

Oracle® CODASYL DBMS

Release Notes

June 2007

Release 7.1.3.1

ORACLE®

Oracle CODASYL DBMS Release Notes, Release 7.1.3.1

Copyright © 1986, 2007 Oracle Corporation. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs is prohibited.

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

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

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software - Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

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

Oracle is a registered trademark of Oracle Corporation. All other company or product names mentioned are used for identification purposes only and may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Preface	vii
1 Installation and Documentation	
1.1 Requirements	1-1
1.2 Installation of Oracle CODASYL DBMS Software	1-1
1.3 Documentation in Adobe Acrobat Format	1-1
2 Problems Corrected	
2.1 Problems with Floating Point Data Types	2-1
2.2 AIJ Backup Operation Aborts With NONAME-F-NOMSG Message Number 00000004	2-1
2.3 DBO/SHOW STATISTICS AIJ ARB:I/O ratio, Blocks-per-I/O ratio Problems	2-2
2.4 Adding a large AIJ file to a DB fails with either an ACCVIO or OPCDEC error	2-2
2.5 DBO Tape Support Added For SDLT600, LTO2, LTO3 Drives	2-2
2.6 Hot Standby Node Failure Recovery When Using DBO/OPEN/CACHE=NOENABLE	2-3
2.7 DBO/RESTORE/MULTITHREAD exits with traceback log when the wrong version is used.	2-3
2.8 State Value Truncated on Hot Standby Statistics Display	2-3
2.9 Possible Shared Memory Corruption When Multiple Databases Attached	2-4
2.10 Bugcheck at COSI\$TIMER_GET_REQIDT With DBMS-F-NOREQIDT ...	2-4
3 Known Problems, Workarounds, and Documentation Errors	
3.1 AIJ Log Server Process May Loop or Bugcheck	3-1
3.2 VMS\$MEM_RESIDENT_USER Rights Identifier	3-1
3.3 DBM\$BIND_MAX_DBR_COUNT Documentation Clarification	3-2
4 New Features and Corrections in Previous Releases	
4.1 Corrections in Release 7.1.3.0	4-1
4.1.1 Problem with Remote Access and FETCH.USING	4-1
4.1.2 Area File not Renamed after DBO/MODIFY/RESTRUCTURE	4-2
4.1.3 DBO/SHOW STATISTICS Hot Standby Statistics State Display Field	4-3
4.1.4 File-System Caching Avoided For DBO /COPY, /MOVE, /BACKUP/MULTI and /RESTORE/MULTI Access To Database	4-3
4.1.5 DBO/BACKUP/AFTER Ignores /EDIT_FILENAME When Backup Filespec Omitted	4-3

4.1.6	Processes Don't Always Terminate After Monitor Terminates	4-4
4.1.7	Latch Hangs Possible From DBO /SHOW STATISTICS	4-4
4.1.8	LRS Shutdown Failure DBM-F-PARTDTXNERR/SYSTEM-F-NOSUCHID	4-4

Purpose of This Manual

The Oracle CODASYL DBMS release 7.1.3.1 release notes summarize new features, software corrections, restrictions, workarounds, and known problems. These release notes cover Oracle CODASYL DBMS for OpenVMS Alpha.

Intended Audience

This document is intended for users responsible for:

- System management
- Database administration
- Application programming

Document Structure

This document consists of three chapters:

Chapter 1	Describes installation requirements and location of documents
Chapter 2	Describes corrected software errors
Chapter 3	Describes known problems, restrictions, and workarounds, as well as documentation errors and omissions
Chapter 4	Describes new features and corrected software errors in release 7.1.3.0

Conventions

Oracle CODASYL DBMS is often referred to as DBMS.

The following conventions are used in this document:

word	A lowercase word in a format example indicates a syntax element that you supply.
[]	Brackets enclose optional clauses from which you can choose one or none.
{ }	Braces enclose clauses from which you must choose one alternative.
...	A horizontal ellipsis means you can repeat the previous item.

.
. .
. . .

A vertical ellipsis in an example means that information not directly related to the example has been omitted.

Installation and Documentation

This chapter contains installation and documentation information for Oracle CODASYL DBMS release 7.1.3.1.

1.1 Requirements

This version of Oracle CODASYL DBMS supports OpenVMS Alpha version 8.3.

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

- OpenVMS Alpha version 7.2 or later

1.2 Installation of Oracle CODASYL DBMS Software

Please refer to the *CODASYL DBMS V7.1 Installation Guide* for detailed Oracle CODASYL DBMS installation instructions. Oracle strongly recommends that you read the installation guide before attempting an installation.

To extract either the PostScript (PS) or text (TXT) version of the installation guide from the kit, use one of the following commands:

```
$ BACKUP <device>:DBM07131A071.A/SAVE/SEL=DBM071_INSTALL_GDE.PS  
$ BACKUP <device>:DBM07131A071.A/SAVE/SEL=DBM071_INSTALL_GDE.TXT
```

The release 7.1 installation guide is available on MetaLink and OTN in Adobe Acrobat PDF format.

1.3 Documentation in Adobe Acrobat Format

You can view the documentation in Adobe Acrobat format using the Acrobat Reader, which allows anyone to view, navigate, and print documents in the Adobe Portable Document Format (PDF). For information about obtaining a free copy of Acrobat Reader and for information on supported platforms, see the Adobe Web site at:

<http://www.adobe.com>

The Oracle CODASYL DBMS and Hot Standby documentation in Adobe Acrobat format is available on MetaLink and OTN.

Problems Corrected

This chapter describes software errors corrected in Oracle CODASYL DBMS release 7.1.3.1.

2.1 Problems with Floating Point Data Types

BUG 5929234 and 6064965

Four separate problems have been discovered with float point data types in Oracle CODASYL DBMS:

- ACCVIO on BIND using Interactive DBQ
- DBMBUGCHK with exception at MOV\$VALIDATE_FLOAT while trying to dump floating point metadata. This only occurs during the processing of another bugcheck dump.
- Bugcheck with exception at MOV\$VALIDATE_FLOAT on DML STORE or MODIFY
- Bugcheck with exception at COSI-F-FLTINV on DBO/LOAD

The first two problems are triggered by having a DDL DEFAULT clause with a value of "0" for any floating data items. The problem would not occur if the default value was anything other than "0".

The third problem is triggered by the use of the DDL CHECK clause comparing a floating point data item to "0". Any value other than "0" would have allowed the DML verb to complete successfully.

The fourth problem can occur during a DBO/LOAD operation when floating point data items have a storage attribute of ALLOCATION IS DYNAMIC and a default value is specified in the schema.

These problems were introduced with DBMS 7.1.2.4 and 7.2.0.1 and do not cause data corruption or invalid data to be stored.

These problems have now been fixed.

2.2 AIJ Backup Operation Aborts With NONAME-F-NOMSG Message Number 00000004

In rare cases, an after-image journal backup operation may fail with an unexpected incorrect status value. The actual value may vary, but at least one customer report of the problem indicated a value of 00000004. A bugcheck dump file "footprint" of this problem is:

```
***** Exception at 0054D94C : AIJBCK$GET_NEXT_JOURNAL + 00000CFC
Saved PC = 005452E8 : AIJBCK$FULL_BACKUP + 00000FF8
Saved PC = 00543C0C : AIJBCK$BACKUP + 0000113C
```

This problem has been corrected. The errant status value was the result of an uninitialized return status being passed back. The correct status is now returned.

2.3 DBO/SHOW STATISTICS AIJ ARB:I/O ratio, Blocks-per-I/O ratio Problems

There was a problem with detecting the warning thresholds set for the "Examine ARB:I/O ratio" and "Examine blocks-per-I/O ratio" options on the DBO/SHOW STATISTICS "AIJ Analysis" screen. This caused the "ARB:I/O ratio" warning:

```
## ARBs per I/O below ## threshold
```

and the the "blocks-per-I/O ratio" warning:

```
## blocks written per I/O below ## threshold
```

to not always be output when the threshold was crossed.

This problem was caused by treating these thresholds as percent values instead of count values. This problem has been fixed and the misleading percent signs have been removed from these warning messages.

2.4 Adding a large AIJ file to a DB fails with either an ACCVIO or OPCDEC error

Creating an after image journal file for a database fails with either an ACCVIO or an OPCDEC error.

The reason was a bug in the code which prematurely cleared a synchronization flag. This allowed the request created by the expired prestarted transactions timer to execute before the current request had completed. The side effect of this was a stack corruption.

As a workaround use a larger prestarted transactions timer value or, disable the prestarted transactions timer completely.

This problem has been corrected.

2.5 DBO Tape Support Added For SDLT600, LTO2, LTO3 Drives

Support for the VMS tape density and compaction values for the Super DLT600, Ultrium460 and Ultrium960 tape drives has been added to the DBO multithreaded (/BACKUP/MULTITHREAD and /RESTORE/MULTITHREAD) utilities.

This will allow the following new density values to be specified with the /DENSITY qualifier for those DBO commands that write to these drives:

```
/DENSITY = (SDLT600,[NO]COMPACTION) - Super DLT600  
/DENSITY = (LTO2,[NO]COMPACTION) - Ultrium460  
/DENSITY = (LTO3,[NO]COMPACTION) - Ultrium960
```

2.6 Hot Standby Node Failure Recovery When Using DBO/OPEN/CACHE=NOENABLE

BUG 5957364

In configurations using the Row Cache and Hot Standby features, row caching must be explicitly disabled on the standby database using the DBO/MODIFY/NOCACHE command prior to starting hot standby for the first time on the database. However, it is also possible (though not recommended) to use the DBO/OPEN/CACHE=NOENABLE command on the standby database in order to suppress row caching.

When using the DBO/OPEN/CACHE=NOENABLE command, if a system failure occurred, it was possible that the database recovery upon reopening the database would attempt to start with a very old last checkpoint location. This location was based on the row cache checkpoint from when the master database had been originally backed up to create the standby database. In some cases, the required AIJ files would be no longer online and the recovery would fail.

This problem has been corrected. The DBR process now ignores the row cache oldest checkpoint location when not recovering from a node failure when the RCS process had been active.

2.7 DBO/RESTORE/MULTITHREAD exits with traceback log when the wrong version is used.

Using DBO/RESTORE/MULTI to restore an incompatible version of a database exits with a register dump instead of simply exiting.

Example:

```
$ dbo/rest/multi/nocdd/dir=sys$disk:[]/log reproducer.dbf
%DBO-F-DB_CVT_FAIL, Cannot convert from version V7.2 to V7.1
%DBO-F-DB_CVT_FAIL, Cannot convert from version V7.2 to V7.1

Improperly handled condition, image exit forced by last chance handler.
Signal arguments:  Number = 0000000000000008
                   Name  = 0000000002C88B94
                   0000000000000004
                   0000000000000007
                   0000000000000002
                   0000000000000007
                   0000000000000001
                   00000000003AEB34
                   100000000000001B
...

```

The reason was to allow this condition to be restartable which it truly is not.

This problem has been corrected.

2.8 State Value Truncated on Hot Standby Statistics Display

Bug 6044632

Previously, it was possible when using a Hot Standby TCP/IP port number greater than 9999 that the port number would be truncated on the DBO /SHOW STATISTICS Hot Standby Statistics display.

For example, when using a TCP/IP port number of 12345, the state display could be shown as "State: TCP/IP:1234"

This problem has been corrected. The state display field now allows a 5-digit TCP/IP port number to be displayed.

2.9 Possible Shared Memory Corruption When Multiple Databases Attached

Starting with Oracle CODASYL DBMS V7.1.2 and Oracle CODASYL DBMS V7.2.0, it was possible for shared memory to become corrupt. The corruption often would appear as (or would be caused by) data from one DBMS root file being written into the shared memory for another database. Once this corruption has occurred, reliability and functionality of the database and database users can be compromised.

Conditions leading to this corruption include:

- Processes accessing multiple databases
- Multiple database users
- Databases accessed from multiple nodes in a cluster
- Databases configured with node count greater than 1.

The memory corruption was caused by incorrect IO buffer synchronization while refreshing root file information in to shared memory.

This problem has been resolved. Oracle strongly recommends that customers with applications or procedures that may attach or more than one database at a time upgrade to this or a later release to avoid this potential memory corruption problem.

2.10 Bugcheck at COSI\$TIMER_GET_REQIDT With DBMS-F-NOREQIDT

Applications using the fast commit feature that periodically detach and reattach to a database within the same program run may eventually run out of DBMS timer blocks and bugcheck with a footprint similar to:

```
***** Exception at 01235994 : COSI$TIMER_GET_REQIDT + 00000294
%DBMS-F-NOREQIDT, reached internal maximum number of simultaneous
timer requests
Saved PC = 01235C38 : COSI_TIMER_SET + 00000288
Saved PC = 01235DA8 : COSI_TIMER_SLEEP + 00000078
Saved PC = 012A7548 : KOD$COMMIT + 00000508
```

This problem has been corrected. The problem was caused by an internal timer data structure being allocated but not being deallocated if the timer had expired. If the timer had not expired, the internal timer data structure was correctly deallocated. Thus, in some cases, the timer data structure was being “leaked” which could eventually lead to the bugcheck exception of “DBMS-F-NOREQIDT, reached internal maximum number of simultaneous timer requests”.

Known Problems, Workarounds, and Documentation Errors

This chapter describes known problems, restrictions, and workarounds, as well as documentation errors and omissions for Oracle CODASYL DBMS release 7.1.3.1.

3.1 AIJ Log Server Process May Loop or Bugcheck

Under unknown but extremely rare conditions, on busy databases where the After Image Journal (AIJ) Log Server process is enabled, the ALS process has been observed to enter a loop condition writing AIJ information to the .AIJ files.

In the worst case, this problem could cause all available journal files to be filled with repeating data. If no remedial action were taken, this condition could cause the database to be shut down, and the AIJ journals to be considered inaccessible.

The database is not corrupted by this problem.

Stopping and restarting the ALS process will clear the looping condition, even if the ALS process must be stopped using the STOP/ID command.

Stopping the ALS process will not impact production as AIJ writes automatically revert to the non-ALS behaviour.

In this release, the behavior has been changed so that if this problem is detected, the ALS process will automatically shut down, producing a bugcheck dump file. This will prevent any danger of filling all available journals and ensure that the database remains available.

ALS may be safely restarted immediately as the conditions that cause such a loop are resolved during recovery of the ALS process.

3.2 VMS\$MEM_RESIDENT_USER Rights Identifier

Oracle CODASYL DBMS version 7.1 introduced additional privilege enforcement for the database or row cache qualifiers MEMORY_MAPPING=SYSTEM and LARGE_MEMORY. If a database utilizes any of these features, the user account that opens the database must be granted the VMS\$MEM_RESIDENT_USER rights identifier. Also, any process attempting to change these attributes, or to convert or restore a database with these attributes enabled must also hold the same right.

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

3.3 DBM\$BIND_MAX_DBR_COUNT Documentation Clarification

The following is an updated description for the DBM\$BIND_MAX_DBR_COUNT logical.

When an entire database is abnormally shut down (for example, due to a system failure), the database must be recovered in a node failure recovery mode. This recovery is performed by another monitor in the cluster if the database is opened on another node or is performed the next time the database is opened.

The DBM\$BIND_MAX_DBR_COUNT logical name and the DBM_BIND_MAX_DBR_COUNT configuration parameter define the maximum number of database recovery (DBR) processes to be simultaneously invoked by the database monitor for each database during a node failure recovery. This logical name and configuration parameter apply only to databases that do not have global buffers enabled. Databases that utilize global buffers have only one recovery process started at a time during a node failure recovery.

In a node failure recovery situation with the Row Cache feature enabled (regardless of the global buffer state), the database monitor starts a single database recovery (DBR) process to recover the Row Cache Server (RCS) process and all user processes from the oldest active checkpoint in the database.

Per-Database Value

The DBM\$BIND_MAX_DBR_COUNT logical name specifies the maximum number of database recovery processes to run at once for each database. For example, if there are 10 databases being recovered and the value for the DBM\$BIND_MAX_DBR_COUNT logical name is 8, up to 80 database recovery processes would be started by the monitor after a node failure.

The DBM\$BIND_MAX_DBR_COUNT logical name is translated when the monitor process opens a database. Databases must be closed and reopened for a new value of the logical to become effective.

New Features and Corrections in Previous Releases

4.1 Corrections in Release 7.1.3.0

4.1.1 Problem with Remote Access and FETCH..USING

A problem has been uncovered with Oracle CODASYL DBMS when using remote database access with either DBQ or DML applications. If you attempt to FETCH a record via a USING clause, the fetch may fail with a DBM-F-END condition, even though the record does exist.

The problem will ONLY occur if one of the data items specified in the USING clause is the last data item defined in that record.

The error does not occur with local database access or with remote access when using the WHERE clause.

For example, given the following schema:

```
AREA NAME IS A1
RECORD NAME IS R1
  WITHIN A1
    ITEM NAME IS I1
      TYPE IS CHARACTER 5
    ITEM NAME IS I2
      TYPE IS CHARACTER 5
    ITEM NAME IS I3
      TYPE IS CHARACTER 5
SET NAME IS ALL_R1
  OWNER IS SYSTEM
  MEMBER IS R1
    INSERTION IS AUTOMATIC
    RETENTION IS FIXED
  ORDER IS SORTED BY
    ASCENDING I3
```

and assuming that there is an R1 record with the following values:

```
I1 = 'AAAAA'
I2 = 'BBBBB'
I3 = 'CCCCC'
```

The following remote query attempting to fetch record R1 will fail:

```
dbq> bind dbmfetrmtdb
dbq> ready
dbq> set noprompt
dbq> move 'CCCCC' TO I3
dbq> fetch first within ALL_R1 using I3
%DBM-F-END, end of collection
```

whereas, the same logical query using a WHERE clause will succeed:

```
dbq> fetch first within ALL_R1 where I3 eq 'CCCCC'  
I1 = AAAAA  
I2 = BBBBB  
I3 = CCCCC
```

This problem has now been fixed. No application programming changes are required.

4.1.2 Area File not Renamed after DBO/MODIFY/RESTRUCTURE

The Oracle CODASYL DBMS reload utility (DBO/MODIFY/RESTRUCTURE) moves database records from a specified target area to a new area.

In versions of DBMS prior to V7.0, the default behavior was to create the new storage area with the same filename (and in the same directory) as the target original area, with an incremented file version number. Note: only offline reload is available in pre-V70 versions.

Starting with DBMS v7.0, these defaults were modified to ensure that the new storage area filename was unique by attempting to append an "_A" (OR "_B", etc) to the storage area name. This was done as part of the work to support online reload (DBO/MODIFY/RESTRUCTURE/ONLINE), where the reload could be stopped and restarted in the middle of execution.

The idea was to make sure that there was no confusion between the original area and the new area, if the reload were stopped for any reason, and to make sure that certain file actions, such as a \$PURGE, would not delete the original area prior to reload completion.

If you wish to retain the old behavior, include the /FILE= qualifier on the DBO/MODIFY/RESTRUCTURE command and specify the original storage area filename as the parameter. This qualifier should be included on the restructure operation that performs the EXECUTE phase for offline reload, or the PREPARE phase in the case of online reload.

For example, assume that you wished to reload the BUY area in the PARTS database.

In pre-V70 offline reloads, the default would be to create a storage area, BUY.DBS;2 (assuming that BUY.DBS;1 was the original area filename). In V70 AND later, the default would be to create BUY_A.DBS;1.

To maintain the old behavior, issue DBO/MODIFY/RESTRUCTURE PARTS BUY/FILE=BUY. Note that you could also specify the /DIRECTORY qualifier to have the new storage area created in a new directory.

To modify the storage area file name of a previously reloaded area, you can rename the file, then use the DBO/ALTER utility and execute:

```
DBALTER> DEPOSIT FILE <area> SPECIFICAION <new-filename>
```

4.1.3 DBO/SHOW STATISTICS Hot Standby Statistics State Display Field

Previously, when using the TCP/IP network transport with the Hot Standby feature, the DBO /SHOW STATISTICS “Hot Standby Statistics” display “State:” field could overwrite the “UserSync:” heading as in the following example:

```
Node: HSVMS (1/1/16) Oracle CODAYSL DBMS V7.1-24 Perf. Monitor 18-JUL-2006 06:17:59.97
Rate: 3.00 Seconds Hot Standby Statistics Elapsed: 00:07:28.63
Page: 1 of 1 $MYDISK:[MYDB]PARTS.R00;1 Mode: Online
-----
State: TCP/IP:72 rSync: Cold Current.Msg: 1 Cl Mstr.AIJ: 1:2
LagTime: 00:00:00 AutoSync: Cold Stalled.Msg: none 1 Stby.AIJ: 1:2
Stby.DB: $MYDISK:[MYDB_STANDBY]PARTS.
```

The line starting with “State:” partly overwrites “UserSync:”.

This problem has been corrected.

4.1.4 File-System Caching Avoided For DBO /COPY, /MOVE, /BACKUP/MULTI and /RESTORE/MULTI Access To Database

In order to reduce CPU consumption and XFC spinlock contention and to help avoid “thrashing” the file system cache and to streamline database file read and write operations during DBO /COPY, /MOVE, /BACKUP and /RESTORE functions, caching by the operating system is disabled when reading from or writing to the database files. There is no effect on caches implemented in storage devices or controllers.

Testing on various configurations indicates that, in general, avoiding the system’s XFC cache for these database operations results in better over-all performance as balanced between CPU and IO costs.

4.1.5 DBO/BACKUP/AFTER Ignores /EDIT_FILENAME When Backup Filespec Omitted

When a **DBO/BACKUP/AFTER** command was issued, if the **/EDIT_FILENAME** qualifier was included but no output filename was given, the default journal filename would be used and the contents of the **/EDIT_FILENAME** qualifier were ignored. For example:

```
$ DBO/BACKUP/AFTER/LOG -
  /EDIT=("_",VNO,"_",YEAR,MONTH,DAY_OF_MONTH,"_QP") -
  DBM$DATABASE ""
.
.
.
%DBO-I-LOGCREBCK, created backup file DEV:[DIR]JOURNAL_1.AIJ;1
```

In the above example, no output filename was specified, that is, "" was given as the output file. The journal that was being backed up had the filename “JOURNAL_1”. The backup filespec constructed by DBO should have been “JOURNAL_1_0_20060829.AIJ”, but the contents of the **/EDIT_FILENAME** qualifier were not incorporated in the output filename.

This problem can be avoided by explicitly providing the backup output filename in the backup command.

This problem has been corrected.

4.1.6 Processes Don't Always Terminate After Monitor Terminates

When the Oracle CODASYL DBMS monitor process terminates abnormally all user processes that are attached to databases on that node should immediately terminate. However, there were cases where that didn't happen, and those user processes would continue to access Oracle CODASYL DBMS resources after the monitor failed. Consider the following example.

User 1, node 1:

```
DBQ> BIND PARTS
```

User 2, node 2:

```
DBQ> BIND PARTS
```

User 3, node 1:

```
$ STOP/ID={pid of monitor process on node 1}
```

In the above sequence of events, the user process on node 1 should have terminated as soon as the monitor process was killed, but it remained active.

This problem can be avoided by using the **DBO/OPEN** command and manually opening databases on all nodes that will have users accessing the database.

This problem has been corrected.

4.1.7 Latch Hangs Possible From DBO /SHOW STATISTICS

Bugs 4397634 and 5842040

In prior release of Oracle CODASYL DBMS, it was possible in a very small timing window for processes running the DBO /SHOW STATISTICS command to become hung while manipulating "latches" during the database attach sequence. Depending on the exact timing and sequence of events, this process may block other users of the database.

This problem has been corrected.

4.1.8 LRS Shutdown Failure DBM-F-PARTDTXNERR/SYSTEM-F-NOSUCHID

Bug 5754461

A possible problem with the Oracle CODASYL DBMS Hot Standby feature has been identified. If the OpenVMS \$GETGTI system service returns a status value of SS\$_NOSUCHID, the LRS process could be unable shutdown cleanly. This could result in an inconsistent standby database.

This problem has been corrected. The LRS process now treats a returned SS\$_NOSUCHID status the same as a SS\$_NOSUCHTID status and will be handled normally and will not be cause the LRS to fail.