediately upon its generation. It is recommended, however, to scan the **REVIEW_AND_GUIDE.REPORT** file created by DBTune and the DBTune log files that are generated during the first eight steps.

These documents will report any errors that may occur during the execution of DBTune as well as provide additional information to ensure successful execution of the transformation procedure. The nine steps of the DBTune process are:

1. Analyze Rdb
2. Load Parameters
3. Read Database Structure
4. Read Customized Analysis/Workload Data
5. Generate Performance Analysis Data File
6. Performance Analysis and Database Tuning
7. Generate Disk Utilization File
8. Generate Rdb Transformation Procedure

9. Transform Rdb

Note DBTune performs the first eight steps; you perform the ninth step.

The first eight steps access the database in READ-ONLY mode and can be executed while other users are accessing the database. The ninth step, which you execute manually, actually tunes the database and requires all users to be out of the database during its execution. Each of these steps is described in more detail in the following sections.

Step 1: Analyze Rdb

The database is scanned to collect Rdb statistics on the logical and physical design. This information is summarized into Rdb component counts, a **COMPLEXITY** rating, and a **TUNE** rating. Additionally, a narrative is produced to explain the statistics and ratings in a format that emphasizes their performance impact.

```

*** DBTune V5.2 ***
Rdb Name: 1$DUA11:[YOUNGYG.TESTDB]MF_PERSONNEL.RDB
Counting Database Items
- Tables
- Views
- Indices
- Domains
- Columns
- Constraints
- Triggers
- Storage Areas
Scan Progress
XXXXXXXXXXXX
Tables      :
Indices    :
Domains    :
Columns    :
Storage Area:
Views      :
Constraints :
Triggers   :
Records    :
Avg Rec/Tbl :
Complexity Rating
Tune Rating

```

This information is gathered during the initial phase of DBTune. The narrative report is written to the current default directory. The DBTune process can be stopped at this point so that the analysis can be reviewed before proceeding with the transformation process. The narrative analysis report is called `<<database_id>>_ANALYSIS.REPORT`, where `<<database_id>>` represents the first ten characters of the .RDB file name. For example, the report for the PERSONNEL database would be called:

PERSONNEL_ANALYSIS.REPORT

Step 2: Load Parameters

DBTune parameters allow you to control the Rdb transformation procedure that is generated. All of the following parameters are pre-set to handle typical scenarios. You are encouraged to tailor these parameters to match their particular environment. To customize parameter settings, you can do one of two things:

1. Edit the default or customized parameter file prior to executing DBTune or
2. Change the parameter online during DBTune execution by choosing the **[SELECT]-Edit Parameters** menu option. Not all parameters are available for editing online and any changes made will affect the current session only—the changes are NOT saved to the parameter file.

DBTune reads the parameter file and parses for exact matches on the parameter key word followed by an “=.” Everything to the immediate right of the “=” is considered the parameter value.

Note If DBTune detects errors when parsing the parameter file, you will be given the opportunity to view a log of the errors found and to make corrections. If corrections are made, the parameter file is actually updated with the new value. You may also elect to continue without making corrections, allowing DBTune to substitute default values for erroneous parameters.

Following are descriptions of each of the (37) DBTune parameters, their effects and their possible values:

(1) **STRATEGY**

The **STRATEGY** parameter controls DBTune's handling of the database storage areas. There are three settings possible:

- E** Use **EXISTING** storage areas and leave them in their current physical locations
- N** Create all **NEW** storage areas, automatically distributing them to new locations across specified disks
- R** A combination of "E" and "N"; use existing storage areas but **RELOCATE** them by automatically distributing to new locations across specified disks

Regardless of the **STRATEGY** setting, all tables and indexes that are to be tuned and are currently stored in RDB\$SYSTEM (the root area) will be moved into a new storage area of their own and will be distributed automatically across specified disks. If tuning via **SQL IMPORT**, segmented strings (lists) will also be moved into new storage areas and will be distributed automatically across disks.

VALUES: **E** Use **EXISTING** storage areas
 N Create all **NEW** storage areas
 R Use existing, but **RELOCATE**

DEFAULT: **N**

(2) *TUNE_TECHNIQUE*

The **TUNE_TECHNIQUE** parameter determines the method used to tune a database. An entire database can be tuned all at once by using a **SQL EXPORT/IMPORT**, or parts of a database can be tuned by using **RMU UNLOAD/LOADs**. If **RMU UNLOAD/LOADs** are chosen to tune a database, you can specify particular tables and/or indexes to be tuned while leaving other tables/indexes alone. In addition, you can specify a time limit (in minutes) for the **UNLOAD/LOAD** process in which only those tables that provide the most performance benefit will be tuned. This time limit can be specified via the DBTune parameter **LOAD_TIME_LIM**.

Note If the database being tuned is a single file database, **RMU UNLOAD/LOADs** will NOT be used, regardless of the setting for the **TUNE_TECHNIQUE** parameter. **SQL EXPORT/IMPORTs** will be used on all single file databases. If **RMU UNLOAD/LOADs** are chosen as the tuning technique, neither the **RDB\$SYSTEM** storage area nor any **LIST** (segmented string) areas will be tuned. Thus, the only way to tune tables/indexes stored in **RDB\$SYSTEM** or to tune segmented strings is to use **SQL EXPORT/IMPORT**.

Note If you choose to use **RMU/UNLOADs** and **RMU/LOADs** to tune selected tables of a database, items which are defined on those tables (views, constraints, triggers, comments) must be dropped and then re-created when the tables are dropped and re-created. If any views, constraints, or triggers were previously defined with **RDO**, the tuning scripts may fail when the items are re-created using **SQL** statement because DBTune uses the original definition of the item to re-create it. If database items were formerly created using **RDO**, you should review the **SQL** scripts generated by DBTune **BEFORE** executing the **MAIN_DRIVER** command file to ensure that no syntactical violations exist for views, constraints, and/or triggers.

Warning The **RMU** setting will create scripts that require an **INTERACTIVE** or **DEVELOPMENT** license for Rdb/VMS and will fail if you only have a **RUNTIME** license! If the scripts that DBTune generates are executed on a system that only has the **RUNTIME** license for Rdb/VMS, you **MUST** set the **TUNE_TECHNIQUE** to **SQL**.

VALUES: **SQL** Use SQL EXPORT/IMPORT
 RMU Use RMU UNLOAD/LOADs
DEFAULT: **SQL**

(3) **DBDISKS**

The **DBDISKS** parameter determines the number of logical storage devices that are to be used for the database being tuned. “New” storage areas (those that reside in RDB\$SYSTEM, those that contain segmented strings, or those being relocated because STRATEGY = N or R) will be spread over the various disks specified (via the **DBDISKnn** parameters below) to reduce the I/O load on any particular disk. The more disk devices specified, the more a database can benefit from the storage area distribution.

VALUES: 1 to 99 logical devices
DEFAULT: 1

(4) DBDISKnn

The **DBDISKnn** parameters are used in conjunction with the **DBDISKS** parameter above to specify actual disks and directories to be used for spreading storage areas. In addition, the available blocks can be specified for a particular disk as well as the type of storage area files you wish to be placed on that disk. The following naming convention should be used when specifying these parameters: “**DBDISKnn**” where “nn” is “01”, “02”, “03” . . . up to “99”. Both available blocks and storage file qualifiers can be specified after the disk and directory specification using “/”s to separate the values. For example, if the previous parameter was specified as “**DBDISKS=4**”, the actual disks and directories to be used could be specified in the following manner:

```
DBDISK01=DISK1 : [MYDATA.RDB] /25000/SYSRDB/SYSRDA/
DBDISK02=DISK2 : [MORE_DATA] /150000/TBLRDA/IDXRDA/
DBDISK03=DISK3 : [EVEN_MORE] /65000/TBLSNP/IDXSNP/SYSSNP/
DBDISK04=DISK4 : [WHOA]
```

where 25000, 150000 and 65000 are the maximum blocks of free space DBTune is allowed to use on the first three disks, respectively. For the fourth disk, DBTune will attempt to determine the available blocks on the physical disk and use that value. If unable to determine the available blocks, DBTune will assume the fourth disk to have “unlimited” space.

For this example, any **DBDISKnn** parameter greater than **DBDISK04** (e.g., **DBDISK05**) would be ignored because the parameter setting of “**DBDISKS=4**” limits the number of **DBDISKnn** parameters that will be used to four.

By default, the first such **DBDISKnn** parameter (**DBDISK01**) is created for the user and left blank. You can add more **DBDISKnn** parameter lines as required by the **DBDISKS** parameter setting.

Valid storage file qualifiers are as follows:

SYSRDB	.RDB file for the database (root area)
SYSRDA	.RDA file for the database system area (RDB\$SYSTEM)
SYSSNP	.SNP file for the database system area (RDB\$SYSTEM)

TBLRDA	.RDA files for TABLE storage areas
TBLSNP	.SNP files for TABLE storage areas
IDXRDA	.RDA files for INDEX storage areas
IDXSNP	.SNP files for INDEX storage areas

If NO file qualifiers exist on a DBDISKnn line, then ANY type of storage file can be stored there. But, if a file qualifier is specified for a DBDISKnn parameter, then ONLY those types of storage files can be stored in the specified location. A typical use of this feature would be to specify a particular location for the database root file(s) or to force all snapshot files to be placed on a particular disk, etc. The storage file qualifiers shown for the previous four disks specify the following:

DBDISK01	allow ONLY the .RDB and .RDA files for the database system area (root) to be stored here
DBDISK02	allow ONLY .RDA files for table and index storage areas to be stored here
DBDISK03	allow ONLY .SNP files for tables, indexes, and the database system area to be stored here
DBDISK04	allow ANY type of file to be stored here

Following is an example of how to prevent all index files (.RDAs and .SNPs) from being stored on a particular disk while allowing any other type of file to be stored there:

```
DBDISK01=DISK1: [RDB] /SYSRDB/SYSRDA/SYSSNP/TBLRDA/TBLSNP/
```

If the database requires more space than is assigned to the DBDISKnn parameters or storage file qualifiers become too restrictive (causing DBTune to run out of locations to place storage area files), DBTune will create an "OVERFLOW" disk whose default location is the current location of the database .RDB file. You can override this default location in the Disk Utilization file during DBTune processing if the EDIT_FILES parameter is set to Y.

Note DBTune assumes that the fastest disk device will be listed first (DBDISK01), the second fastest disk second (DBDISK02), etc.

Important DBTune considers each of the **DBDISKnn** parameters to be a separate device, even though you may assign them all to the same physical disk. Thus, if the following assignments are made:

```
DBDISK01=DISK1:[MYDATA.RDB] / 10000 /  
DBDISK02=DISK1:[MYDATA.MORE]/ 10000 /
```

and there are only 10000 blocks of free space on the physical disk "DISK1:", DBTune may try to allocate 20000 blocks on this disk because it considers DBDISK01 and DBDISK02 to be separate devices. To correctly limit the allocation to 10000 blocks, you would have to specify 5000 blocks for each DBDISK.

Caution If after-image journaling is enabled for your database, it is highly recommended NOT to specify the disk on which the AIJ file resides as a value for one of the DBDISKnn parameters.

(5) *EDIT_FILES*

The **EDIT_FILES** parameter controls whether or not the DBTune process will pause and allow you to edit the Performance Analysis Data (PAD) file and later, the Disk Utilization file. DBTune automatically generates these files, but you can edit them during the DBTune process to further tailor the information to affect tuning and storage area spreading.

Note If you plan to tailor the **PAD** file and use DBTune to maintain a database, it is recommended to create a ModPAD file rather than repeatedly editing the **PAD** file online during DBTune execution. The ModPAD file facilitates continuous maintenance of the database by “seeding” the online PAD file with values each time DBTune is executed. To repeatedly use the same ModPAD file for a database, the **MODPAD_FILE** parameter below can be utilized.

VALUES:	N	Do NOT Edit Performance Analysis Data File during DBTune session
	Y	Edit Performance Analysis Data File during DBTune session
DEFAULT:	N	

(6) *FREND_EDITOR*

The **FREND_EDITOR** parameter is used to specify the editor to be invoked if the **EDIT_FILES** parameter is set to **Y**. If no editor is specified, then **EDIT/EDT** will be used.

EXAMPLES: EDIT/EDT, EDIT/TPU, etc.

DEFAULT: EDIT/EDT

(7) *MODPAD_FILE*

The **MODPAD_FILE** parameter, if not blank, indicates that DBTune will generate its online **Performance Analysis Data** (PAD) file based on the specified ModPAD file. The ModPAD file is used to “seed” the online **PAD** file, allowing you to consistently use the same performance data to facilitate continuous maintenance of a particular database (in much the same way the **MODPARAMS.DAT** file is used to seed the **PARAMS.DAT** file for the VMS AUTOGEN utility). If the file specified for **MODPAD_FILE** does not yet exist, then DBTune issues a warning message, but attempts to generate a brand new ModPAD file with the name and location specified by this parameter. However, if the file specified for **MODPAD_FILE** does exist, it is used to pre-set values in the online **PAD** file. In addition, if there have been any changes in the database (tables/indexes added or dropped), the specified ModPAD file will be updated with the new items.

EXAMPLE: MODPAD_FILE=PERSONNEL.MODPAD

Important If an “@” character is entered in a changeable column in the ModPAD file, DBTune will interpret the “@” character to mean “replace with current database value.” Thus, if an “@” character is entered in the Number of Recs column for a particular table, DBTune will convert the “@” into the actual cardinality that exists for that table in the database. If a value in a column is a “hard-coded” value (e.g., 25000 records), DBTune will keep that value and NOT override it with the actual database value. Thus, the ModPAD file can contain both variable and static information to be used for the database tuning. See the next page for an example.

Note Changes made to the online **PAD** file during execution of DBTune are **NOT** saved to the ModPAD file!

Note Any **DBDISK** information entered into the ModPAD file is ignored.

Note The filename entered for the **MODPAD** parameter **CANNOT** have the file extension **.PAD** because this may conflict with creation of the associated **PAD** file. To be safe, name ModPAD files with the extension **.MODPAD** or **.MOD**.

The example below shows a ModPAD file with explanations of the various symbols and syntax used:

```
! ModPAD (used to modify online PAD file)
! TABLE Section:
! =====
! "Grw %" column : Table Growth Percentage (values: 0...999)
! "Acc Bia" column : Access Bias (values: 0..100; 0=100% Write, 100=100%Read)
! "Act Lvl" column : Table Activity Level (values: 1..9 with 1=Low,9=Hi)
! "Snp %" column : Snapshot Percentage (values: 0...999)
! "Tun Tbl" column : Tune Table? (values: Y-Yes, N-No; only for RMU/LOADs)
! "Ena Cmp" column : Enable Compression? (values: Y-Yes, N-No)
!
! Database Table Name          Number   Grw Acc Act Snp Tun Ena
! of Recs      %   Bia Lvl % Tbl Cmp
! -----*-----*-*-*-*-*
[ CANDIDATES ]                /    @    /850/ 50/ 5 / 25/ Y / @ /
[ COLLEGES ]                  /    @    / @ / 30/ 1 / 75/ N / @ /
[ DEGREES ]                    /   10000/  0/ 70/ 9 / @ / Y / @ /
[ DEPARTMENTS ]               /    @    / 30/ @ / 7 / 10/ Y / @ /
```

Explanation of the above TABLE entries:

Row 1: Use the values that are “hard-coded”, but replace “Number of Recs” with the actual cardinality for this table.

Row 2: Replace “Number of Recs” and replace “Growth %” with the value specified for the DBTune “GROWTH” parameter. Do NOT tune this table if **TUNE_TECHNIQUE = RMU**.

Row 3: Use “10000” instead of using the actual cardinality for the table and replace “Snapshot %” with the value specified for the DBTune “SNP_PERC” parameter.

Row 4: Replace “Number of Recs” and replace “Access Bias” with the value specified for the DBTune “BIAS” parameter.

```
! INDEX Section:
! =====
! "Idx Typ" column : Index Type (values: S-Sorted, H-Hashed)
! "Act Lvl" column : Index Activity Level (values: 1..9 with 1=Low,9=Hi)
! "Avg Dups" column: Average Duplicates for an index (values: 0..9999)
! "Key Nod" column : Index Key Values Per Node (values: 3...999)
! "Fil %" column   : Index Node Fill Percentage (values: 33...100)
! "Snp %" column   : Snapshot Percentage (values: 0 to 999)
! "Tun Idx" column : Tune Index? (values: Y-Yes, N-No; only for RMU/LOADs)
!
! Database Index Name      Idx Act Avg Key Fil Snp Tun
! -----
! { DEG_COLLEGE_CODE }    / H / 2 / @ / @ / 90/ 25/ Y /
! { DEG_EMP_ID }         / S / 7 / @ / @ / @ /150/ Y /
! { EMP_EMPLOYEE_ID }    / H / 3 / @ / @ / 90/ @ / Y /
! { EMP_LAST_NAME }      / S / 8 / 10 / 17/ 90/400/ Y /
```

Explanation of the above INDEX entries:

Row 1: Use the values that are “hard-coded”, but replace “Avg Dups” with the current average number of duplicates for this index and re-calculate a new value for “Key Values Per Node”.

Row 2: Replace “Avg Dups”, calculate “Key Values Per Node”, and replace “Node Fill %” with the value specified for the DBTune “FILL” parameter.

Row 3: Replace “Avg Dups”, calculate “Key Values Per Node”, and replace “Snapshot %” with the value specified for the DBTune “SNP_PERC” parameter.

Row 4: Use all “hard-coded” values, do not substitute or re-calculate any value for this index.

```
! CLUSTER Section:
! =====
! [DEGREES] & { DEG_COLLEGE_CODE }
! [DEGREES] = { DEG_COLLEGE_CODE }
! CONTRA Section:
! =====
! [ CANDIDATES ]~[ COLLEGES ]
```

(8) DYNAMIC_WORKLOAD_FILE

The **DYNAMIC_WORKLOAD_FILE** parameter, if not blank, indicates that DBTune is to use the Dynamic Workload data file produced by a database monitor process. This file should contain activity weightings for tables and/or indexes and can be created automatically via DBXAct, manually by the user, or by a third-party tool. The information contained is the result of a transaction analysis on a database. It weights the tables and indexes on a scale of 1 to 9, where 9 indicates the highest activity and 1 indicates the lowest. This information is merged with the Performance Analysis Data file during execution of DBTune. The Performance Analysis considers the activity information when optimizing the distribution of storage areas across the available disks. If a file specification is entered for this parameter, the file will be validated for existence and access privilege. The file is expected to be an RMS sequential file (CONTROL=CARRIAGE_RETURN, VARIABLE LENGTH) and if the file is unable to be accessed, DBTune will ignore it.

EXAMPLE: **DYNAMIC_WORKLOAD_FILE=**
 PERSONNEL.DYN_ACT

Note Activity Level values in the **DYNAMIC_WORKLOAD_FILE** override Activity Level values specified in the **MODPAD_FILE**.

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! Dynamic Database Activity
!
! For Database: MY_DATABASE.RDB
!
*SOURCE=MANUAL
*BEGAN = 07-Feb-1995 11:43:27
*ENDED = 14-Feb-1995 15:21:56
!
!
!
!
!
!
! Item Name: [TABLE] or {INDEX}    Activity
!                               Level    Growth%
!                               (9-HI,1-LO) (0-999)    Access Bias
!-----
[EMPLOYEE_TABLE]            /    9    /    10    /    50    /
[INVOICE_TBL]               /    7    /    117   /    25    /
{EMP_NUM_IDX}               /    6    /
{INV_NUM_KEY}               /    2    /
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

(9) SQL_DIR

The **SQL_DIR** parameter assigns the disk and directory where DBTune will create the SQL scripts and DCL command procedures to tune the database.

VALUES: CURRENT (use current default directory)
 disk:[dir] (specify some other directory)

DEFAULT: CURRENT

(10) BACKUP_DIR

The **BACKUP_DIR** parameter assigns a disk and directory that will contain the RMU/backup of the database. You can prevent the backup from occurring by specifying NONE. It is strongly recommended to allow DBTune to perform a backup prior to tuning the database so that recovery is possible in the unlikely event of system failure or some other event that prevents the tuning process from completing successfully. It is recommended you use a scratch disk with abundant free blocks to ensure the tuning process has sufficient space to store the existing database. The **REVIEW_AND_GUIDE.REPORT** created by DBTune contains information on the estimated storage requirements for the backup. Your system manager may be able to provide a more precise storage requirement based on previous backups.

VALUES: NONE (no backup performed)
 CURRENT (use current default directory)
 disk:[dir] (specify some other directory)

DEFAULT: CURRENT

Note You can specify a tape device for this parameter. However, in order for DBTune to validate the parameter, it will attempt to create a file on the specified device. For this to occur, a tape must already be loaded and mounted into the tape drive **PRIOR** to running DBTune. The tape should already have been **INITIALIZE**d and should be **MOUNT**ed as a files device (without the **/FOREIGN** qualifier).

(11) **EXPORT_UNLOAD_DIR**

The **EXPORT_UNLOAD_DIR** parameter assigns a disk and directory that will contain the database **EXPORT** file if performing a SQL **EXPORT/IMPORT** or the **UNLOAD** data files if performing **RMU/UNLOADs** and **LOADs**. The export/unload files are temporary files used to transform the database into the new design. You should you use a scratch disk with abundant free blocks because the export/unload files are an uncompressed copy of the data and may be larger than the actual database.

The **REVIEW_AND_GUIDE.REPORT** file created by DBTune contains information on the estimated storage requirements for the export/unload files. Your system manager may be able to provide a more precise storage requirement based on previous exports/unloads.

VALUES: CURRENT (use current default directory)
 disk:[dir] (use some other directory)

DEFAULT: CURRENT

Note You can specify a tape device for this parameter. However, in order for DBTune to validate the parameter, it will attempt to create a file on the specified device. For this to occur, a tape must already be loaded and mounted into the tape drive **PRIOR** to running DBTune. The tape should already have been **INITIALIZE**d and should be **MOUNT**ed as a files device (without the **/FOREIGN** qualifier).

(12) RUJ_DIR

The **RUJ_DIR** parameter assigns a disk and directory that will contain the **RUJ** (Recovery Unit Journal) file for the **MAIN_DRIVER** process that actually tunes the database. The **RUJ** file is a temporary file used to transform the database into the new design. You should use a scratch disk with abundant free blocks as the **RUJ** file must store a copy of each table row that is being loaded/unloaded or altered and can grow quite large. If the disk that contains the **RUJ** file runs out of free blocks during execution of the tuning procedure, the procedure will fail. To approximate how much disk space is required for the **RUJ** file, estimate the size of the largest table in the database (in blocks) and multiply by two. At a minimum, there should be at least 20000 blocks of free space on the disk assigned to the **RUJ_DIR** parameter.

Note If a VMS logical is used in the **RUJ_DIR** specification, it **MUST** be assigned at the **SYSTEM** level to ensure proper access by the database. If a non-SYSTEM logical is used, the SQL tuning scripts may fail.

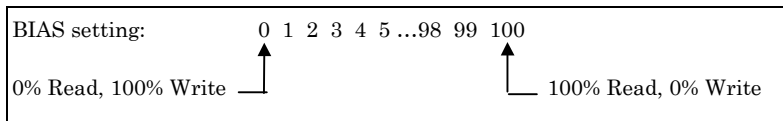
VALUES: CURRENT (use current default directory)
 disk:[dir] (use some other directory)

DEFAULT: CURRENT

(13) BIAS

The **BIAS** parameter controls whether the Performance Analysis Data file is initially loaded with **READ** or **WRITE** biased tables and indexes. **READ** bias will create larger **NODE** sizes and **PAGE** sizes for those storage areas that contain a sorted index, reducing the number of Direct I/Os. **WRITE** bias will create smaller **NODE** and **PAGE** sizes, reducing lock contention, but increasing the number of Direct I/Os. The **BIAS** setting has a global effect on the database, but you can change individual table settings in the **PAD** and/or **ModPAD** files.

The **BIAS** setting represents a sliding scale with 100 percent **WRITE** bias on one end and 100 percent **READ** bias on the other end. For example:



If a table is accessed 50 percent of the time with **READ ONLY** transactions, its rating would be 50. If a table is accessed 80 percent of the time with **READ WRITE** (update/insert/delete) transactions, its rating would be 20 (it is **READ ONLY** 20 percent of the time).

VALUES: 0 to 100

DEFAULT: 50 (50% Read, 50% Write)

(14) **FILL**

The **FILL** parameter value is used to control the initial placement of index records in each index node (for **SORTED** indexes only). Generally, for **READ** biased (retrieval-intensive) indexes, a **FILL** factor greater than or equal to 70 is recommended. For **WRITE** biased (update-intensive) indexes, a **FILL** factor of less than or equal to 70 is recommended. This **FILL** factor is a global setting for the database, but you can change this value for individual indexes in the **PAD** and/or **ModPAD** files.

VALUES: 33 to 100

DEFAULT: 90

(15) **GROWTH**

The **GROWTH** parameter sets the percentage of additional space that will be allocated to existing tables and indexes. For example, a **GROWTH** parameter value of 30 would cause an additional 30 percent to be allocated to the existing cardinality of each table (i.e., allocation would be calculated for 130 records if the table currently contained 100 records). This parameter is a global setting that affects all tables and indexes in the database. You can override this setting for individual tables, however, via the **PAD** and/or **ModPAD** files.

VALUES: 0 to 999

DEFAULT: 10 [recommended to be 30 if ample disk space available]

(16) SNP_PERC

The **SNP_PERC** parameter controls how much space is allocated to the snapshot file for each storage area. **SNP_PERC** represents a percentage of the allocation for the associated **.RDA** file. For example, if an **.RDA** file is allocated 200 pages and **SNP_PERC** is set to 25, then the **.SNP** file will be allocated 50 pages (25 percent of 200). This parameter is a global setting that affects all tables and indexes in the database. You can override this setting for individual tables/indexes, however, via the **PAD** and/or **ModPAD** files.

VALUES: 0 to 500

DEFAULT: 5 [recommended to be 25 if ample disk space available]

(17) MIN_PAGE_SIZE

The **MIN_PAGE_SIZE** parameter controls the minimum page size (in blocks) for database storage areas that are tuned via DBTune. The value may range from 1 to 32 blocks. For certain **READ**-intensive applications, larger page sizes may improve direct I/O performance. For **WRITE**-intensive applications, however, larger page sizes may increase locking contention. The default **MIN_PAGE_SIZE** will produce optimal performance for applications that are divided equally between both **READs** and **WRITEs**. If **TUNE_TECHNIQUE = RMU**, this parameter cannot exceed the existing buffer size of the database.

VALUES: 1 to 32

DEFAULT: 1

(18) MAX_PAGE_SIZE

The **MAX_PAGE_SIZE** parameter controls the maximum page size (in blocks) for database storage areas that are tuned via DBTune. The value may range from 2 to 32 blocks, but cannot be less than **MIN_PAGE_SIZE** and cannot exceed **MAX_BUFFER_SIZE**. For certain READ-intensive applications, larger page sizes may improve direct I/O performance. For WRITE-intensive applications, however, larger page sizes may increase locking contention. The default **MAX_PAGE_SIZE** will produce optimal performance for applications that are divided equally between both READs and WRITEs.

VALUES: 2 to 32

DEFAULT: 32

(19) MIN_BUFFER_SIZE

The **MIN_BUFFER_SIZE** parameter controls the minimum buffer size for a database tuned via DBTune. Buffer size is only calculated for a database when **TUNE_TECHNIQUE = SQL**. The value may range from 3 to 64 blocks. Buffer size impacts the size of process working sets required for database users. A large buffer size or a large number of buffers may require larger working sets, while a small buffer size or a small number of buffers may result in poor direct I/O performance for certain applications. DBTune will automatically calculate a buffer size for a database based on the sizes of its tuned storage areas. The **REVIEW_AND_GUIDE.REPORT** created by DBTune provides information on resources impacted by the resulting buffer usage.

VALUES: 3 to 64

DEFAULT: 6

(20) **MAX_BUFFER_SIZE**

The **MAX_BUFFER_SIZE** parameter controls the maximum buffer size for a database tuned via DBTune. Buffer size is only calculated for a database when **TUNE_TECHNIQUE = SQL**. The value may range from 3 to 64 blocks, but cannot be less than **MIN_BUFFER_SIZE** or **MAX_PAGE_SIZE**. Buffer size impacts the size of process working sets required for database users. A large buffer size or a large number of buffers may require larger working sets, while a small buffer size or a small number of buffers may result in poor direct I/O performance for certain applications. DBTune will automatically calculate a buffer size for a database based on the sizes of its tuned storage areas. The **REVIEW_AND_GUIDE.REPORT** provides information on resources impacted by the resulting buffer usage.

VALUES: 3 to 64

DEFAULT: 64

(21) **MIN_BUFFERS**

The **MIN_BUFFERS** parameter controls the minimum number of buffers assigned to the database; the value may range from 0 to 250. DBTune will start with this number of buffers and add to it based on the resulting physical design. Buffers can reduce disk I/O by utilizing data retrieved via previous I/Os but they also require additional memory from the process working set. DBTune's **REVIEW_AND_GUIDE.REPORT** provides information on resources impacted by the resulting buffer usage.

VALUES: 0 to 250

DEFAULT: 20

(22) **MAX_BUFFERS**

The **MAX_BUFFERS** parameter controls the maximum number of buffers assigned to the database; the value may range from 20 to 1000. DBTune ensures that **MAX_BUFFERS** is greater than or equal to **MIN_BUFFERS**. DBTune's **REVIEW_AND_GUIDE.REPORT** provides information on resources impacted by the resulting buffer usage.

VALUES: 20 to 1000

DEFAULT: 100

(23) **SYS_MEM_PAGES**

The **SYS_MEM_PAGES** parameter controls the number of physical memory pages to be used by a database for global buffers. Global buffers can improve performance by reducing I/O operations and by better utilizing memory. To see how many memory pages are available or “free” on a system, type the following command at the DCL prompt: **\$ SHOW MEM/PHYS**. If the database being tuned can be accessed from more than one node or system, the value entered for this parameter should not exceed the LEAST number of free pages available on the candidate systems. For example, if MY_DATABASE can be accessed from Node A (20000 free pages), Node B (15000 free pages), and Node C (50000 free pages), **SYS_MEM_PAGES** should be set no higher than 15000. If this parameter is set to 0, the existing global buffer settings for the database will be maintained. If this parameter is set to a value greater than 0, global buffers will be enabled for the database and will be set to maximize usage of the specified system memory pages. If the value entered for this parameter is insufficient for a minimum number of buffers (5) to be allocated to the maximum number of database users, an error message will be given and the parameter will be reset to 0. DBTune’s **REVIEW_AND_GUIDE.REPORT** provides information on resources impacted by any resulting buffer usage.

Note Do NOT assign ALL available pages of memory to the **SYS_MEM_PAGES** variable. It is safest to leave some pages in reserve in the event of a system emergency or an abnormally heavy user load where extra pages may be needed.

VALUES: 0 to 999999999999

DEFAULT: 0

(24) MAX_DB_USERS

The **MAX_DB_USERS** parameter controls the maximum number of users allowed to access the database at one time. Each process (online and batch) that attaches to the database is considered a “user.” The number of database users, in conjunction with the **SYS_MEM_PAGES** parameter, affects the number of global buffers allocated to each user. If this parameter is set to 0, the existing number of users for the database will be maintained. If this parameter is set to a value greater than 0, it will override the existing database setting.

VALUES: 0 to 2032

DEFAULT: 0

(25) STOR_AREA_SPREAD

The **STOR_AREA_SPREAD** parameter controls what value will be used to spread database storage areas over the disks specified above. Each storage area is weighted according to this parameter. The weighting is used along with **CONTRA** and **CLUSTER** information to place storage areas on available disks to enable maximum I/O throughput and reduce I/O bottlenecks.

VALUES: **A** Spread areas based on **ACTIVITY**
V Spread areas based on **VOLUME**
B Spread areas based on a factor of **BOTH** Activity and Volume

DEFAULT: **B**

(26) LOGICALS

The **LOGICALS** parameter determines if VMS logicals will be generated for “new” and “relocated” storage area files (RDA and SNP). The logicals will be named using the following format:

```
<<db_id>>_<<storage_area_name>>_RDA for .RDA files and
<<db_id>>_<<storage_area_name>>_SNP
```

for .SNP files

(28) CONCEAL_LOGS

The **CONCEAL_LOGS** parameter indicates whether the logicals that are generated are **CONCEALED** and **ROOTED**.

Note If **CONCEALED** logicals are not used, then RDB translates the logicals to their physical device. In this case, later reassignment of storage areas may require you to execute **RMU/MOVE** manually for each storage area with a non-concealed logical. **CONCEALED** logicals can only be used if they are assigned to rooted directory specifications.

VALUES: **Y** (YES, use CONCEALED, ROOTED logicals)
 N (NO, do NOT use CONCEALED, ROOTED logicals)

DEFAULT: **N**

(29) LOAD_TIME_LIM

The **LOAD_TIME_LIM** parameter determines the number of minutes allowed for performing **RMU UNLOAD/LOADs** on a database. If **LOAD_TIME_LIM** equals zero (0), DBTune assumes it has “unlimited” time to perform the **UNLOAD/LOADs**. If **LOAD_TIME_LIM** is greater than zero, DBTune will determine which tables/indexes will provide the most performance benefit when tuned and then tune as many as possible in the number of minutes specified. A common use for this parameter is to set it for the specific window of time that your system can be down and then see how much of the database can be tuned during that period. You can “pre-select” which tables/indexes will be considered for tuning via a **MODPAD_FILE**. DBTune will ignore all tables whose **Tbl Tun** value is **N** and all indexes whose “**Idx Tun**” value is **N** in the **MODPAD_FILE**. Only those tables/indexes whose “Tune” value is **Y** will be considered for tuning. If the **LOAD_TIME_LIM** parameter is set to a value greater than zero, a Cost/Benefit analysis for items chosen to be tuned within the specified time limit will be printed in the **dbname_DBTUNE.LOG** output file.

Note The **LOAD_TIME_LIM** parameter is used only if **TUNE_TECHNIQUE = RMU**.

Note If you wish to see how long it will take to tune the entire database using **RMU/LOADS**, set **LOAD_TIME_LIM** to 30000, run DBTune, and then read the <<DB_ID>>_DBTUNE.LOG file that DBTune generates.

VALUES: 0 (unlimited) to 30000 minutes

DEFAULT: 0

(30) **MACHINE_VUPS**

The **MACHINE_VUPS** parameter is used **ONLY** when performing **RMU UNLOADS/LOADS** to estimate the total tuning time required. This parameter should represent the **VUPS** rating for a single processor on the system from which tuning will occur. For example, if a database will be tuned on Node A (a VAXstation 3100 with a VUPS rating of 2.5), tuning may take longer than if it occurs on Node B (a MicroVAX 4000-200 with a VUPS rating of 5.0). If a system has multiple processors (i.e., VAX 6440 has four processors of 6 VUPS each for a total of 24 VUPS), it is assumed that the tuning process can only take advantage of a single processor. So, in this case, the **MACHINE_VUPS** parameter would be set to 6 not 24. The more VUPS a processor has, the quicker the tuning process will complete. Disk I/O rates and disk fragmentation can also affect the time required to tune a database; however, these variables are considered uniform at the present time and are not allowed to be specified.

VALUES: 1.0 to 999.9 VUPS

DEFAULT: 2.5

(31) TABLE_COMMIT

The **TABLE_COMMIT** parameter is used **ONLY** when performing **RMU UNLOADs/LOADs**. One of the intermediate steps involved in tuning via **RMU/LOADs** is to delete (drop) the tuned tables from the database and then add them back. When a table is dropped from the database, the **RUJ** file can grow quite large. To keep the size of the **RUJ** file down to a manageable size, a **SQL COMMIT** can be performed after every **DROP TABLE**. This has the disadvantage, however, of not allowing a complete rollback if a later table drop fails. On the other hand, if all the tables are dropped and then a **COMMIT** is performed once at the end, a complete rollback can be performed but the **RUJ** file may grow so large that it runs out of disk space, causing the tuning scripts to fail. Thus, if the tables being tuned are very large (e.g., have over 500000 records) and/or disk space is limited, it might be a good idea to **COMMIT** after **EVERY** table is dropped. However, if tables are not unusually large, or if disk space is abundant, or if you want all changes to be rolled back if an error occurs, drop all the tables and perform a **COMMIT** once at the end.

VALUES:	N	NO , do NOT commit after every table drop (results in large RUJ)
	Y	YES , commit after EVERY table drop (unable to rollback fully in case of error)
DEFAULT:	Y	

(34) SMALL_TABLE

The **SMALL_TABLE** parameter is used to set the common storage area name for tables that have a record count less than the **SA_MIN_CARD** parameter. **IT IS HIGHLY RECOMMENDED THAT THE USER PRECEDE THE “SMALL_TABLE” NAME WITH THE NAME OF THE DATABASE IF STORING MORE THAN ONE DATABASE IN THE SAME DIRECTORY!** For example, if storing the **INVOICE** database and the **ORDERS** database in the same directory, the values for this parameter might be “**INVOICE_SMALL_TABLE**” or “**ORDERS_SMALL_TABLE**”, depending on which database was being tuned during a session.

Note The storage area name must start with an alpha character (A . . Z).

VALUES: Any valid 31-character Rdb storage area name

DEFAULT: **SMALL_TABLE_AREA**

(35) SMALL_SORTED

The **SMALL_SORTED** parameter is used to set the common storage area name for sorted indexes of tables that have a record count less than the **SA_MIN_CARD** parameter. **IT IS HIGHLY RECOMMENDED THAT YOU PRECEDE THE “SMALL_SORTED” NAME WITH THE NAME OF THE DATABASE IF STORING MORE THAN ONE DATABASE IN THE SAME DIRECTORY!** For example, if storing the **INVOICE** database and the **ORDERS** database in the same directory, the values for this parameter might be “**INVOICE_SMALL_SORTED**” or “**ORDERS_SMALL_SORTED**,” depending on which database was being tuned during a session.

Note The storage area name must start with an alpha character (A . . Z).

VALUES: Any valid 31-character Rdb storage area name

DEFAULT: **SMALL_SORTED_AREA**

(36) SMALL_HASHED

The **SMALL_HASHED** parameter is used to set the common storage area name for hashed indexes of tables that have a record count less than the **SA_MIN_CARD** parameter. **IT IS HIGHLY RECOMMENDED THAT YOU PRECEDE THE “SMALL_HASHED” NAME WITH THE NAME OF THE DATABASE IF STORING MORE THAN ONE DATABASE IN THE SAME DIRECTORY!** For example, if storing the **INVOICE** database and the **ORDERS** database in the same directory, the values for this parameter might be “**INVOICE_SMALL_HASHED**” or “**ORDERS_SMALL_HASHED**,” depending on which database was being tuned during a session.

Note The storage area name must start with an alpha character (A . . Z).

VALUES: Any valid 31 character Rdb storage area name

DEFAULT: **SMALL_HASHED_AREA**

(37) TUNE_FOR_COMPRESSION

The **TUNE_FOR_COMPRESSION** parameter controls the values used for tuning storage areas: compressed or uncompressed data values. The default is to use uncompressed data values. If you want compressed values to be used for tuning, this parameter can be changed to **Y**. Tuning using compressed values may save disk space and will use the average compressed record sizes for those tables that have compression enabled. However, tuning based on compressed values may also result in increased record fragmentation as the database is used. In addition, using compressed values requires more processing time by DBTune than using uncompressed values (see the note below). Tuning using uncompressed values will reduce record fragmentation during database usage and is quicker to calculate.

Note Setting the value of this parameter to **Y** will require more processing time for DBTune to determine actual compression values of data stored in the database. For databases that are several gigabytes in size, this option could add several hours of processing time to DBTune.

VALUES: **Y** **YES**, Tune using COMPRESSED data values
 N **NO**, Tune using UNCOMPRESSED data values

DEFAULT: **N**

After the DBTune parameter file has been read in and validated, the values are displayed for acceptance or online editing:

```

*** DBTune V5.2 ***
Rdb Name: DSA105: [NRCTANDEM] PROD_DB.RDB
Current DBTune Parameter Settings
--- PRODB.PARAMS ---
Strategy = N, Min Card = 200
Technique = RMU, Growth % = 20
# DBDisks = 5, Snapshot % = 10
Edit Files = N, AccessBias = 70
RMULoadTime = 0, NodeFill % = 90
MachineVUPs = 2.5, Logicals = N
TableCommit = Y, Logcl Type = PROCESS
SaveComment = Y, Concealed = N
MinPageSz = 1, MaxPageSz = 32
MinBuffSz = 6, MaxBuffSz = 64
Min Buffs = 20, Max Buffs = 100
SysMemPages = 0, MaxDBUsers = 0
ModPAD file = HRCDSK07: [PROD.DATABASE.DBTUNE.APR
SQL Dir = DB$DISK: [PROD.DATABASE.DBTUNE.APR
Backup Dir = NONE
Exp/Unl Dir = DSA32: [UNLOAD]

EXIT_SELECTION_MENU
Strategy
Tune_Technique
DBDisks
Edit_Files
ModPad_File
Dynamic_Workload_File
SQL_Dir
Backup_Dir
Export_Unload_Dir
RUJ_Dir
Bias
Fill
Growth
SNP_Perc
Min_Page_Size
Max_Page_Size
Min_Buffer_Size
Max_Buffer_Size
Min_Buffers

[DD]-Create Tuning Scripts [SELECT]-Edit Parameters [HELP]-Help
[ESC]-Exit

```

Note The actual parameter values used for a DBTune session are recorded in the DBTune LOG file.

Step 3: Read Database Structure

DBTune automatically loads the database structure. There is no input required for this phase. The status window provides progress information during this process.

```

*** DBTune V5.2 ***
Rdb Name: DSA105:[NRCTANDEM]PROD_DB.RDB
Reading Rdb Database
Reading Rdb Structures...
Reading 107 tables and views...
(6) Table: AR_SUPPORT_DATA

Scan Progress
#####

Tables      : 107
Indices     : 123
Domains     : 69
Columns     : 1695
Storage Area: 167
Views       : 0
Constraints : 134
Triggers    : 2
Records     : 129,518,679
Avg Rec/Tbl : 1,210,455

Complexity Rating: 38
#####

Tune Rating: 98
#####

```

Step 4: Read Customized Analysis (MODPAD) / Workload Data

If a valid filename was specified for the DBTune parameter **MODPAD_FILE** or the parameter **DYNAMIC_WORKLOAD_FILE**, those files are parsed for validity. In the case of the **MODPAD_FILE**, any variable data (columns containing an “@” character) is filled in with the actual values for the database. (For more information on variable data, see explanation of the **MODPAD_FILE** parameter on page 88.) The results of this step are used to create DBTune’s Performance Analysis Data (PAD) file. Any errors that occur during this step will be placed in the DBTune log file.

Note If DBTune detects errors when parsing the Dynamic Workload or ModPAD file you will be given the opportunity to view a log of the errors found and to make corrections. You may also elect to continue without making corrections, allowing DBTune to substitute default values for erroneous settings.

Step 5: Generate Performance Analysis Data (PAD) File

DBTune creates the **Performance Analysis Data (PAD)** file using the logical and physical database design, dynamic activity, transaction analysis results, DBTune parameters, and ModPAD file. You can dynamically change all column values in the PAD file during an editing session except for those columns below that explicitly state “DO NOT CHANGE!” The PAD file consists of five sections:

DBDISKS

Assignment of physical disk devices to be used when spreading storage areas. It contains three values:

Logical Disk Name DO NOT CHANGE!

Physical Specification Disk and directory specification for the logical disk.

Available Blocks Disk blocks available for database storage areas.

TABLE

Information for database tables used during tuning. It contains seven values:

Table Name DO NOT CHANGE!

Cardinality Current number of rows in table. Can be modified to desired volume.

Growth Cardinality will increase (grow) by this percentage.

Access Bias 0..100, 0 = 100% Write bias, 100 = 100% Read bias. 50 = 50% Write, 50% Read.

Activity Level 1-9, 1 indicates the lowest activity, 9 indicates the highest activity. Used to optimize disk utilization (if STOR_AREA_SPREAD is A or B).

Snapshot % Allocation for table’s snapshot file will be set to a percentage of the table’s data (.RDA) file allocation.

Tune Table Y/N, Tune the table? Used only when performing RMU/LOADs (**TUNE_TECHNIQUE** parameter = **RMU**). If performing **RMU/LOADs** and Tune Table is set to **N**, table will be skipped during the tuning process.

Enable Compression Y/N, Enable Compression? Will enable (Y)/disable (N) compression for the individual table. By default, this is set to the current compression setting for the table in the database.

INDEX

Information for database indexes used during tuning. It contains eight values:

Index Name DO NOT CHANGE!

Index Type S/H (Sorted or Hashed). This value may be changed to alter the index type.

Activity Level 1-9, 1 indicates the lowest activity, 9 indicates the highest activity. Used to optimize disk utilization (if **STOR_AREA_SPREAD** is A or B).

Average Duplicates Current number of average duplicates for the index.

Key Values/Node Indicates the number of records that will fit in the node size that DBTune will calculate for this sorted index. Can be changed to increase or decrease node size.

Node Fill Fill factor to be used for the sorted index.

Snapshot % Allocation for index's snapshot file will be set to a percentage of the index's data (**.RDA**) file allocation.

Tune Index Y/N, Tune the index? Used only when performing **RMU/LOADs** (**TUNE_TECHNIQUE** parameter = **RMU**). If performing **RMU/LOADs** and Tune Index is set to **N**, index will be skipped during the tuning process.

CLUSTER

Used to indicate tables and/or indexes to be clustered together in the same storage area. May also specify that a table is to be **PLACED VIA** an index.

To Cluster two objects together:
 [EMPLOYEES] & {EMP_EMPLOYEE_ID}

To have a table PLACED VIA an index:
 [EMPLOYEES] = {EMP_EMPLOYEE_ID}

Note It is not necessary to cluster a table and an index together to use the **PLACE VIA** option. When a table is **PLACED VIA** an index, the page size is calculated to hold the number of table records equal to the average number of duplicates for the specified index. For example, if the EMPLOYEE_HISTORY table is PLACED VIA the EMPLOYEE_HIST_IDX and the EMPLOYEE_HIST_IDX has an average duplicate value of 15, then the page for the EMPLOYEE_HISTORY table will be sized to hold at least 15 records.

CONTRA

Used to indicate tables and/or indexes that are NOT to be stored on the same disk.

To indicate that the storage areas for 2 objects should be on separate disks:
 [EMPLOYEES]~[RESUMES]

```

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
! PAD (On-Line Performance Analysis Data--DBTune V5.2)
!
! ( ModPAD Input File : NRCDSK07:[PROD.DATABASE.DBTUNE.APR01]PRODDB.MDD_PAD )
!
! Listed below are values that will be used for tuning storage areas
! for database: PROD_DB
! The user may change values for those columns underlined with "---".
! When finished, exit and save this file to continue the DBTune process.
! NOTE: Changes made in this file are NOT copied to the ModPAD file.
!
! Following are RESERVED characters and their interpretations:
! "!" : comment line (if ! is the 1st char, entire line ignored)
! "/" : column separator, MUST exist btw column values, e.g. col1/col2/.../
! "[ ]": the enclosed item is a database TABLE, e.g., [EMPLOYEE_TABLE]
! "{ }": the enclosed item is a database INDEX, e.g., [EMPLOYEE_INDEX]
! "&" : CLUSTER symbol, the item following this symbol will be stored in the
! same storage area as the item preceding this symbol
! "=" : PLACE VIA symbol, [table_name]={index_name} indicates to PLACE the
! table VIA the index (index must belong to the table)
! "~" : CONTRA symbol, the item following this symbol will NOT be stored on
!
! Buffer: PROD_DB.PAD ; Write ; Insert ; Forward
!
374 lines read from file NRCDSK07:[PROD.DATABASE.DBTUNE]PROD_DB.PAD:13

```

You may edit the **PAD** file online during the DBTune process to customize the performance analysis criteria for their particular applications and their environment. This can be accomplished by setting the **EDIT_FILES** parameter to **Y**. In addition to the typical transaction parameters such as workload and volume data, the performance analysis considers the logical design to create an optimized Rdb physical design. Thus, the database is physically structured to achieve optimal performance according to its planned and actual usage.

Note If DBTune detects errors when parsing the **PAD** file, you will be given the opportunity to view a log of the errors found and to make corrections. You may also elect to continue without making corrections, allowing DBTune to substitute default values for erroneous settings.

The **PAD** file may be generated anew every time DBTune is run or the **PAD** file can be “saved” in a **ModPAD** file in order to reproduce a given tuning strategy. To control how much of a **PAD** is generated each time DBTune is run, the following parameters can be set:

MODPAD_FILE

If left blank, DBTune generates the **PAD** file using default or parameter values. If assigned a valid file specification, DBTune uses the named ModPAD file to “seed” the values in the new **PAD** file, substituting actual database values where variable data is indicated (“@”) and using hard-coded values otherwise.

DYNAMIC_WORKLOAD_FILE

If a valid file specification is assigned to this parameter, activity information in the named file will be used as input to the **PAD** file and will override any activity information found in the **MODPAD_FILE**.

The following characters are considered **RESERVED**, and DBTune uses them to interpret the **Performance Analysis Data** file:

RESERVED	INTERPRETATION

!	Comment line, entire line ignored
/	Column separator, e.g., col1/col2/.../
*DISK	Indicates this is a physical disk assignment line, e.g., *DISK01/DISK1:[MYDATA.RDB]/2500/
[]	Enclosed item is an Rdb table name (except in the case of a disk assignment), e.g., [EMP_TABLE]
{ }	Enclosed item is an Rdb index name, e.g., {EMP_INDEX}
&	Used to CLUSTER two items in the same storage space e.g., [TABLE_1] & [INDEX_2]
=	Used to specify that a table is to be PLACED VIA an index, e.g., [INVOICE_TABLE]=[INDEX_IDX]
~	Used to force two items to be stored on separate disks (a "CONTRA"), e.g., [TABLE_1]~-[TABLE_2]

```

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
! Performance Analysis Data
!
! Listed below are values that will be used for tuning storage areas
! for database: PERSONNEL
! The user may change values for those columns underlined with "---".
! When finished, exit and save this file to continue the DBTune process.
!
! Following are RESERVED characters and their interpretations:
! "!" : comment line (if ! is the 1st char, entire line ignored)
! "/" : column separator, MUST exist btw column values, e.g., col1/col2/.../
! "[ ]": the enclosed item is a database TABLE, e.g., [EMPLOYEE TABLE]
! "{ }": the enclosed item is a database INDEX, e.g., {EMPLOYEE_INDEX}
! "&": CLUSTER symbol, the item following this symbol will be stored in the
! same storage area as the item preceding this symbol
! "=" : PLACE VIA symbol, [table name] = {index_name} indicates to PLACE
! the table VIA the index (the index MUST belong to the table)
! "~" : CONTRA symbol, the item following this symbol will NOT be stored on
! the same disk as the item preceding this symbol
! "*DISK" : indicates this is a physical disk assignment line
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
! DBDISKS Section:
! =====
!
! The parameter 'DBDISKS' indicates 1 disk is available for use.
! For each disk, a physical DISK:[DIR] specification and available free
! blocks can be entered. If a disk and directory are NOT entered for each
! disk, the disk and directory where the .RDB file resides will be used
! (i.e., ISG$DATA1:[FREND.RDB] ).
! The number of free blocks for each disk is set to 'unlimited' by default.
! If the physical disk specification you enter is accessible to this
! process and is set to 'unlimited', the free space will be derived
! automatically. Or, you can set your own limit to the free space that

```

```

! will be used by entering a number in the 'Free Blocks' column below.
!
! Following are two examples:
! DISK01 / $1$DUA1:[DATA.RDB] / unlimited / <-> will be calc'd during tuning
! DISK02 / $1$DUA2:[MORE_DATA] / 150000 /
!
! * * * * * NOTE * * * * *
! Enter physical disk and directory specifications in order of disk speed
! and disk utilization: the fastest disk that has the least amount
! of I/O requests should be entered for 'DISK01', the next fastest disk
! should be 'DISK02', etc.
!
!
! Physical Disk & Directory Free Blocks
! -----*-----*-----
!*DISK01 /ISG$DATA2:[MAINT.ACCTNG.RDB] / 104971/
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

```

(Continued)

```

!
! TABLE Section:
! =====
!
! Any of the following table information (except the TABLE NAME) can be
! changed to customize the Performance Analysis and tuning processes.
!
! "Grw %" column : Table Growth Percentage (values: 0...999)
! "Acc Bia" column : Access Bias (values: 0..100; 0=100% Write, 100=100% Read)
! "Act Lvl" column : Table Activity Level (values: 1..9 with 1=Low,9=Hi)
! "Snp %" column : Snapshot Percentage (values: 0...999)
! "Tun Tbl" column : Tune Table? (values: Y-Yes, N-No; only for RMU/LOADs)
! "Ena Cmp" column : Enable Compression? (values: Y-Yes, N-No)
!
!
! Database Table Name          Number  Grw Acc Act Snp Tun Ena
! of Recs  %  Bia Lvl %  Tbl Cmp
! -----*-----*-*-*-*-*
[ CANDIDATES ]                /      3/ 10/ 50/ 5 / 10/ Y / N /
[ COLLEGES ]                  /      15/ 10/ 50/ 5 / 10/ Y / Y /
[ DEGREES ]                   /     165/ 10/ 50/ 5 / 10/ Y / N /
[ DEPARTMENTS ]              /      26/ 10/ 50/ 5 / 10/ Y / Y /
[ EMPLOYEES ]                 /     102/ 10/ 50/ 5 / 10/ Y / Y /
[ JOBS ]                      /      15/ 10/ 50/ 5 / 10/ Y / Y /
[ JOB_HISTORY ]              /     274/ 10/ 50/ 5 / 10/ Y / Y /
[ RESUMES ]                   /       3/ 10/ 50/ 5 / 10/ Y / Y /
[ SALARY_HISTORY ]           /     729/ 10/ 50/ 5 / 10/ Y / Y /
[ WORK_STATUS ]              /       3/ 10/ 50/ 5 / 10/ Y / Y /
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
! INDEX Section:
! =====
!
! Any of the following index information (except the INDEX NAME) can be
! changed to customize the Performance Analysis and tuning processes.
!
! "Idx Typ" column : Index Type (values: S-Sorted, H-Hashed)
! "Act Lvl" column : Index Activity Level (values: 1..9 with 1=Low,9=Hi)
! "Avg Dups" column: Average Duplicates for an index (values: 0...9999)
! "Key Nod" column : Index Key Values Per Node (values: 3...999)
! "Fil %" column : Index Node Fill Percentage (values: 33...100)
! "Snp %" column : Snapshot Percentage (values: 0...999)
! "Tun Idx" column : Tune Index? (values: Y-Yes, N-No; only for RMU/LOADs)
!
!
! Database Index Name          Idx Act Avg  Key Fil Snp Tun
! Typ Lvl Dups Nod %  % Idx
! -----*-----*-*-*-*-*
{ COLL_COLLEGE_CODE }        / S / 5 / 0/ 25/ 90/ 10/ Y /
{ DEG_COLLEGE_CODE }        / S / 5 / 12/ 25/ 90/ 10/ Y /
{ DEG_EMP_ID }               / S / 5 / 1/ 25/ 90/ 10/ Y /
{ DEPARTMENTS_INDEX }       / S / 5 / 0/ 25/ 90/ 10/ Y /
{ EMPLOYEES_HASH }          / H / 5 / 0/ - / - / 10/ Y /
{ EMP_EMPLOYEE_ID }         / S / 5 / 0/ 25/ 90/ 10/ Y /
{ EMP_LAST_NAME }           / S / 5 / 1/ 17/ 90/ 10/ Y /
{ JH_EMPLOYEE_ID }          / S / 5 / 2/ 25/ 90/ 10/ Y /
{ JOB_HISTORY_HASH }        / H / 5 / 2/ - / - / 10/ Y /
{ SH_EMPLOYEE_ID }          / S / 5 / 7/ 25/ 90/ 10/ Y /

```

(Continued)

```

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
! CLUSTER Section:
! =====
!
! Enter below specific pairs of items that you wish to be clustered together
! in the same storage area. The first item in the pair will determine the
! storage area to be used. So, to cluster three tables together in the
! same storage area, the following two lines could be entered:
!   [TABLE_A] & [TABLE_B]   meaning "store TABLE_B in TABLE_A's storage area"
!   [TABLE_A] & [TABLE_C]   meaning "store TABLE_C in TABLE_A's storage area"
! Only two items can be listed per line and they must be separated by a "&".
! REMEMBER: Tables must be surrounded by []'s and indexes surrounded by {}'s.
!
! For each pair of items listed, cluster 'item2' in the 'item1' storage area:
! -----
[EMPLOYEES] & {EMPLOYEES_HASH}
[EMPLOYEES] = {EMPLOYEES_HASH}
[EMPLOYEES] & {JOB_HISTORY_HASH}
[JOB_HISTORY] = {JOB_HISTORY_HASH}
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
! CONTRA Section:
! =====
!
! Enter below specific pairs of items that you do NOT wish to be stored on
! the same disk (a pair of such items is called a "CONTRA"). For example,
! to ensure that a particular table and index are not stored on the same
! disk, the following line could be entered (without the "!"):
!   [TABLE_A] ~ {INDEX_B}
! meaning "do not store TABLE_A and INDEX_B on the same disk."
! RDA's and SNP's are automatically stored on separate disks, if possible,
! so the user does not have to specify a CONTRA for this type of condition.
! Only two items can be listed per line and they must be separated by a "~".
! REMEMBER: Tables must be surrounded by []'s and indexes surrounded by {}'s.
!
! Do NOT store the following pairs of items on the same disk:
! -----
[EMPLOYEES]~[RESUMES]
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

```

Step 6: Performance Analysis

The **Performance Analysis** reads the **PAD** file and the logical and physical design of Rdb. Then it combines these inputs to create a new physical design for the Rdb database. This design is optimized to increase performance by considering the interrelations of the database and its application usage and operating environment. The applications can be used as they are without application programming changes. Recommendations for setting affected **VMS AUTHORIZE** or **SYSGEN** quotas are provided in the **REVIEW_AND_GUIDE.REPORT** created by DBTune for this design. There is no input required for this phase. The status window provides progress information on the tuning process.

```
*** DBTune V5.2 ***
Rdb Name: DSA105:[NACTANDEM]PROD_DB.RDB
Summarizing Results
Incorporating Disk Utilization data...
Summarizing Database Modifications...
-Index changes (node size, type)
-Preserve table items
  Modules to be dropped...

Tables      : 107
Indices     : 123
Domains    : 69
Columns     : 1695
Storage Area: 167
Views      : 0
Constraints : 134
Triggers    : 2
Records     : 129,531,060
Avg Rec/Thl : 1,210,571

Complexity Rating:30
#####

Tune Rating:98
#####
```

Step 7: Generate Disk Utilization (DISKUTIL) File

The **Disk Utilization** file is created after the performance analysis and database tuning have optimized the new design. All “new” or “relocated” storage areas are assigned to one of the logical (DBDISK) devices while “existing” areas are placed in their original locations and will not show up in this file. In addition to listing which storage areas are assigned to which disk, the total blocks assigned to each device is given.

If the **EDIT_FILES** parameter is set to **Y**, DBTune pauses and allows you to edit the device assignments. If **EDIT_FILES = N**, the editing session will be skipped.

Note If DBTune detects errors when parsing the Disk Utilization file, you will have the opportunity to view a log of the errors found and to make corrections. You may also elect to continue without making corrections, allowing DBTune to substitute default values for erroneous settings.

Note After tuning calculations are complete, the disk utilization file is edited. You must ensure that there will be sufficient space to hold any storage areas that have been reassigned. DBTune will not override any assignments that are made during this edit session.


```

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
! Disk Utilization Data
!
! Listed below are values used to control the placement of "new" and "relocated" storage areas
! for the database: "MF_PERSONNEL" The Physical Disk & Directory specification and Free Blocks
! are listed again for the DBDISKS, but the values in these columns cannot be changed.
! The only values that can be changed in this file are those in the 'Disk' column, which can be
! found in the STORAGE AREA ASSIGNMENT section below. After editing, exit and save this file to
! continue the DBTune process.
!
! Following are the disk specifications chosen previously:
!
!
!           Physical Disk & Directory           Free Blocks
!-----*-----
!DISK01    /ISG$DATA1:[MAINT.ACCTNG.RDB]      /      104971/
!
! STORAGE AREA ASSIGNMENT Section:
! =====
!
! Following are disk assignments for storage areas that are new or are being relocated.
! The 'Disk' column may be changed by the user, but the user is responsible for ensuring
! adequate disk space exists for any changes made. Each line that begins with "!" is
! considered a comment and will be ignored. Every other line will be considered a storage
! area disk assignment.
!
!
!           Area Allocation
! Storage Area File      Type      (blocks)      Disk
!-----*-----
MF_PERSONNEL            /RDB /           276/ DISK01    /
RDB$SYSTEM              /RDA /          3084/ DISK01    /
EMPIDS_OVER             /RDA /           184/ DISK01    /
SALARY_HIS_TBL         /RDA /           120/ DISK01    /
SH_EMPLOYE_IDX         /RDA /           102/ DISK01    /
JH_EMPLOYE_IDX         /RDA /            96/ DISK01    /
DEGREES_TBL            /RDA /            64/ DISK01    /
DEG_COLLEG_IDX         /RDA /            48/ DISK01    /
DEG_EMP_ID_IDX        /RDA /            42/ DISK01    /
ACCTNG_SMALL_TABLE_AREA /RDA /           260/ DISK01    /
EMP_EMPLOY_IDX         /RDA /            39/ DISK01    /
EMP_LAST_N_IDX        /RDA /            42/ DISK01    /
ACCTNG_SMALL_SORTED_AREA /RDA /            36/ DISK01    /
COLL_COLLE_IDX        /RDA /            33/ DISK01    /
RDB$SYSTEM             /SNP /           312/ DISK01    /
EMPIDS_OVER            /SNP /            19/ DISK01    /
SALARY_HIS_TBL        /SNP /            24/ DISK01    /
SH_EMPLOYE_IDX        /SNP /            18/ DISK01    /
JH_EMPLOYE_IDX        /SNP /            18/ DISK01    /
DEGREES_TBL           /SNP /            24/ DISK01    /
DEG_COLLEG_IDX        /SNP /            18/ DISK01    /
DEG_EMP_ID_IDX        /SNP /            18/ DISK01    /
ACCTNG_SMALL_TABLE_AREA /SNP /            32/ DISK01    /
EMP_EMPLOY_IDX        /SNP /            18/ DISK01    /
EMP_LAST_N_IDX        /SNP /            18/ DISK01    /
ACCTNG_SMALL_SORTED_AREA /SNP /            18/ DISK01    /
COLL_COLLE_IDX        /SNP /            18/ DISK01    /
!
! * DISK01:  27 Stor Files,  Used:          4981 out of 104971 blocks
!
!
! ** Placed:  27 Stor Files,  Used:          4981 total blocks
!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

```

Step 8: Generate Rdb Transformation Procedure

DBTune automatically generates a procedure that you can then execute at a later time to transform the original database into its tuned structure. There is no input required for this phase. The status window provides progress information on the generation process.

```

*** DBTune V5.2 ***
Rdb Name: DSA105:[NRACTANDEM]PROD.DB.RDB
Summarizing Results
Incorporating Disk Utilization data...
Summarizing Database Modifications...
-Index changes (node size, type)
-Preserve table items
-Storage area changes
-Sequencing database changes ( 269)
Generating Transformation Procedures...
-Generating transformation MAIN_DRIVER
-Generating SQL optimization scripts

Scan Progress
#####
Tables      : 107
Indices     : 123
Domains     : 69
Columns     : 1695
Storage Area: 167
Views       : 0
Constraints : 134
Triggers    : 2
Records     : 129,531,062
Avg Rec/Tbl : 1,210,571

Complexity Rating: 38
#####

Tune Rating: 98
#####

```

Upon completion, the transformation procedure, advice, tuning, and documentation files are presented.

```

*** DBTune V5.2 ***
Rdb Name: DSA105:[NRCTANDEM]PROD_DB.RDB          00/00/00 23:43
          DBTune Process Complete          Statistics
      DBTune successfully created the
      following transformation procedure:
      PROD_DB_TRANSFORM.MAIN_DRIVER
      in DB$DISK:[PROD.DATABASE.DBTUNE.APR01.SQL2]
      =====
      Four reports were created in the directory:
      DB$DISK:[PROD.DATABASE.DBTUNE]
      PROD_DB_REVIEW_AND_GUIDE.REPORT
      (Tuning: results, advice, errors)
      PROD_DB_TUNING_DETAIL.REPORT
      (Tuning detail, B-trees and more)
      PROD_DB_ANALYSIS.REPORT
      (Narrative analysis of database )
      PROD_DB_DBTUNE.LOG
      (DBTune log file: params used, errors)

      Tables      : 187
      Indices     : 123
      Domains     : 69
      Columns     : 1695
      Storage Area: 167
      Views       : 0
      Constraints : 134
      Triggers    : 2
      Records     : 129,518,679
      Avg Rec/Tbl : 1,210,455

      Complexity Rating:38
      #####

      Tune Rating:98
      #####

Cleaning up scratch area ...

```

Step 9: Transform Rdb

The output created by the DBTune process consists of a set of reports and a transformation procedure that tunes the database. **IT IS STRONGLY SUGGESTED THAT THE REPORTS BE REVIEWED PRIOR TO EXECUTION OF THE TRANSFORMATION PROCEDURE.** These reports provide an analysis of the database as well as tuning advice, warnings about potential problems, and instructions on how to execute the transformation procedure. Examples of these reports can be found on the following pages.

When all the reports have been reviewed, you must ensure that any VMS logicals currently required by the database have been assigned prior to the database transformation or the transformation procedure may fail!

After reviewing the reports and assigning any necessary database logicals, the database can be tuned by executing the transformation **MAIN_DRIVER** command file. This command file is called **<<database_id>>_TRANSFORM.MAIN_DRIVER** and can be found in the directory assigned to the **SQL_DIR** parameter. The **MAIN_DRIVER** is a **DCL** command file that will assign necessary VMS logicals and execute DCL and SQL scripts to tune the target database.

The **MAIN_DRIVER** command file can be executed either online or in batch. If executed online, the procedure asks a series of questions and provides warnings if privilege is insufficient or if special actions need to be taken prior to tuning.

ALL USERS MUST BE OUT OF THE DATABASE WHEN THE MAIN_DRIVER IS EXECUTED OR THE SCRIPTS WILL FAIL!

Following is an example of an execution of the **MAIN_DRIVER** procedure for the MF_PERSONNEL database:

```

@mf_personn_transform.main_driver
=====
You are executing the MF_PERSONN_TRANSFORM.MAIN DRIVER command file.
This command file will execute scripts to modify the following database:
    MF_PERSONNEL

If this database was previously defined using VMS logicals, those
logicals MUST exist for this process or the scripts will fail.
=====
Press <<RETURN>> to continue, <<CTRL>>-Z to exit...

=====

For this tuning session, an SQL EXPORT/IMPORT will be performed
on your database. Before doing so, however, an RMU/BACKUP of
your database will be taken and placed in the backup directory:
    ISG$DATA1:[GISBCK]

NOTE:
-----
Because many tuning changes are not journaled and the
After-Image Journal (.AIJ) may slow down execution of the
tuning process, the current AIJ file for this database will
be DISABLED prior to tuning and then re-enabled after tuning
has completed.
=====
Press <<RETURN>> to continue, <<CTRL>>-Z to exit...

=====

WARNING:
-----
This database was specified with the 'OPEN IS MANUAL' option. If
it was manually opened with an RMU/OPEN command, it must be closed
on all nodes with an RMU/CLOSE/CLUSTER command before continuing,
otherwise database modifications may fail. The modifications
may also fail if there are users or BATCH processes that are
accessing the database during execution of these scripts.
If either of these situations exist, exit this MAIN_DRIVER
procedure now and correct before continuing.
=====
Press <<RETURN>> to continue execution, <<CTRL>>-Z to exit...

=====

No more input is required from the user.

Execution of scripts beginning...

=====

```

(Continued)

The database will now be closed with an RMU/CLOSE command.
Any processes currently attached to the database will be exited.

```

Performing an RMU/BACKUP of your database...
$ rmu/backup MF_PERSONNEL -
  ISG$DATA1:[GISBCK]MF_PERSONNEL.RBF
%RMU-I-BCKTXT_01, Thread 1 uses devices ISG$DATA1:
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]MF_PERS_DEFAULT.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]SALARY_HISTORY.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]RESUME_LISTS.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]EMPIDS_OVER.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]EMPIDS_LOW.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]EMPIDS_MID.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]EMP_INFO.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]RESUMES.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]JOBS.RDA;1
%RMU-I-BCKTXT_08, Thread 1 was assigned file ISG$DATA1:[RDB]DEPARTMENTS.RDA;1
%RMU-I-BCKTXT_00, Backed up root file ISG$DATA1:[RDB]MF_PERSONNEL.RDB;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]MF_PERS_DEFAULT.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMPIDS_LOW.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMPIDS_MID.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMPIDS_OVER.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]DEPARTMENTS.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]SALARY_HISTORY.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]JOBS.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMP_INFO.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]RESUME_LISTS.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]RESUMES.RDA;1
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]MF_PERS_DEFAULT.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 6 inventory pages
%RMU-I-BCKTXT_06, backed up 129 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 582 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]SALARY_HISTORY.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 126 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]RESUME_LISTS.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 30 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMPIDS_OVER.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 51 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMPIDS_LOW.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 51 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMPIDS_MID.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 51 data pages

```

(Continued)

```

%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]EMP_INFO.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 30 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]RESUMES.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 30 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]JOBS.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 27 data pages
%RMU-I-BCKTXT_02, Full backup of storage area ISG$DATA1:[RDB]DEPARTMENTS.RDA;1
%RMU-I-BCKTXT_04, ignored 1 space management page
%RMU-I-BCKTXT_05, backed up 0 inventory pages
%RMU-I-BCKTXT_06, backed up 0 logical area bitmap pages
%RMU-I-BCKTXT_07, backed up 27 data pages
$ set nover

```

Before exporting, prepare the database for the import by making some preliminary changes...

Executing the first set of ALTERS...
 *** SQL script successfully executed. ***

Now, export the existing database and then delete it so that the newly altered database can be imported...

EXPORTING the database...
 *** SQL script successfully executed. ***

DROPPING (deleting) the database...

Now, import the database with the necessary changes...

```

IMPORTING the database...
Exported by Rdb/VMS V4.2-0 Import/Export utility
A component of VAX SQL V4.2-0
Previous name was MF_PERSONNEL
It was logically exported on 11-MAY-1995 15:45
Multischema mode is DISABLED
Database NUMBER OF USERS was 50, now is 50
Database NUMBER OF VAXCLUSTER NODES was 16, now is 16
Database NUMBER OF DBR BUFFERS was 20, now is 46
Database SNAPSHOT was ENABLED
Database SNAPSHOT is ENABLED
Database SNAPSHOT was IMMEDIATE
Database SNAPSHOT is IMMEDIATE
Database JOURNAL ALLOCATION is 51
Database JOURNAL EXTENSION is 51
Database BUFFER SIZE was 6, now is 12 blocks
Database NUMBER OF BUFFERS was 20, now is 46
Adjustable lock granularity was ENABLED
Adjustable lock granularity is ENABLED
Database global buffering was DISABLED
Database global buffering is DISABLED
Journal fast commit is DISABLED
Journal fast commit checkpoint interval is 0 blocks

```

(Continued)

```

Journal fast commit checkpoint time is 0 seconds
Commit to journal optimization is Disabled

```

```

Journal fast commit TRANSACTION INTERVAL is 256
LOCK TIMEOUT is 0 seconds
Definition of STORAGE AREA RDB$SYSTEM overridden
Definition of STORAGE AREA EMPIDS_LOW overridden
Definition of STORAGE AREA EMPIDS_MID overridden
Definition of STORAGE AREA EMPIDS_OVER overridden
Definition of STORAGE AREA DEPARTMENTS overridden
Definition of STORAGE AREA SALARY_HISTORY overridden
Definition of STORAGE AREA JOBS overridden
Definition of STORAGE AREA EMP_INFO overridden
IMPORTing STORAGE AREA: RESUME_LISTS
Definition of STORAGE AREA RESUMES overridden
IMPORTing table WORK_STATUS
Definition of STORAGE MAP WORK_STATUS_MAP deleted
IMPORTing table EMPLOYEES
Definition of STORAGE MAP EMPLOYEES_MAP overridden
Definition of INDEX EMP_LAST_NAME overridden
Definition of INDEX EMP_EMPLOYEE_ID overridden
IMPORTing table JOBS
Definition of STORAGE MAP JOBS_MAP overridden
IMPORTing table DEPARTMENTS
Definition of STORAGE MAP DEPARTMENTS_MAP deleted
IMPORTing table JOB_HISTORY
Definition of STORAGE MAP JOB_HISTORY_MAP deleted
Definition of INDEX JH_EMPLOYEE_ID overridden

IMPORTing table SALARY_HISTORY
Definition of STORAGE MAP SALARY_HISTORY_MAP deleted
Definition of INDEX SH_EMPLOYEE_ID overridden
IMPORTing table COLLEGES
Definition of STORAGE MAP COLLEGES_MAP overridden
Definition of INDEX COLL_COLLEGE_CODE overridden
IMPORTing table DEGREES
Definition of STORAGE MAP DEGREES_MAP overridden
Definition of INDEX DEG_EMP_ID overridden
Definition of INDEX DEG_COLLEGE_CODE overridden
IMPORTing table CANDIDATES
Definition of STORAGE MAP CANDIDATES_MAP overridden
IMPORTing table RESUMES
Definition of STORAGE MAP RESUMES_MAP overridden
IMPORTing view CURRENT_SALARY
IMPORTing view CURRENT_JOB
IMPORTing view CURRENT_INFO
  *** SQL script successfully executed. ***

Perform any additional alters which need to be done after the import...

Executing the second set of ALTERS...
%RDMS-W-DOFULLBCK, full database backup should be done to ensure future
recovery
  *** SQL script successfully executed. ***

*****
  *** Execution of SQL scripts completed successfully. ***

Tuning process began: 11-MAY-1995 15:44:06.83
Tuning process ended: 11-MAY-1995 15:48:20.54

NOTE: Because the After-Image Journal (.AIJ)
      was re-enabled for this database, a full
      RMU/BACKUP should now be performed to
      ensure future database recovery.
*****

```

DBTune Reports

Four reports are created by the DBTune procedure and they are placed in your default directory. The four reports are:

1. <<database_id>>_REVIEW_AND_GUIDE.REPORT
2. <<database_id>>_ANALYSIS.REPORT
3. <<database_id>>_TUNING_DETAIL.REPORT
4. <<database_id>>_DBTUNE.LOG

where <<database_id>> is the first ten characters of the .RDB file name.

The **REVIEW_AND_GUIDE.REPORT** provides tuning advice, instructions for executing the transformation procedure, and any warnings about errors that occurred or requirements that must be met for the procedure to execute properly. It also includes recommendations for setting certain **VMS AUTHORIZE** or **SYSGEN** parameters that may need to be changed. It is **STRONGLY** suggested that this report be read prior to executing the transformation **MAIN_DRIVER** command file.

The **ANALYSIS** report provides a narrative of the database's tuning and complexity status prior to executing the DBTune transformation procedure.

The **TUNING_DETAIL** report lists values used or calculated during the tuning process. B-tree information is also included for storage areas containing a single sorted index.

The **DBTUNE.LOG** report is a log of the DBTune process, showing the tables and indexes read as well as the storage areas that were tuned. In addition, results or errors encountered while parsing the PAD and Disk Utilization data files can be found in this file.

Following is an example of each of these reports.

DATABASE PARAMETER CHANGES:

```
-----
- Old buffer size =      6  New buffer size =      12
- Old buffer count=    40  New buffer count=     44
- Old global buffs= DISABLED  New global buffs= DISABLED
- Old number users=    50  New number users=    50
```

Because of tuning changes, certain AUTHORIZE and SYSGEN parameters may need to be changed for users or systems which access this database.

AUTHORIZE parameter Recommended MINIMUM settings...

```
-----
FILLM          124 (cannot exceed SYSGEN param 'CHANNELCNT')
DIOLM          100
BIOLM          100
ASTLM          113
ENQLM          8000
BYTLM          31000
WSDEF *        1024
WSQUO *        1024
WSEXTENT *     3072 (cannot exceed SYSGEN param 'WSMAX')
PGFLQOTA      160000 (cannot exceed SYSTEM param 'VIRTUALPAGECNT')
```

* How to calculate WORKING SET sizes:

- ```

1) $ SHOW MEM/PHYS
2) AVAIL_MEM = 'TOTAL' - 'PERMANENTLY ALLOC TO VMS' - 3000
3) Determine the MAX number of users on your system (online, batch, etc)
4) WSQUO = AVAIL_MEM / MAX_USERS
5) WSDEF = 512, 1024 or same as WSQUO
6) WSEXT = multiple of WSQUO; can go as high as AVAIL_MEM but cannot
 exceed WSMAX (SYSGEN param)
```

## INDEX ANALYSIS:

```

Following are the results of a physical index analysis for this database
which scans for key words in the index attributes in an attempt to highlight
possible improvements in index design:
```

```
=====
* The following indices are SORTED and contain a keyword which typically
represents a unique value. Because hashed keys require an exact match and
sorted keys are better suited for range retrievals, these indices may
provide better performance as HASHED keys rather than as sorted keys...
```

| INDEX NAME        | ATTRIBUTE ID    | KEY WORD |
|-------------------|-----------------|----------|
| COLL_COLLEGE_CODE | COLLEGE_CODE    | CODE     |
| DEG_COLLEGE_CODE  | COLLEGE_CODE    | CODE     |
| DEG_EMP_ID        | EMPLOYEE_ID     | ID       |
| DEPARTMENTS_INDEX | DEPARTMENT_CODE | CODE     |
| EMP_EMPLOYEE_ID   | EMPLOYEE_ID     | ID       |
| JH_EMPLOYEE_ID    | EMPLOYEE_ID     | ID       |
| SH_EMPLOYEE_ID    | EMPLOYEE_ID     | ID       |

Out of 10 indices scanned, 7 were found whose design could potentially be improved according to the key-word analysis.

(Continued)

```

=====
Following are indices with an average number of duplicates that exceed 10.
Indices with many duplicates can seriously degrade I/O performance.
To reduce the number of duplicates for an index, additional key fields can be
added or two or more indices can be combined to form a larger index.

```

```

INDEX NAME AVERAGE DUPLICATES
=====
DEG_COLLEGE_CODE 12

```

```

*
* Following is a summary of the database changes found:
* -----
* - Number of INDICES altered : 10
* - Number of STORAGE MAPS altered : 10
* - Number of STORAGE MAPS dropped : 4
* - Number of STORAGE AREAS added : 10
* - Number of STORAGE AREAS altered: 2
* - Number of STORAGE AREAS dropped: 7
* - Number of DATABASE PARAMS specified: 11
*
* TOTAL DATABASE CHANGES FOUND 54
*

```

DATABASE TUNING SUMMARY:

\*\*\*\*\* Disk Space Requirements \*\*\*\*\*

```

o Approx size of tuned database : 5240 blocks
o Approx size of the BACKUP file : 1309 blocks
o Approx size of the SQL EXPORT file : 203 blocks

```

Before executing the MAIN\_DRIVER procedure, ensure that adequate free space exists on the disks to which the database and its associated files are assigned:

```

o RDB root area : ISG$DATA1:[RDB]
o RMU/BACKUP area: ISG$DATA1:[GISBCK]
o SQL EXPORT area: ISG$DATA1:[INSTALL]

```

\*\*\*\*\*

*(Continued)*

TO EXECUTE THE GENERATED SQL SCRIPTS:

-----

All of the files listed below are located in the SQL directory: ISG\$DATA1:[INSTALL]

The scripts that are created are driven by... MF\_PERSONN\_TRANSFORM.MAIN\_DRIVER

It executes the following command files:
RMU/BACKUP of the database
MF\_PERSONN.ALTERS1\_SQL
MF\_PERSONN.EXPORT\_SQL
DROP the database
MF\_PERSONN.IMPORT\_SQL
MF\_PERSONN.ALTERS2\_SQL

During execution of the MAIN\_DRIVER command, a log file will be created in the SQL directory and will be named MF\_PERSONN.SQL\_LOG. This file can be viewed for SQL messages or errors that may have occurred during execution of the scripts and can be deleted after viewing.

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

IMPORTANT:
For the database transformation procedure to succeed, any VMS logicals which are required to invoke this database must be assigned PRIOR to the transformation!

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

AUTOSQL: Automatic SQL Creation process ended successfully at 05/11/97 13:28:25:09

## Analysis Report

```

*
* FILE : MF_PERSONN_ANALYSIS.REPORT
* CREATED : 05/11/97 10:10
*
* Narrative analysis of current database characteristics
*

```

Database Complexity  
-----

The complexity rating is a weighted measure of the database design and its stored records. A rating of 8 indicates a relatively small database. Tuning requirements are simple. The largest factor in this rating is the domain count. It accounts for 13% of the complexity rating. The next largest component of this rating is the column count, which is also 13%. The complexity will increase as records and Rdb items (e.g., tables, columns) are added.

Database Tune Rating  
-----

The tune rating is a composite measure of the physical storage design as it applies to the logical structure of this database. The complexity rating of 8 and the tune rating of 37 indicate that the physical storage strategy should allow this database to perform reasonably well and allow for some growth in complexity without noticeable performance degradation. The tune rating measures significant factors that affect the physical storage design. It does not consider every factor, but does objectively measure factors critical to successful Rdb tuning. Remember that overall performance is a factor of many things, including system load, system tuning, and application design, in addition to the physical storage strategy.

Database Storage Area Allocation  
-----

The 'Storage Area Allocation' data indicate the percentage of allocated space that has been extended for both RDA and SNP files. Of the total RDA pages, a large percentage (60%) have been extended as these storage areas have been loaded. RDA pages are extended when the data requirements for tables and/or indices exceed the existing allocation of pages. The RDA areas have been extended 0 times.

Of the total SNP pages, a large percentage (55%) have been created as these storage areas have been utilized. SNP files are used to enable READ-only transactions to access data concurrently while WRITE transactions are active. SNP pages are only used if SNAPSHOTS ARE ENABLED.

*(Continued)*

## Database Index Analysis

-----

The index analysis data indicate that 20% of your indices are HASHED. Thus, 80% of your indices are SORTED. DBTune reviewed your indices and found that five of the SORTED indices are candidates to be HASHED and none of the HASHED indices are candidates to be SORTED. Review the REVIEW and GUIDE report to see which indices have been selected.

DBTune looks for certain key words within the index columns. It assumes certain types of queries will be made based on these key words. The person responsible for maintaining the database should review actual usage to determine whether to modify the index.

NOTE: HASHED indices facilitate exact match queries.  
They incur narrow locks for updates.  
SORTED indices facilitate sequential retrievals.  
They incur broader lock contention for updates  
than HASHED indices.

## Tuning Detail Report

```

*
* FILE : MF_PERSONN_TUNING_DETAIL.REPORT
* CREATED : 05/11/97 13:25
*
* Detail of parameters and calculated values resulting from tuning
* the storage areas for MF_PERSONNEL.
* For tuning summaries and advice, see REVIEW & GUIDE report.
*

*
NOTE: For those storage areas containing a single SORTED index, B-tree
information will be printed (tree levels, duplicate nodes, etc).
Be aware that each B-tree represents a non-compressed index and
is based on the record count, index fill percentage, growth rate,
access bias, and node size listed for the particular index involved.
If compression is specified for an index, its actual B-tree may be
smaller than represented below.

Values used to calculate page size for: DEGREES_TBL

STORED ITEMS : 1 table
TABLE BYTE COUNT/SEGS : 29 / 5
OF RECS (10% growth): 182
TABLE RECORDS PER PAGE : 41
ACCESS BIAS : R
PAGE SIZE/ALLOCATION : 4 / 16
TOTAL RDA BLOCKS : 64
TOTAL SNP BLOCKS : 24

Values used to calculate page size for: DEG_COLLEG_IDX

* B-TREE CONSISTS OF ONLY ONE NODE

STORED ITEMS : 1 sorted index
SORT KEY SIZE/SEGMENTS : 4 / 1
OF RECS (10% growth): 190
NODE SIZE (bias : R) : 432
INDEX RECORDS PER NODE : 25 (22 with 90% fill)
NODES PER PAGE : 3
AVERAGE DUPLICATES ... : 12
PAGE SIZE/ALLOCATION : 3 / 16
TOTAL RDA BLOCKS : 48
TOTAL SNP BLOCKS : 18

```

*(Continued)*



```

Values used to calculate page size for: DEG_EMP_ID_IDX

* B-TREE LEVELS FOR INDEX: DEG_EMP_ID ...
 - LEVEL 1 : 9 nodes (BOTTOM of the tree)
 - LEVEL 2 : 1 node (TOP of the tree)
 =====
 TOTAL B-TREE NODES : 10 nodes
 LEVELS IN THE TREE : 2 levels

 STORED ITEMS : 1 sorted index
 SORT KEY SIZE/SEGMENTS : 5 / 1
 # OF RECS (10% growth): 190
 NODE SIZE (bias : R) : 457
 INDEX RECORDS PER NODE : 25 (22 with 90% fill)
 NODES PER PAGE : 3
 AVERAGE DUPLICATES ... : 1
 PAGE SIZE/ALLOCATION : 3 / 14
 TOTAL RDA BLOCKS : 42
 TOTAL SNP BLOCKS : 18

Values used to calculate page size for: EMPIDS_OVER

 N-CLUSTER ITEMS : 2 table(s)
 : 2 hashed index(es)
 LARGEST RECORD + OH : 140
 SMALLEST RECORD + OH : 40
 'GROUP' RECORD + OH : 361
 SPAM THRESHOLDS : 17, 67, 90
 PAGE SIZE/ALLOCATION : 1 / 183
 TOTAL RDA BLOCKS : 183
 TOTAL SNP BLOCKS : 10

Values used to calculate page size for: EMP_EMPLOY_IDX

* B-TREE LEVELS FOR INDEX: EMP_EMPLOYEE_ID ...
 - LEVEL 1 : 6 nodes (BOTTOM of the tree)
 - LEVEL 2 : 1 node (TOP of the tree)
 =====
 TOTAL B-TREE NODES : 7 nodes
 LEVELS IN THE TREE : 2 levels

 STORED ITEMS : 1 sorted index
 SORT KEY SIZE/SEGMENTS : 5 / 1
 # OF RECS (10% growth): 115
 NODE SIZE (bias : R) : 457
 INDEX RECORDS PER NODE : 25 (22 with 90% fill)
 NODES PER PAGE : 3
 AVERAGE DUPLICATES ... : 0
 PAGE SIZE/ALLOCATION : 3 / 13
 TOTAL RDA BLOCKS : 39
 TOTAL SNP BLOCKS : 18

```

(Continued)

Values used to calculate page size for: EMP\_LAST\_N\_IDX

```

* B-TREE LEVELS FOR INDEX: EMP_LAST_NAME ...
- LEVEL 1 : 8 nodes (BOTTOM of the tree)
- LEVEL 2 : 1 node (TOP of the tree)
=====
TOTAL B-TREE NODES : 9 nodes
LEVELS IN THE TREE : 2 levels

STORED ITEMS : 1 sorted index
SORT KEY SIZE/SEGMENTS : 14 / 1
OF RECS (10% growth): 115
NODE SIZE (bias : R) : 474
INDEX RECORDS PER NODE : 17 (15 with 90% fill)
NODES PER PAGE : 3
AVERAGE DUPLICATES ... : 1
PAGE SIZE/ALLOCATION : 3 / 14
TOTAL RDA BLOCKS : 42
TOTAL SNP BLOCKS : 18
```

Values used to calculate page size for: JH\_EMPLOYEE\_IDX

```

* B-TREE LEVELS FOR INDEX: JH_EMPLOYEE_ID ...
- LEVEL 1 : 8 nodes (BOTTOM of the tree)
- LEVEL 2 : 1 node (TOP of the tree)
=====
TOTAL B-TREE NODES : 9 nodes
LEVELS IN THE TREE : 2 levels
LARGE DUPLICATE NODES : 16 nodes
SMALL DUPLICATE NODES : 141 nodes

STORED ITEMS : 1 sorted index
SORT KEY SIZE/SEGMENTS : 5 / 1
OF RECS (10% growth): 315
NODE SIZE (bias : R) : 457
INDEX RECORDS PER NODE : 25 (22 with 90% fill)
NODES PER PAGE : 3
AVERAGE DUPLICATES ... : 2
PAGE SIZE/ALLOCATION : 3 / 32
TOTAL RDA BLOCKS : 96
TOTAL SNP BLOCKS : 18
```

Values used to calculate page size for: RDB\$SYSTEM

```

STORED ITEMS : 0 table(s)
: 0 sorted index(es)
: 0 hashed index(es)
PAGE SIZE/ALLOCATION : 4 / 763
TOTAL RDA BLOCKS : 3052
TOTAL SNP BLOCKS : 156
```

Values used to calculate page size for: SALARY\_HIS\_TBL

```

STORED ITEMS : 1 table
TABLE BYTE COUNT/SEGS : 25 / 4
OF RECS (10% growth): 802
TABLE RECORDS PER PAGE : 45
ACCESS BIAS : R
PAGE SIZE/ALLOCATION : 4 / 30
TOTAL RDA BLOCKS : 120
TOTAL SNP BLOCKS : 24
```

(Continued)

Values used to calculate page size for: SH\_EMPLOYEE\_IDX

```

* B-TREE LEVELS FOR INDEX: SH_EMPLOYEE_ID ...
- LEVEL 1 : 6 nodes (BOTTOM of the tree)
- LEVEL 2 : 1 node (TOP of the tree)
=====
TOTAL B-TREE NODES : 7 nodes
LEVELS IN THE TREE : 2 levels
LARGE DUPLICATE NODES : 42 nodes
SMALL DUPLICATE NODES : 77 nodes

STORED ITEMS : 1 sorted index
SORT KEY SIZE/SEGMENTS : 5 / 1
OF RECS (10% growth): 838
NODE SIZE (bias : R) : 457
INDEX RECORDS PER NODE : 25 (22 with 90% fill)
NODES PER PAGE : 3
AVERAGE DUPLICATES ... : 7
PAGE SIZE/ALLOCATION : 3 / 34
TOTAL RDA BLOCKS : 102
TOTAL SNP BLOCKS : 18
```

Values used to calculate page size for: SMALL\_SORTED\_AREA

```

STORED ITEMS : 0 table(s)
 : 2 sorted index(es)
 : 0 hashed index(es)
'GROUP' RECORD + OH : 1326
PAGE SIZE/ALLOCATION : 4 / 20
TOTAL RDA BLOCKS : 80
TOTAL SNP BLOCKS : 24
```

Values used to calculate page size for: SMALL\_TABLE\_AREA

```

STORED ITEMS : 6 table(s)
 : 0 sorted index(es)
 : 0 hashed index(es)
'GROUP' RECORD + OH : 301
PAGE SIZE/ALLOCATION : 4 / 65
TOTAL RDA BLOCKS : 260
TOTAL SNP BLOCKS : 24
```

## DBTune Process Log Report

```

*
* FILE : MF_PERSONN_DBTUNE.LOG
* CREATED : 05/11/97 10:10
*
* Log of the parameter settings and execution of DBTune procedure for...
*
* Database: MF_PERSONNEL
*

*** Total BUFFER-PAGE records successfully parsed from FRENDS$BUFFER$PAGE: 29
*** Total BIAS-SCALE records successfully parsed from FRENDS$BIAS$SCALE : 101

Following are the parameter settings used for this execution of DBTune:

Strategy = N -->> (Create all NEW storage areas)
TuneTechnique= SQL -->> (Use SQL Export/Import to tune)
DBDisks = 1 -->> (See next 2 lines)
DBDISK01 = ISG$DATA1:[RDB]
 (DBDISK01 file types: /TBLRDA/TBLSNP/IDX RDA/IDXSNP/SYSRDA/SYSSNP/SYSRDB/)
Edit Files = Y -->> (Edit PAD & DiskUtil during process)
FrEnd Editor = EDIT/EDT
ModPAD file = -->> (No file specified)
DynWork File = -->> (No file specified)
Access Bias = R -->> (READ biased)
Index Fill % = 90% -->> (fill for sorted index nodes)
Growth % = 10% -->> (estimated database growth)
Snapshot % = 5% -->> (% of RDA allocated for SNP)
Min Page Size= 1 -->> (Min storage area page size in blocks)
Max Page Size= 32 -->> (Max storage area page size in blocks)
Min Buff Size= 6 -->> (Min database buffer size in blocks)
Max Buff Size= 63 -->> (Max database buffer size in blocks)
Min Buffers = 20 -->> (Minimum # of database buffers)
Max Buffers = 100 -->> (Maximum # of database buffers)
Sys Mem Pages= 0 -->> (use existing db global buffer settings)
Max DB Users = 0 -->> (use existing db users setting)
Spread Areas = B -->> (Based on BOTH Volume & Activity)
Logicals = N -->> (No logs; use physical locations)
Logical Type = PROCESS -->> (No logs; use physical locations)
Conceal Logs = N -->> (No logs; use physical locations)
Load Time Lim= 0 -->> (has no effect for SQL Export/Import)
Machine VUPs = 2.5 -->> (has no effect for SQL Export/Import)
Table Commit = Y -->> (has no effect for SQL Export/Import)
Save Comments= Y -->> (has no effect for SQL Export/Import)
Tune for Comp= N -->> (Use UNCOMPRESSED data values for tuning)
Min Card = 100 -->> (See next section below)
-- Storage Areas for Items with a Cardinality Below: 100 ('Min Card')
Tables = SMALL_TABLE_AREA
Sorted = SMALL_SORTED_AREA
Hashed = SMALL_HASHED_AREA
----- Assigned Directories -----
SQL Dir = ISG$DATA1:[INSTALL]
RUJ Dir = ISG$DATA1:[TEST]
BackupDir = ISG$DATA1:[GISBCK]
ExpUnlDir = ISG$DATA1:[INSTALL]

```

(Continued)

```

----- Assigned Logicals -----
FRIENDRDBIMPORT = ISG$DATA1:[RDB]MF_PERSONNEL.RDB
FRIEND$DBTUNE$PARAMS = ISG$DATA3:[DBTUNE]DBTUNE_DEFAULT.PARAMS
FRIEND$DBTUNE$HOME = ISG$DATA3:[DBTUNE]
FRIEND$DBTUNE$SCRATCH = ISG$DATA3:[DBTUNE].SCRATCH]
FRIEND$BUFFER$PAGE = FRIEND$DBTUNE$HOME:BUFFER_PAGE.DAT
FRIEND$BIAS$SCALE = FRIEND$DBTUNE$HOME:BIAS_SCALE.DAT

DBTUNE: Initializing Metadata for DBTune...

DBTUNE: Reading Rdb Tables...

 Reading table: CANDIDATES
 Reading table: COLLEGES
 Reading table: DEGREES
 Reading table: DEPARTMENTS
 Reading table: EMPLOYEES
 Reading table: JOBS
 Reading table: JOB_HISTORY
 Reading table: RESUMES
 Reading table: SALARY_HISTORY
 Reading table: WORK_STATUS

DBTUNE: Reading Rdb Indices...

 Reading index: COLL_COLLEGE_CODE
 Reading index: DEG_COLLEGE_CODE
 Reading index: DEG_EMP_ID
 Reading index: DEPARTMENTS_INDEX
 Reading index: EMPLOYEES_HASH
 Reading index: EMP_EMPLOYEE_ID
 Reading index: EMP_LAST_NAME
 Reading index: JH_EMPLOYEE_ID
 Reading index: JOB_HISTORY_HASH
 Reading index: SH_EMPLOYEE_ID

DBTUNE: Searching for N-clustered items...

*** NOTE: The EMPLOYEES table is a member of
 the N-cluster area: EMPIDS_OVER
*** NOTE: The EMPLOYEES_HASH index is a member of
 the N-cluster area: EMPIDS_OVER
*** NOTE: The JOB_HISTORY table is a member of
 the N-cluster area: EMPIDS_OVER
*** NOTE: The JOB_HISTORY_HASH index is a member of
 the N-cluster area: EMPIDS_OVER

PAD WARNING:
- Free Blocks specification for DISK01 (104971)
 exceeds the actual number of free blocks for the disk: 99210.

DBTUNE: Gathering information for the tuning process...

```

*(Continued)*

```
DBTUNE: Tuning database storage areas...

Tuning storage area: ACCTNG_SMALL_SORTED_AREA
Tuning storage area: ACCTNG_SMALL_TABLE_AREA
Tuning storage area: COLL_COLLE_IDX
Tuning storage area: DEGREES_TBL
Tuning storage area: DEG_COLLEG_IDX
Tuning storage area: DEG_EMP_ID_IDX
Tuning storage area: EMPIDS_OVER
Tuning storage area: EMP_EMPLOY_IDX
Tuning storage area: EMP_LAST_N_IDX
Tuning storage area: JH_EMPLOYE_IDX
Tuning storage area: RDB$SYSTEM
Tuning storage area: SALARY_HIS_TBL
Tuning storage area: SH_EMPLOYE_IDX

DBTUNE: Tuning process SUCCESSFUL.

DBTUNE: Spreading storage areas over available disks

DBTUNE: Creating Disk Utilization data file

DBTUNE: Interpreting Disk Utilization data file

DBTUNE: Determining file specifications for storage areas

DBTUNE: Scanning database for modifications...
DBTUNE: - Storage map changes
DBTUNE: - Index changes (node size, type)
DBTUNE: - Storage area changes
DBTUNE: - Sequencing database changes
DBTUNE: Modification scan SUCCESSFUL.

DBTUNE: Generating transformation procedure...
DBTUNE: - Generating transformation MAIN_DRIVER
DBTUNE: - Generating EXPORT/IMPORT commands
DBTUNE: - Generating SQL optimization scripts
DBTUNE: Generation of transformation procedure SUCCESSFUL.
```

## DBTune Help

Online HELP is available within DBTune by pressing the **HELP** key or by selecting the **HELP** option in the DBTune menu and pressing **Return**. To obtain DBTune help outside of the DBTune utility, type the following command at the **DCL** prompt after installation of DBTune:

```
$ HELP/LIBRARY=FREND$DBTUNE$HOME:DBTUNE.HLB
```

Following is an example of the DBTune HELP window:

```

*** DBTune V5.2 ***
Rdb Name: 1$DUA11:[YOUNGVC.TESTDB]MF_PERSONNEL.RDB 09/05/98 17:10
HELP
You are currently executing DBTune V5.2, a read-only utility that
analyzes a database and produces SQL scripts that a user can later
execute to tune the database. For a list of available topics, type
"?" at the "Topic?" prompt or type in the specific item for which help
is required and press [RETURN]. To EXIT Help at any time, press [ESC]
or press [RETURN] at the prompt.

Press RETURN to continue..._

ModPAD file=
SQL Dir = ISG$DATA0:[DBTRDBVMS 52]
Backup Dir = ISG$DATA0:[DBTRDBVMS 52]
Exp/Unl Dir= ISG$DATA0:[DBTRDBVMS 52]

Tune Rating: 70

[DO]-Create Tuning Scripts [SELECT]-Edit Parameters [HELP]-Help
[ESC]-Exit

```





# Chapter 3

---

## DBXAct for Rdb

### What is DBXAct?

---

**D**BXAct for Rdb is a monitoring performance tool designed to generate baseline statistics and reports on Rdb activity. DBXAct for Rdb allows you to clearly understand the hundreds of data points available for monitoring database activity. The detailed statistics and reports generated by DBXAct will allow you to perform precise tuning and optimization of your Rdb database.

### Rdb Monitoring and Performance Tuning with DBXAct

Understanding Rdb activity is an important step in tuning your Rdb database. What type of activity, how much activity, and where the activity is occurring is vital information you need when optimizing the Rdb database and the system on which it runs.

Establishing benchmarks to accurately reflect the activity of your database before and after changes are made is critical in order to understand the effects of the changes made. Tuning your Rdb database without measurements is like working in the dark. You must be able to see the results of your changes in order to determine if your tuning efforts have been successful.

Tuning has been described as a balancing process. Resources available for a system are finite. Finding the optimal balance between available resources such as CPU, memory, and disk I/O is the key to successful Rdb database tuning and optimization. DBXAct for Rdb will accurately provide you with the information you need to successfully tune and optimize your Rdb database.

## DBXAct Features

**BaseLine Statistics** A customizable report provides baseline measurements of activity.

**Automatic Feed to DBTune by Tables, Indexes, and Clusters** Information on total activity, I/O bias (read/write), and growth are available.

**Monitoring** Reporting of critical statistics for individual databases.

- DBXAct 5.0/5.1 supports Rdb versions through 7.0-x.
- DBXAct 5.2 supports Rdb versions through 7.1.x
- Memory leaking has been eliminated, allowing DBXAct to execute for a long duration without resource conflicts.
- DBXAct collects extensive information on the Rdb database tables, indexes, clusters, Rdb version, average memory usage, buffer ratio, and session information.
- Multiple report files are routed to one output file. This feature allows you to monitor several periods and project growth based on long-term monitoring. A useful column, **Standard Deviation**, is a newer feature of this file.
- Average value column in the **GIS.REPORT** file is based on the average value per second instead of per scan.
- A DBXAct batch parameter selection is a recently added feature of DBXAct. A parameter is available to pass active hours per day for projection instead of 24 hours a day.

- You can shut down a specified DBXAct executing. This allows the remaining DBXAct executions to continue operating.
- Logical Symbols have now been corrected in **DBXACT.COM**.
- A registration ID code to compile and link is incorporated in the product.

## DBXAct Overview and Concepts

DBXAct can monitor any active Rdb database on a system. An “active” database is one that either was opened manually via **RMU/OPEN** or that currently has users attached.

When executed, DBXAct generates a customizable summary report, highlighting critical statistics. Both summary and detailed information are available for the database.

A database activity file is also produced. This activity file can later be used by DBTune (a tuning tool) to evaluate the tuning needs of the database and create SQL scripts for its physical redesign. The activity file recommends a growth percentage based on the changes in record counts of the tables in the database, factored over a number of days that you specify. In addition to the growth percentage, tables are ranked according to the number and type of I/Os performed against them. A **READ** or **WRITE** bias affects the tuning calculations performed by DBTune. An example of one of these DBTune activity files can be found on the following page. The name of the DBTune file is always:

**<database name>.DBTUNE**

For example, if you have an Rdb database called MYDB, the activity file produced from a DBXAct run would be called:

**MYDB.DBTUNE**

## Example of Database Activity

```

!!
!
! Filename: MYDB.DBTUNE
!
!
! DBTune Activity Analysis File for Database:
! 1dka0:[rdb]mydb.rdb
!
! File generated at 10/12/94 12:52:40
!
! Tuning Interval is 30 days.
!
!*SOURCE=DBXACT
!*BEGAN = 10/12/94 10:50:55
!*ENDED = 10/12/94 12:50:57
!
! Process monitored for 2 seconds.
!
!*TOTAL_MEM = 131072
!*AVAIL_MEM = 17636
!
! Listed below are values that can be incorporated in tuning:
! The user may change values for those columns underlined with
! "--*--". Listed below are reserved characters used by
! DBTune when interpreting this file:
! "!" : Comment line (if ! is the 1st char, line is ignored).
! "/" : Column separator, MUST exist between column values.
! "{}" : The enclosed item is a database TABLE.
! "{}" : The enclosed item is a database INDEX.
! "<>" : The enclosed item is a database CLUSTER.
!!
!
! Activity %Growth BIAS
! [TABLE] or {INDEX} or <CLUSTER> (1-9) (0-999) (1..100)
! -----*-----*-----*-----
[ACCOUNTS_PAYABLE] / 2 / 19 / 100
[CALL_LOG] / 5 / 7 / 100
[CALL_STATUS_CODES] / 3 / 0 / 44
[CLIENTS] / 1 / 2 / 100
[CLIENT_PROD] / 1 / 340 / 100
[CONTRACTS] / 5 / 33 / 21
[CONTRACT_ITEMS] / 5 / 0 / 100
[DBTUNE_USERS] / 7 / 0 / 26
[LEADS] / 1 / 83 / 100
[LICENSE] / 5 / 0 / 100
[PRODUCT_ID] / 1 / 1 / 86
[PRODUCT_VERS] / 9 / 0 / 100
[PURCHASE_ORDER] / 2 / 0 / 100
[RESELLER_CONTACTS] / 1 / 73 / 0
[SALESMAN] / 6 / 0 / 100
[TAPE_MANAGEMENT] / 5 / 0 / 1
[TIME_RECORDS] / 5 / 0 / 67
{ CALL_STATUS_CD_IDX } / 1 /
{ CLIENT_NAME } / 5 /
{ LASTNAME } / 5 /
{ LEAD_ID } / 1 /
{ MEDIA_CODES } / 2 /
{ PRODUCT_CLIENT_IDX } / 9 /
{ QUERY_ID } / 1 /
!!

```

## Getting Started

---

This section provides the information you need to quickly install and run DBXAct for Rdb. If you are familiar with installing and running DBXAct for Rdb, you may prefer to go directly to the **Quick Start** section below. Less experienced users should carefully read the remaining sections in this chapter, which provide detailed system requirements and installation instructions.

### Quick Start

► **To install and run DBXAct for Rdb:**

1. Install DBXAct for Rdb. (If necessary, see the following sections on system requirements and installation.)
2. Go to the section on page 153, titled **Running DBXAct**.

## System Requirements

---

| Item               | Requirement                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operating System   | Open VMS 6.1 or higher. For sites where the version of VMS is lower than 6.1, possible solutions may be provided. Please contact ALI's technical support for assistance. Phone: (866) 257-8970 |
| Disk Space         | 45,000 blocks of disk space for VAX/VMS and approximately 60,000 blocks for Alpha AXP                                                                                                          |
| Authorize Settings | See the following <b>AUTHORIZE</b> description.                                                                                                                                                |

Recommended **minimum AUTHORIZE** settings for a DBXAct user account. (Medium and large databases may need to increase these numbers.)

|                        |                             |              |                 |
|------------------------|-----------------------------|--------------|-----------------|
| Username:              | USER                        | Owner:       | USER            |
| Account:               | USER                        | UIC:         | [Group, Member] |
| CLI:                   | DCL                         | Tables:      | DCLTABLES       |
| Default:               | <disk>:[dir]                |              |                 |
| LGIMD:                 | LOGIN                       |              |                 |
| Login Flags:           |                             |              |                 |
| Primary Days:          | Mon. Tues. Wed. Thurs. Fri. |              |                 |
| Secondary Days:        | Sat. Sun.                   |              |                 |
| No access restrictions |                             |              |                 |
| Expiration:            | (none)                      | Pwdminimum:  | 0               |
| Pwdlifetime:           | (none)                      | Pwdchange:   |                 |
| Last Login:            |                             | Login Fails: | 0               |
| Maxjobs:               | 0                           | Fillm:       | 512             |
| Maxacctjobs:           | 0                           | Shrfillm:    | 0               |
| Maxdetach:             | 0                           | BIOfm:       | 100             |
| Prclm:                 | 4                           | DIOIm:       | 100             |
| Prio:                  | 4                           | ASTIm:       | 113             |
| Queprio:               | 0                           | TQEIm:       | 10              |
| CPU:                   | (none)                      | Enqlm:       | 8000            |
| Authorized Privileges: | GROUP TMPMBX<br>NETMBX      | Bytlm:       | 10000           |
| Default Privileges:    | GROUP MPMBX<br>NETMBX       | Pbytlm:      | 0               |
|                        |                             | Jtquota:     | 1024            |
|                        |                             | Wsdef:       | 1024            |
|                        |                             | Wsquo:       | 5120            |
|                        |                             | Wsexent:     | 5120            |
|                        |                             | Pqlfquo:     | 80000           |

**Warning** If these minimums are not in place when DBXAct is executed, the analysis may fail!

## Installing DBXAct for Rdb

**Caution** If the minimum requirements listed in the System Requirements section are not available, DBXAct may fail when executed.

► **To install DBXAct for Rdb:**

1. Back up your system disk (optional).
2. Log in under the **SYSTEM** account or an account that has the VMS privilege **SYSPRV**.
3. Place the DBXAct distribution tape in the tape drive.
4. Type the following command to invoke the VMS install facility to install DBXAct on your system:

**For VAX/VMS**

```
$ @SYS$UPDATE:VMSINSTAL DBXRDBVMS050 <tape-drive>:
```

**For Alpha AXP:**

```
$ @SYS$UPDATE:VMSINSTAL DBXRDBAXP050 <tape-drive>:
```

where <tape-drive> is the name of the device where the DBXAct distribution tape has been mounted (such as MUA6:).

**Caution** DBXAct V5.0/5.1/5.2 should not be installed in the same directory with any other product from ALI (such as DBTune).

5. After the VMS installation has completed, place the following line into the system startup command file

```
(SYS$MANAGER:SYSTARTUP_VMS.COM)
```

so that the required logical is set up when the system is rebooted:

```
$ DEFINE/SYSTEM/EXEC FRENDDDBXHOME <disk>:[dir]
```

where <disk> and [dir] are the names of the disk and directory to which DBXAct was installed, such as \$1\$DUA1:[DBXRDBVMS50] or \$1\$DUA1:[DBXRDBAXP50].

6. Edit the **SYS\$MANAGER:SYLOGIN.COM** file and add the following symbol:

```
$ DBXACT ::= @FRENDDDBXHOME:DBXACT.COM
```

7. To obtain a license pak for DBXAct, type the following commands:

```
$ SET DEFAULT FRENDDDBXHOME
```

```
$ EDIT DBXACT.LICENSE
```

8. For each node (“machine”) on which you wish to run DBXAct, do the following:

- Replace “your node name” with the node name of the machine on which you have installed DBXAct. To get this information, type:

```
$ WRITE SYS$OUTPUT F$GETSYI (“nodename”)
```

- If the “operating system” value supplied with your license is not accurate for your system, replace it with the output generated from the following command:

```
$ WRITE SYS$OUTPUT F$GETSYI (“node_swtype”)
```

- Replace “your company name” with your company’s full name.
  - Exit and save the file.
9. To obtain the appropriate REGISTRATION ID for each machine entered, call ALI at (866) 257-8970 [or (803) 648-5931] or fax a copy of the altered **DBXACT.LICENSE** file to (803) 641-0345. Be sure to indicate the version of the product you are using.



**Note** International clients may obtain registration IDs or support through their local distributor's office.

## Running DBXAct

---

This section describes how to start and run DBXAct, and explains the significance of all DBXAct variables.

### Starting DBXAct

To invoke DBXAct, type the following:

```
$@FREND$DBX$HOME:DBXACT
```

A DCL symbol can be created to make this easier:

```
DBXACT ::= "@FRENDDBXHOME:DBXACT.COM"
```

This allows DBXAct to be executed by the following command:

```
$DBXACT
```

Check with your system manager to see if a symbol was created or create one in your own **LOGIN.COM**.

If you have several different versions of Rdb currently running on your system, you will now see the following paragraph:

```

WELCOME TO DBXACT VERSION 5.0 FOR RDB!

The logical FRENDSDBX$DATABASE is assigned to the database:
 <Value assigned to this logical.>
This will be the database monitored.
If the version of the database is not: <version #>, then enter
the proper version now. Otherwise hit <RETURN>.
Rdb Version:
```

If you only have one version running, you will see this:

```

WELCOME TO DBXACT VERSION 5.0 FOR RDB!

The logical FRENDSDBX$DATABASE is assigned to the database:
 <value assigned to this logical>
What version of Rdb is the database:
 <database name?>
```

After receiving the answers to these questions, the command file displays a paragraph describing the default parameter settings for the procedure:

```
Current DBXAct
Parameter Settings
DBXAct will record statistics for 8 hours.
While executing, it will scan for statistics every 10 seconds.
DBTune activity files will project growth for 30 days.
DBXAct will execute in batch on queue SYS$BATCH.
Do you wish to make changes to any of these settings (Y/N) [Y]:
```

Entering a **CTRL-Z** at this prompt lets you exit the procedure without running DBXAct. By entering a **Y** the procedure will bring up a menu of parameter options that you can change:

DBXAct Batch  
Parameter Selection

1. Monitoring Duration
2. Monitoring Scan Rate
3. Growth Projection Interval
4. Batch Queue

Enter the number you wish to change [0]:

Entering **CTRL-Z** or a zero at this prompt returns you to the parameter paragraph. By selecting a number from one to four, the procedure will prompt you for changes to be made to parameters that govern DBXAct's execution. Following is a sample of the prompts that you will see for each option. Please note that entering a simple **Return** will always result in the use of the current value.

1. Monitoring Duration  
DBXAct will monitor database activity for 8 hours.  
How many hours would you like to monitor activity (0,720) [8]:
2. Monitoring Scan Rate  
Current scan rate is 10 seconds.  
The valid range is 5 to 3600.  
Enter new scan rate (5,3600) [10]:
3. Growth Projection Interval  
Current growth projection interval is for 30 days.  
The valid range is 1 to 365.  
Enter new growth interval (1,365) [30]:
4. Batch Queue  
DBXAct will be submitted to execute in queue SYS\$BATCH.  
Enter the name of the queue, or 'HELP' for a list of queues.

Any batch run of DBXAct will produce a DBTune dynamic activity file called **<database\_name>.DBTUNE** and a report file called **<database name>.REPORT** in **SYS\$LOGIN**. Two log files, **DBX\_BATCH.LOG** and **DBXACT.LOG**, will also be created in **SYS\$LOGIN**. If you suspect something may have gone wrong during the DBXAct process, be sure to check the log file in this directory.

---

## Explanation of DBXAct Variables

---

**Monitoring Duration** The default duration period for DBXAct monitoring is eight hours. This can be changed to any desired amount of time, given that the specified time is greater than or equal to one hour and that only full hour variables are given. (Two and a half hours is not allowed, but two hours or three hours is allowed.)

**Monitoring Scan Rate** This value specifies the interval at which to scan. The default value is to scan every ten seconds, but this may be changed to every thirty seconds or every five seconds. By specifying a value of x seconds, DBXAct will scan the database, then “sleep” for x seconds, then scan the database, then “sleep” for x seconds, and so on. A value of zero is not allowed.

**Growth Projection Interval** The default value for the growth projection interval is set to thirty days. This means that the DBTune file will project the results of the monitoring session over thirty days, rather than simply the monitoring duration period. For example, if DBXAct is set to monitor a database for 24 hours and a table grows by 1 percent during that time period, then the 1 percent growth would be projected over a 30-day time period to result in 30 percent growth for that table. Use this parameter to indicate how many days will elapse between database tuning sessions so that appropriate growth can be planned for and used to anticipate database storage needs.

**Batch Queue** This is the name of the batch queue where DBXAct will run. The default value is **SYS\$BATCH**, but can be changed to any logical name for a batch queue.

## Using DBXAct

---

**Note** DBXAct can only be run in batch mode.

The DBXAct process collects database statistics over a specified period of time, and later these statistics are reported in DBXAct's customizable summary report.

The performance of DBXAct can be controlled with a number of runtime logicals or variables. The section about DBXAct Logicals beginning on page 168 lists the logical names that affect DBXAct.

**Important** The account used to run DBXAct must have sufficient privileges to perform an **RMU/SHOW STATISTICS** for the database and to select data from **SYSTEM** relations. The VMS privilege **SYSPRV** will be sufficient or a combination of database privileges.

The current version of DBXAct supports Rdb 4.2-13 through Rdb 7.0-x.

---

## Analyzing Data Gathered with DBXAct

---

This section provides information on the various database statistics gathered by DBXAct for Rdb. These statistics will provide you with important information that will allow you to accurately tune and optimize your Rdb database.

### Storage Area Statistics

Storage area statistics are always reported in the DBXAct customizable report. These are used to determine which storage areas are more active than others, and where any trouble areas may arise. The following statistics for each storage area are reported on a per I/O basis and can be used for comparison purposes:

- Read IOs
- Write IOs
- Extend IOs
- Read Blocks
- Write Blocks
- Extend Blocks
- Read Stall Time
- Write Stall Time
- Extend Stall Time

### General Database Statistics

- Transactions
- Verb Successes
- Verb Failures
- in GB, Rt Ver, Need Lck

### Stall Statistics

- Data Read Time
- Data Write Time

## Locking Summary Statistics

Locks Requested  
Locks Unlocked  
Locks Promoted  
Locks Demoted  
Requests Stalled  
Request Deadlocks  
Proms Stalled  
Prom Deadlocks  
Stall Time

## PIO Data Fetch Statistics

**PIO Fetches for IO:** Sum of PIO fetches for reads and PIO fetches for writes

**PIO Spam Fetches for IO:** Sum of PIO Spam fetches for reads and PIO Spam fetches for writes

## Record Statistics

Records Fetched  
Records Fragmented  
Records Stored  
Pages Checked

## Index Statistics

**B-tree Scans** Number of B-tree scans during monitoring period

**Hash Scans** Number of hash scans during monitoring period

## Checkpoint Statistics

Ckpts Due to Txn Lim  
Ckpts Due to AIJ Growth  
Ckpts Due to Time Lim

## Transaction Statistics

Tot Txn Durtn for Txns  
Tot Txn Durtns

## Summary All File IOs

All Files Read IOs  
All Files Write IOs  
All Files Extend IOs  
All Files Read Blks  
All Files Write Blks  
All Files Extend IOs  
All Files Read Stall Time  
All Files Write Stall Time  
All Files Extend Stall Time



## Generating and Understanding Reports

**D**BXAct for Rdb produces a report titled `<database name>.REPORT`. This outlines summary and detailed information for the various statistics it monitors. The report is produced automatically at the end of each monitoring run, after data has been gathered. You can customize the report to generate information about statistics of interest using the following file:

```
FRENDDBXHOME:STAT_PROFILE.EDT
```

### File Statistics

By default, ten important statistics are selected in the `<database name>.REPORT` file. In addition to these ten statistics, information detailing file I/O, file blocks transferred, and file stall time is always reported. Information about reads, writes, and extends is recorded about each of these three categories. An example report file will look like the following:

```
File Name: gis.report

 DBXAct - Summary Report for isg$data2:[gis]gis.rdb

Start Date/Time: 10/12/94 12:34:57 End Date/Time: 10/12/94 13:40:33

Total System Memory: 131072 Available System Memory: 32628
Buffer Hit Rate (%): 92.96 GB Direct IO Savings (%): 7.08

 Database Statistics

Statistic Name Total Avg Value Peak Value Day/Time

Transactions 17.00 0.00 5 12/12:40
Data Read Time 4.00 0.00 1 12/12:40
Data Write Time 2.00 0.00 1 12/12:40
Locks Requested 3653.00 125.00 841 12/12:40
Stall Time 4.00 0.00 2 12/12:40
Records Fetched 1762.00 60.00 354 12/12:40
Records Fragmented 0.00 0.00 0 12/12:40
All Files Read IOs 1154.00 39.00 115 12/12:40
All Files Write IOs 336.00 11.00 42 12/12:40
```

## File I/O Statistics

| File Name            | Read<br>I/O | Write<br>I/O | Extend<br>I/O | Total<br>I/O | % of<br>Total I/O |
|----------------------|-------------|--------------|---------------|--------------|-------------------|
| RDB\$SYSTEM          | 24          | 2            | 0             | 26           | 16.56             |
| AP_COMP_NM_IDX       | 2           | 2            | 0             | 4            | 2.55              |
| AP_STATUS_IDX        | 1           | 3            | 0             | 4            | 2.55              |
| AP_TABLE             | 4           | 6            | 0             | 10           | 6.37              |
| BUDGET_ITE_IDX       | 2           | 0            | 0             | 2            | 1.27              |
| BUDGET_ITE_TBL       | 2           | 0            | 0             | 2            | 1.27              |
| CHECK_AP_N_IDX       | 2           | 2            | 0             | 4            | 2.55              |
| CHECK_BUDG_IDX       | 1           | 3            | 0             | 4            | 2.55              |
| CHECK_NUMB_IDX       | 0           | 2            | 0             | 2            | 1.27              |
| CK_TABLE             | 3           | 3            | 0             | 6            | 3.82              |
| GIS_SMALL_TABLE_AREA | 3           | 0            | 0             | 3            | 1.91              |
| LEADS_TABLE          | 19          | 1            | 0             | 20           | 12.74             |
| LEAD_HISTO_TBL       | 13          | 0            | 0             | 13           | 8.28              |
| LEDGER_ENT_IDX       | 0           | 2            | 0             | 2            | 1.27              |
| PO_LINE_IT_IDX       | 0           | 2            | 0             | 2            | 1.27              |
| QUOTE_COMP_IDX       | 3           | 0            | 0             | 3            | 1.91              |
| QUOTE_MAST_TBL       | 20          | 0            | 0             | 20           | 12.74             |
| VENDOR_NAM_IDX       | 3           | 2            | 0             | 5            | 3.18              |
| VENDOR_SEA_IDX       | 2           | 1            | 0             | 3            | 1.91              |
| VENDOR_TAB_TBL       | 3           | 1            | 0             | 4            | 2.55              |
| Total                | 115         | 42           | 0             | 157          |                   |
| Pct of Total         | 73.25%      | 26.75%       | 0.00%         |              |                   |

## File Blocks Transferred

| File Name            | Read<br>Blocks | Write<br>Blocks | Extend<br>Blocks | Total<br>Blocks | % of<br>Total Blocks |
|----------------------|----------------|-----------------|------------------|-----------------|----------------------|
| RDB\$SYSTEM          | 288            | 6               | 0                | 294             | 19.39                |
| AP_COMP_NM_IDX       | 24             | 8               | 0                | 32              | 2.11                 |
| AP_STATUS_IDX        | 12             | 12              | 0                | 24              | 1.58                 |
| AP_TABLE             | 39             | 30              | 0                | 69              | 4.55                 |
| BUDGET_ITE_IDX       | 24             | 0               | 0                | 24              | 1.58                 |
| BUDGET_ITE_TBL       | 24             | 0               | 0                | 24              | 1.58                 |
| CHECK_AP_N_IDX       | 24             | 6               | 0                | 30              | 1.98                 |
| CHECK_BUDG_IDX       | 12             | 12              | 0                | 24              | 1.58                 |
| CHECK_NUMB_IDX       | 0              | 12              | 0                | 12              | 0.79                 |
| CK_TABLE             | 27             | 9               | 0                | 36              | 2.37                 |
| GIS_SMALL_TABLE_AREA | 36             | 0               | 0                | 36              | 2.37                 |
| LEADS_TABLE          | 228            | 3               | 0                | 231             | 15.24                |
| LEAD_HISTO_TBL       | 156            | 0               | 0                | 156             | 10.29                |
| LEDGER_ENT_IDX       | 0              | 12              | 0                | 12              | 0.79                 |
| PO_LINE_IT_IDX       | 0              | 8               | 0                | 8               | 0.53                 |
| QUOTE_COMP_IDX       | 36             | 0               | 0                | 36              | 2.37                 |
| QUOTE_MAST_TBL       | 240            | 0               | 0                | 240             | 15.83                |
| VENDOR_NAM_IDX       | 36             | 8               | 0                | 44              | 2.90                 |
| VENDOR_SEA_IDX       | 24             | 4               | 0                | 28              | 1.85                 |
| VENDOR_TAB_TBL       | 27             | 3               | 0                | 30              | 1.98                 |
| Total                | 1349           | 167             | 0                | 1516            |                      |
| Pct of Total         | 88.98%         | 11.02%          | 0.00%            |                 |                      |

## Customized Reports

DBXAct monitors many types of statistics. These statistics provide valuable information regarding database activity and can be used as a tool in improving performance. DBXAct's summary report furnishes detailed and summary information on any statistic selected in the file:

```
FRENDDBXHOME:STAT_PROFILE.EDT.
```

Ten critical statistics are selected by default, but DBXAct provides a simple method for altering these statistics. Simply editing the file **STAT\_PROFILE.EDT** and placing an **\*** directly to the left of any statistic causes DBXAct to report this statistic. The file **STAT\_PROFILE.EDT** that was used to produce the above report is on the next page in the section titled **"Report Data."** Notice that there are no spaces between the **\*** and the statistic name, and that the **\*** is the only character that needs to be added. Removing the **\*** causes the statistic to *not* be included in the reports.

For each selected statistic, DBXAct reports its total change value over the monitoring period, its average value, peak value, and the time at which that peak occurred.

## Report Data

Since DBXAct records cumulative values, taking the difference between data points during a time increment and adding the differences creates the report.

Example of the file `STAT_PROFILE.EDT` :

**Note** Only those statistics preceded by an “\*” character will be collected during the monitoring session.

```
*Transactions
Verb Successes
Verb Failures
Ckpts Due to Txn Lim
Ckpts Due to AIJ Growth
Ckpts Due to Time Lim
In GB, Rt Ver, Need Lck
Tot Txn Durtn for Txns
Tot Txn Durtns
*Data Read Time
*Data Write Time
*Locks Requested
Locks Unlocked
Locks Promoted
Locks Demoted
Requests Stalled
Request Deadlocks
Proms Stalled
Prom Deadlocks
*Stall Time
PIO Fetches For IO
PIO Spam Fetches For IO
*Records Fetched
*Records Fragmented
Records Stored
Pages Checked
B-tree Scans
Hash Scans
*All Files Read IOs
*All Files Write IOs
All Files Read Blks
All Files Write Blks
All Files Read Stall Time
All Files Write Stall Time
```

## Using C/S Control Center with DBXAct

You can use C/S Control Center to replay DBXAct for Rdb activity files and create reports.

► **To replay DBXAct activity files and create reports:**

1. Run DBXAct on the host machine where the database resides. The default run time is eight hours with a ten-second scan interval. When running DBXAct for this long a period, you may wish to increase the scan interval (to 60 seconds or higher, for example) to reduce the system load.
2. After completing its run, DBXAct will produce a file with a CAE extension. This file can be replayed using C/S Control Center.
3. Use any File Transfer utility (such as FTP) to move the CAE file from the host down to the PC and place it in the ARCHIVE directory under the CC installation.

(ex. C:\Cntrlr\16bit\CC\ARCHIVE)

**Note** C/S Control Center may not be able to read the file that DBXAct generates. Originally, the imported file is not formatted in a method that Windows will recognize. This occurs because this file may not have the <CR>s (carriage returns) properly placed in the file. You must save the generated file in the DOS 8.3T [filename.ext] format. If you encounter any errors while C/S Control Center is attempting to open or read the file, you should follow the procedure below to ensure that the file is correctly formatted.

To insert the <CR>s, open the file in Windows Notepad or the DOS Editor. Then save the file. Do **NOT** make any changes! This process of opening and saving the file will insert the necessary <CR>s in the file, which can then be read by C/S Control Center.

4. Start C/S Control Center.

5. Click the **File** menu and select **Archive**. The new screen will show:
  - a) Date
  - b) Time
  - c) Number of Samples
  - d) File Size
  
6. Highlight the **CAE** file and then click **Replay**. The new screen will show:
  - a) Statistics collected
  - b) Object [database] from which statistics were collected
  
7. To choose the statistics you want to display:
  - a) Select individual statistics by clicking the box to put an **X** in it, or
  - b) Click **Select All** to view all statistics at once.
  
8. Click **OK**.
  
9. Click **Start** in the C/S Control Center / Archive Replay - Graph screen. C/S Control Center will then load the data, but you must click **Start** in order to view it. Other capabilities that are available while viewing the data are:
  - a) Pause
  - b) Speed up
  - c) Slow down
  - d) Choose a specific point and read the statistics gathered for that point in the scan by pausing the graph movement and double-clicking that point in the graph.
  
10. After viewing, click **Close** to exit the **Graph** screen.
  
11. Click **Close** to exit the **Statistics Selection** screen.
  
12. Click **Report**. This feature will generate a small report for any of the collected statistics.

13. To choose the statistics you want to display:
  - a) Select individual statistics by clicking the box to put an **X** in it, or
  - b) Click **Select All** to view all statistics.
  
14. Click **Generate** to create the report. The report will display:
  - a) Average and Peak values during the monitoring period
  - b) Date and Time the peak values were reached
  - c) Date and Time the minimum values were reached
  - d) Peak Transactions and Average Transactions

**Note** All information that you have created and viewed using C/S Control Center is printable for other means of reporting.

## DBXAct Logical Names

---

### **FREND\$DBX\$HOME**

This logical points to the “home” area for DBXAct. By home, we mean the disk and directory where the DBXAct executable and other distribution files reside. By modifying the system startup command file as specified in the installation instructions, this logical will be reassigned each time the system is rebooted.

### **FREND\$DBX\$DATABASE**

This logical points to the database that will be monitored by DBXAct. It should be set at the **PROCESS** level. The value assigned to this logical should be of the form:

```
<DISK> : [<DIRECTORY>] <DATABASE> .RDB .
```

### **DBXACT\$SHUTDOWN\_FLAGS**

This logical can be set at the **SYSTEM** level to force a particular DBXAct job to abort before its specified end time. If for some reason you or the system manager needs to stop DBXAct while it is in the middle of a long monitoring operation, assigning the value of the **PID** for the DBXAct process to this logical will cause the DBXAct process to complete processing whenever it wakes up from its hibernation or wait state, after the next scan. Therefore, although the DBXAct process may not stop immediately after setting this logical, it will begin to shut down after completing the next database scan.

Use of this logical is preferable to stopping the process and allows DBXAct to complete the creation of the DBTune activity file.



---

## Answers to Commonly Asked Questions

---

### Why does DBXAct run only in batch mode?

DBXAct simply collects statistics and reports them after the monitoring period in both the report file and the DBTune output file. Therefore, it only needs to be executed in batch mode.

### Can DBXAct monitor multiple databases at the same time?

Not in the current release, but future releases will enable this characteristic. Several different DBXAct processes can be started, however. In this way, each process can monitor a different database.

### What happens if the process is shut down while DBXAct is running?

If you need to stop DBXAct while it is in the middle of a long monitoring operation, you should follow these steps:

1. Use the **DCL** command

```
show sys/page
```

to locate the **PID** of the DBXAct process to be stopped. Note the **PID** number.

2. Use the **DCL** command

```
define/system DBXAct$shutdown_flags #
```

where # is the **PID** of the DBXAct process you wish to stop.

Example assignment of the DBXACT\$SHUTDOWN\_FLAGS:

```
$ define/system DBXAct$shutdown_flags 25A02703
```

### **How can I use Activity values to tune my database?**

Philosophies vary in this area; some suggest tuning based on 50 to 150 percent of average activity, while others suggest 60 to 100 percent of peak, and still others suggest tuning to available resources. No single answer is more correct than the others are, and either of the first two choices must be weighed against the resources of the system.

### **What Rdb versions will DBXAct monitor?**

DBXAct supports versions of Rdb from 4.2-13 through Rdb 7.1.

# Appendix A

---

## Customer Support

If you experience problems installing or using Rdb Controller for Rdb, please contact Customer Support at ALI.

### How to Contact Customer Support:

**Telephone:** (866) 257-8970  
(803) 648-5931

**Fax:** (803) 641-0345

**Web Address:** [www.aliconsultants.com](http://www.aliconsultants.com)

**Support E-Mail:** [support@aliconsultants.com](mailto:support@aliconsultants.com)

International clients may also obtain support through their local distributor's office.



# Appendix B

---

## Release Notes

Release notes and major changes since the last release may be found on the VMS CD in the file VMS\_RELEASE\_NOTES.TXT.



# Appendix C

---

## Glossary of Technical Terms

This appendix provides definitions of database terms and network terms that you may encounter in this manual or that may generally apply to the use of our products.

### Database Terms

| Term                | Description                                                                                                                                                                                                                                                                                |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AIJ</b>          | In Rdb, image journaling stores copies of database rows after they are updated and committed. Records are stored in one or more AIJ files. After a system failure, AIJ can be used to reconstruct a database to include the last successfully completed transaction.                       |
| <b>Availability</b> | A measure of the percentage of time an application is up or down, or the percentage of time the application must be up and running (0 - 100 percent).                                                                                                                                      |
| <b>Cluster</b>      | Tables that are frequently accessed together may be physically stored together. To store them together, a <i>cluster</i> is created to hold the tables. The data in the tables is then stored together to minimize the number of I/Os that must be performed and thus improve performance. |

---

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Data (DML) Locks</b>                    | The maximum number of DML locks is one for each table modified in a transaction. The value should equal the grand total of locks on tables referenced by all users.                                                                                                                                                                                                                                                                                                                                |
| <b>Data Buffer Cache</b>                   | The portion of reserved database memory that stores copies of physical data blocks. By holding data in the Data Buffer Cache, data can be more speedily accessed than from physical storage.                                                                                                                                                                                                                                                                                                       |
| <b>Data Buffer Hit Ratio</b>               | Determined by dividing the number of cache hits (logical reads that require no physical reads) by the number of logical reads and multiplying by one hundred. The resulting percentage, when monitored over an extended period of time, is a valuable indication of how successful a database is in maintaining frequently accessed data in memory.                                                                                                                                                |
| <b>Data File</b>                           | An Oracle tablespace is comprised of one or more physical files. The data files associated with a tablespace store a database's data in that tablespace.                                                                                                                                                                                                                                                                                                                                           |
| <b>Database File</b>                       | Physical files that contain all database information. For Oracle, this would include control files, redo log files, and data files. For Sybase and SQL Server, this would include devices. These files can be manipulated by the operating system.                                                                                                                                                                                                                                                 |
| <b>Database File Multiblock Read Count</b> | Used to multi-block I/O. This is the maximum number of blocks readable in one I/O operation during a sequential scan. Values in the range of 4 to 32 are reasonable. The actual maximums are operating system specific. This term does not apply to Rdb.                                                                                                                                                                                                                                           |
| <b>Data Source</b>                         | These are the data conduits between applications and Empirical Director. As you define applications in Empirical Director, you indicate the data source for each application.                                                                                                                                                                                                                                                                                                                      |
| <b>DB Block Buffers</b>                    | This is an INIT.ORA parameter for Oracle that determines the number of database blocks cached in memory (one buffer equals one block). Because this parameter affects how frequently a block is stored in memory, it has a significant impact on Oracle database performance. Increasing the value for this parameter will increase the likelihood that data will be stored in memory (consequently increasing the Data Buffer Cache Hit Ratio), but at the expense of greater memory consumption. |



---

|                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DB Block Size</b>               | The number of bytes per database block. Typical values are 2K (2048 bytes) and 4K (4096 bytes). This value typically does not change after being set at the time of database creation.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Device</b>                      | A Sybase or SQL Server database is comprised of one or more physical files called devices. The database engine uses devices to store database data.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Dictionary (DDL) Lock</b>       | Protects the definition of a schema object while that object is being accessed by an ongoing DDL transaction. Whereas a DDL lock is automatically acquired by the database during any DDL operation, users cannot explicitly request DDL locks. DDL Locks fall into the following categories: exclusive, shared, and breakable parse (for Oracle), exclusive, shared, intent, update, and demand (for Sybase and SQL Server).                                                                                                                  |
| <b>Dictionary Cache</b>            | Used by Oracle during database operation to ascertain that objects exist and that users have accessed them properly. Oracle also updates the Data Dictionary continuously to reflect changes in database structures, auditing, granting, and data. Data Dictionary Cache is the part of the Data Dictionary cached in SGA using the LRU (least recently used) algorithm for fast access.                                                                                                                                                       |
| <b>Dictionary Cache Read Count</b> | Related to the following statistics stored in the V\$ROWCACHE table: GETS shows the total number of requests for information on the corresponding items; GETMISSES shows the number of data requests resulting in cache misses. For frequently accessed Dictionary Caches, the Dictionary Cache Hit Ratio should be higher than 90 percent. To increase the Dictionary Cache Hit Ratio, increase the value of the initialization parameter SHARED_POOL_SIZE. This increases the memory available to the data dictionary cache. Used by Oracle. |
| <b>Distributed Lock</b>            | Used by Oracle to ensure that the data and other resources distributed among the various instances of an Oracle Parallel Server remain consistent.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Execution (Explain) Plan</b>    | To execute a DML statement, a database may have to perform many steps. Each of these steps either physically retrieves rows of data from the database or prepares them in some way for the user issuing the statement.                                                                                                                                                                                                                                                                                                                         |

---

|                                |                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Export</b>                  | An Oracle utility used to write data from an Oracle instance to the file system. These data, stored as export dump files, contain schema definitions and their contents that can be later imported back into the database. In Rdb, export is a SQL command rather than a separate utility.                                                                                       |
| <b>FreeLists</b>               | Used by Oracle to point to the available free blocks. The parameter, FreeLists, is used to specify the number of freelists for a segment. This parameter is particularly important when a large number of inserts have to be carried out in parallel. In this case the freelists parameter should be set to be greater than one.                                                 |
| <b>Import</b>                  | An Oracle utility used to read data from an Oracle dump file on the file system into an Oracle instance. In Rdb, import is a SQL command rather than a separate utility.                                                                                                                                                                                                         |
| <b>Index</b>                   | An object created on a table to increase performance of data retrieval from the indexed table. The index is logically and physically independent of the table data.                                                                                                                                                                                                              |
| <b>Initial</b>                 | Used by Oracle to set the size of the first extent of a segment. This storage area is allocated when the segment is first created, and additional extents are added on later as needed.                                                                                                                                                                                          |
| <b>Latch (Internal Lock)</b>   | In Oracle, a simple low-level serialization mechanism to protect shared data structures in the SGA.                                                                                                                                                                                                                                                                              |
| <b>Library Cache</b>           | In Oracle, the library cache contains shared SQL and PL/SQL areas. The Library Cache is contained in the shared pool (Shared SQL Area) within the SGA and is available to multiple, concurrent users.                                                                                                                                                                            |
| <b>Library Cache Hit Ratio</b> | In Oracle, the Library Cache Hit Ratio is the ratio of shared SQL and PL/SQL items found in the Library Cache versus physical storage. Related to the following statistics stored in the V\$LIBRARYCACHE table: PINS, which shows the number of times an item in the library cache was executed; and RELOADS, which shows the number of library cache misses on execution steps. |

---

|                                             |                                                                                                                                                                                                                                                                                                  |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Locking Mechanism</b>                    | A database automatically locks a resource on behalf of a transaction to prevent other transactions from performing a task that requires exclusive access to the same resource. The lock is automatically released when certain events occur and the transaction no longer requires the resource. |
| <b>MaxExtents</b>                           | Used by Oracle to specify the maximum number of extents that can be allocated for a single segment.                                                                                                                                                                                              |
| <b>MinExtents</b>                           | Used by Oracle to specify the number of extents that should be allocated when the segment is first created. If <code>minextents = 3</code> , for example, the initial extent plus two times the next extent are allocated when the segment is first created.                                     |
| <b>Next</b>                                 | Used by Oracle to specify the size of the extents to be created after the first (initial) extent. This area is not allocated until necessary.                                                                                                                                                    |
| <b>Optimal</b>                              | Used by Oracle to specify the optimum size for a rollback segment. This parameter can only be set for rollback segments.                                                                                                                                                                         |
| <b>Parallel Cache Management (PCM) Lock</b> | A distributed lock used by Oracle that covers one or more data blocks (table and index blocks) in the buffer cache.                                                                                                                                                                              |
| <b>PctFree</b>                              | Sets the percentage of free space within a data block that will be reserved for possible updates to rows already stored within each data block of the segment. This term does not apply to Rdb.                                                                                                  |
| <b>PctIncrease</b>                          | The percentage by which each incremental extent of a segment grows over the previous incremental extent allocated. This term does not apply to Rdb.                                                                                                                                              |
| <b>PctUsed</b>                              | Sets the percentage of free space within a data block that must be available before row insertion into the data block will be allowed. This term does not apply to Rdb.                                                                                                                          |
| <b>Pipes</b>                                | These are the data conduits between applications and Empirical Director, also commonly referred to as “data sources.” As you define applications in Empirical Director, you indicate the data source for each application.                                                                       |

---

|                                      |                                                                                                                                                                                                                                                                       |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Redo Log File</b>                 | Contains committed transactions that have not yet been written to the respective data files. This term does not apply to Rdb.                                                                                                                                         |
| <b>Referential Integrity</b>         | A rule defined on a column or set of columns in one table that allows you to insert or update a row only if the value for that column or set of columns (in the child table) matches the value in a column of a related table (parent table).                         |
| <b>Response Time</b>                 | The average time required to perform a transaction. Slow response time may be attributed to system-wide bottlenecks or to individual application problems.                                                                                                            |
| <b>Rollback Segment</b>              | A logical structure used by Oracle responsible for undoing uncommitted transactions.                                                                                                                                                                                  |
| <b>RUJ</b>                           | In Rdb, the recovery unit journaling stores copies of database rows before they are updated. The RUJ files are used to undo uncommitted updates to a database when a rollback is performed or a system failure occurs.                                                |
| <b>SGA Size</b>                      | This is an Oracle parameter that sets the size of main memory allocated for exclusive use by Oracle.                                                                                                                                                                  |
| <b>SQL</b>                           | SQL is short for Structured Query Language. This is a specialized programming language for sending queries to databases.                                                                                                                                              |
| <b>Sysop</b>                         | Sysop means System Operator. The Sysop is anyone responsible for the physical operations of a computer system or network resource. A System Administrator decides how often backups and maintenance should be performed and the System Operator performs those tasks. |
| <b>Table</b>                         | The basic unit of data storage in a database. Within a table, data is stored in rows and columns. Each column is given a column name, a datatype, and a width and precision or scale. A row is a collection of column information corresponding to a single record.   |
| <b>Table Scans<br/>(Long Tables)</b> | The total number of full table scans performed on tables with more than 5 db_blocks. When the number of full table scans is greater than 0 per transaction, the SQL statements in the application would benefit from tuning. This term does not apply to Rdb.         |

|                                   |                                                                                                                                                                                                                                                                  |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Table Scans (Short Tables)</b> | The number of full table scans performed on tables with less than 5 db_blocks. It is optimal to perform full table scans on short tables rather than using indexes. This term does not apply to Rdb.                                                             |
| <b>Tablespace</b>                 | Logical storage space for database objects, such as tables, indexes, clusters, etc. Physically, a tablespace consists of one or more database files. Using multiple tablespaces allows for the logical separation of user data. This term does not apply to Rdb. |
| <b>Throughput</b>                 | The number of transactions per minute (0- 1,000,000) that the system is handling. This “workload” monitoring may include information on the number and type of transactions being executed, who is executing them, and the applications being used.              |
| <b>Transaction Efficiency</b>     | A ratio of the total number of transactions compared with the number of transactions that complete successfully. Also expressed as the percentage of successful database transactions (0 - 100 percent).                                                         |
| <b>Trigger</b>                    | A procedure that is implicitly executed when an INSERT, UPDATE, OR DELETE statement is issued against the associated table.                                                                                                                                      |

## Network Terms

| <b>Term</b>      | <b>Description</b>                                                                                                                                                                                                                                          |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Backbone</b>  | Usually a high-performance network cable, such as fiber optic or thick wire, that provides the attachment point for branching network segments. For instance, a backbone may run completely through a building while hubs are attached at different points. |
| <b>Bandwidth</b> | Bandwidth refers to how much information can be sent through a connection. Usually bandwidth is measured in bits-per-second. A full page of English text is about 16,000 bits.                                                                              |
| <b>Baud</b>      | In common usage the baud rate of a modem is how many bits it can send or receive per second.                                                                                                                                                                |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Bit</b>         | An abbreviation of Binary DigIT. A bit is the smallest unit of computerized data. Bandwidth is usually measured in bits-per-second.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>CGI</b>         | CGI means Common Gateway Interface. CGI is a set of rules that describe how a Web server communicates with another piece of software on the same machine, and how the other piece of software (the CGI program) talks to the Web server. Usually a CGI program takes data from a Web server and does something with it, like putting the content of a form into an e-mail message, or turning the data into a database query.                                                                                                                                                                   |
| <b>Client</b>      | Any device in a network that uses services from a server or servers. A PC is the most typical client on a network. Client may also refer to a software program that is used to contact and obtain data from a Server software program on another computer, usually across a great distance. Each Client program is designed to work with one or more specific kinds of Server programs, and each Server requires a specific kind of Client. (Also see "Server.")                                                                                                                                |
| <b>Ethernet</b>    | Ethernet is the most popular network medium in use today and is most popular in VAX, UNIX, and PC networks. It is a common method of networking computers in a LAN.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hub</b>         | Typically a device used to connect network cables. Twisted-pair systems use hubs to connect multiple cable segments to a backbone or other cable segment.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>LAN</b>         | Local Area Network. Network connecting systems that are usually in a single department, single building, or group of buildings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>MIB Modules</b> | MIB modules usually contain object definitions, may contain definitions of notifications, and sometimes include compliance statements specified in terms of appropriate object groups. MIB modules define the management information maintained by the instrumentation in managed nodes, made remotely accessible by management agents, conveyed by the management protocol, and manipulated by management applications. In general, management information defined in any MIB module, regardless of the version of the data definition language, can be used with any version of the protocol. |

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SNMP</b>       | SNMP stands for Simple Network Management Protocol, a set of network communication specifications that cover the basics of network management in a method that poses little stress on the existing network. SNMP is a set of standards for communication with devices connected to a TCP/IP. Empirical Director uses SNMP to collect performance data.                                                                                                                                                         |
| <b>Server</b>     | A computer, or a software package, that provides a specific kind of service to client software running on other computers. The term can refer to a particular piece of software or to the machine on which the software is running. Any device that provides services to a network is considered to be a server. This may include file and print servers (such as UNIX, Novell NetWare, Windows NT Server, or VMS), terminal servers, application servers, and other specialized devices. (Also see “Client.”) |
| <b>TCP/IP</b>     | TCP/IP is short for Transmission Control Protocol/Internet Protocol. This is the suite of protocols that defines the Internet.                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Telnet</b>     | The command and program used to login from one Internet site to another. The telnet command/program gets you to the login: prompt of another host.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Token Ring</b> | A newer network medium (relative to Ethernet) that is predominantly used in IBM shops.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Topology</b>   | The physical architecture of the network, including cabling and configuration (such as a star, bus, or ring). Topologies have blurred during the last few years because most networks now consist of multiple topologies and cable types.                                                                                                                                                                                                                                                                      |
| <b>UNIX</b>       | A computer operating system (the basic software running on a computer, underneath things like word processors and spreadsheets). UNIX is designed to be used by many people at the same time and has TCP/IP built in. It is the most common operating system for servers on the Internet.                                                                                                                                                                                                                      |
| <b>WAN</b>        | Wide Area Network. A network that usually interconnects LANs in different buildings or areas. For instance, a LAN in Building A connected to a LAN in Building B may be called a WAN.                                                                                                                                                                                                                                                                                                                          |

# Index

- 
- Analyzing data gathered with DBXAct, 158
  - C/S Control Center, 165, 166, 167
  - Customer support, 171
  - DBAnalyzer
    - Available queues, 59
    - Batch mode, 5, 13
    - Batch use, 15
    - Complexity rating, 21
    - Database
      - Macro view, 5
      - Micro-Index view, 6
      - Micro-Storage view, 6
      - Micro-Table view, 5
    - Database views, 5
    - DEC AXP/Open VMS, 9
    - DEC VAX/Open VMS, 8
    - Domain sort order, 53
    - Execution options, 37, 38, 57
    - Full format report, 45
    - Getting started, 8
    - Hashed Index percentage, 22
    - Hash-to-Sort/Sort-to-Hash ratio, 23
    - Installation, 10, 11
    - Integrity rating, 22
    - Keystrokes for online execution, 19
    - Macro Mode windows, 39, 40, 41, 42, 43
    - Macro View windows, 23, 24, 25, 26
    - Micro-Index view, 28, 29
    - Micro-Storage area view, 29, 30
    - Micro-Table view, 27, 28
    - Online use, 14
    - Rdb statistics, 21
    - Report components, 53
    - Report generation
      - keystrokes, 20
    - Report output options, 60
    - Report parameters, 52
    - Report printing, 58
    - Reports, 31, 32, 33, 34, 35, 36
    - Selected storage areas, 54
    - Selected tables, 57
    - Selected views, 55
    - Storage area allocation, 22
    - Storage area selection, 54
    - Table detail options, 56
    - Table selections, 56
    - Tune and Complexity ratings, 6
    - Tune rating, 21
    - View selections, 55
  - DBTune
    - Analyze Rdb, 79
    - Batch use, 74
    - Database structure, 109
    - DEC AXP/Open VMS, 68
    - DEC VAX/Open VMS, 65
    - Disk utilization, 120
    - Getting Started, 65
    - Help, 143
    - Installation, 70
    - Keystrokes for online execution, 77
    - Load parameters, 80
    - Online use, 73
    - Parameters, 76
    - Performance analysis, 119
    - Performance Analysis Data (PAD) file, 111
    - Rdb database transformation, 122
    - Reports, 128, 130, 131, 132, 133, 134, 135, 136, 137, **138**, **139**, **140**, **141**, **142**
    - Row cache, 63
    - Sorted ranked indexes, 63
    - Temporary tables, 62
    - Transform main driver, 123
    - Transformation process, 78, 123
    - Workload data, 110
  - DBTune PAD file
    - Cluster, 112
    - Contra, 113
    - DBDISKS, 111
    - DYNAMIC\_WORKLOAD\_FILE, 114
    - Index, 112



- MODPAD\_FILE, 114
  - Table, 111
- DBTune Parameters
  - BACKUP\_DIR, 92
  - BIAS, 94
  - CONCEAL\_LOGS, 102
  - DBDISKnn, 84
  - DBDISKS, 83
  - DYNAMIC\_WORKLOAD\_FILE, 91
  - EDIT\_FILES, 87
  - EXPORT\_UNLOAD\_DIR, 93
  - FILL, 95
  - FREND\_EDITOR, 88
  - GROWTH, 95
  - LOAD\_TIME\_LIM, 102
  - LOGICAL\_TYPE, 101
  - LOGICALS, 100
  - MACHINE\_VUPS, 103
  - MAX\_BUFFER\_SIZE, 98
  - MAX\_BUFFERS, 98
  - MAX\_DB\_USERS, 100
  - MAX\_PAGE\_SIZE, 97
  - MIN\_BUFFER\_SIZE, 97
  - MIN\_BUFFERS, 98
  - MIN\_PAGE\_SIZE, 96
  - MODPAD\_FILE, 88
  - RUJ\_DIR, 94
  - SA\_MIN\_CARD, 105
  - SAVE\_COMMENTS, 105
  - SMALL\_HASHED, 107
  - SMALL\_SORTED, 106
  - SMALL\_TABLE, 106
  - SNP\_PERC, 96
  - SQL\_DIR, 92
  - STOR\_AREA\_SPREAD, 100
  - STRATEGY, 81
  - SYS\_MEM\_PAGES, 99
  - SYSGEN settings, 69
  - TABLE\_COMMIT, 104
  - TUNE\_FOR\_COMPRESSION, 108
  - TUNE\_TECHNIQUE, 82
- DBTune Reports, 128, 130, 131, 132, 133, 134, 135, 136, 137, **138, 139, 140, 141, 142**
- DBXAct
  - All file IOs summary, 160
  - Analyzing data, 158
  - Authorize settings, 150
  - Batch mode, 157, 169
  - Batch queue, 156
  - Checkpoint statistics, 160
  - Customized reports, 163
  - Features, 146
  - File statistics, 161, 162
  - General database statistics, 158
  - Generating reports, 161
  - Getting started, 149
  - Growth projection interval, 156
  - Index statistics, 159
  - Installing, 151
  - Locking summary statistics, 159
  - Logical names, 168
  - Monitoring duration, 156
  - Monitoring scan rate, 156
  - Overview, 147
  - PIO data fetch statistics, 159
  - Record statistics, 159
  - Report data, 164
  - Stall statistics, 158
  - Start up, 153
  - Storage area statistics, 158
  - System requirements, 149
  - Transaction statistics, 160
  - Understanding reports, 161
  - Using with C/S Control Center, 165, 166, 167
  - Variables, 156
- DEC AXP/OpenVMS, 9, 68
- DEC VAX/OpenVMS, 8, 65
- Installing DBAnalyzer, 10
- Installing DBTune, 70
- Installing DBXAct, 151
- Macro Window, 23, 24, 25, 26, 39, 40, 41, 42, 43
- Rdb Version 7.0, 6, 62, 63, 73, 146, 157
- Release notes, 173
- REVIEW\_AND\_GUIDE.REPORT**, 66, 68, 69, 78, 92, 93, 97, 98, 99, 119, 129, 130
- Tune and Complexity Ratings, 6
- Using C/S Control Center with DBXAct, 165, 166, 167