Oracle® Database

Readme

11*g* Release 2 (11.2)

E22488-04

May 2012

Purpose of this Readme

This Readme file is relevant only to the delivered Oracle Database 11*g* Release 2 (11.2.0.1, 11.2.0.2 and 11.2.0.3) product and its integral parts, such as SQL, PL/SQL, the Oracle Call Interface (OCI), SQL*Loader, Import/Export utilities, and so on.

This Readme documents differences between the server and its integral parts and its documented functionality, as well as known problems and workarounds.

A complete list of open known bugs is available on the product CD.

Operating system releases, such as Linux, UNIX and Windows, often provide readme documents specific to that operating system. Additional Readme files may also exist. This Readme file is provided instead of system bulletins or similar publications.

For licensing information, refer to the *Oracle Database Licensing Information*.

The major sections of this Readme file are:

- Readme Information for Oracle Database 11g Release 2 (11.2.0.1)
- Readme Information for Oracle Database 11g Release 2 (11.2.0.2)
- Readme Information for Oracle Database 11g Release 2 (11.2.0.3)
- Documentation Accessibility

1 Readme Information for Oracle Database 11*g* Release 2 (11.2.0.1)

> **Note:** If you are on Oracle Database 11*g* Release 2 (11.2.0.1), this is the Readme section that you need to read.

This section of the Readme contains the following sub-sections:

Section 1.1, "Nomenclature Changes"

Section 1.2, "Compatibility, Upgrading, Downgrading, and Installation"

Section 1.3, "Features Not Available or Restricted in 11.2.0.1"

Section 1.4, "Default Behavior Changes"

Section 1.5, "Oracle Automatic Storage Management (Oracle ASM)"

Section 1.6, "Oracle Enterprise Manager Database Control"

Section 1.7, "Database Security"

ORACLE®

1.1 Nomenclature Changes

Note the following nomenclature changes:

- Flash Recovery Area has been renamed to Fast Recovery Area.

- The name Oracle *inter*Media was changed to Oracle Multimedia in Oracle Database 11*g* Release 1 (11.1). The feature remains the same, only the name has changed.

1.2 Compatibility, Upgrading, Downgrading, and Installation

For late-breaking updates and best practices about preupgrade, post-upgrade, compatibility, and interoperability discussions, see Note 785351.1 on My Oracle Support (formerly Oracle*MetaLink* at https://metalink.oracle.com/) that links to the "Upgrade Companion" web site for Oracle Database 11*g* Release 2.

> **Caution:** After installation is complete, do not manually remove or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running. If you remove these files, then Oracle software can encounter intermittent hangs. Oracle Grid Infrastructure for a cluster and Oracle Restart installations fail with the following error:
>
> ```
> CRS-0184: Cannot communicate with the CRS daemon.
> ```

1.2.1 Deinstallation Restrictions

The following sections describe deinstallation and deconfiguration restrictions. See Section 1.38.2, "Deinstallation Tool Known Bugs" for additional information.

1.2.1.1 Deinstallation Using OUI Starting with Oracle Database 11*g* Release 2 (11.2), a separate deinstallation and deconfiguration tool ships with the product (as a separate download). Use the `deinstall` tool to deinstall and deconfigure the software. This tool can be used to deinstall and deconfigure Oracle Grid Infrastructure for a cluster homes, Oracle Real Application Clusters (Oracle RAC) database homes, single-instance databases, database clients, and Oracle Grid Infrastructure for standalone server (Oracle Restart) homes.

The `deinstall` tool is also installed in all Oracle homes. To use the tool, run it from the Oracle home. The `deinstall` tool has built-in intelligence to check installed software, and access the files that it needs to complete the deinstallation. If the tool detects missing files, it prompts you to download a standalone version of the `deinstall` tool to complete a deinstallation successfully.

Refer to the Readme that is included with the `deinstall` tool download for more information.

1.2.1.2 Error When Running Deinstallation from an Existing Oracle Home With the -home Option
If you try to run the `deinstall` tool from the product home with the `-home` option, then the deinstallation fails with the following error message:

```
$ ./deinstall -home /scratch/user_dir/oracle/product/11.2.0/dbhome_1
Error: invalid argument -home.
Since the tool is run from inside an Oracle Home it will deinstall the home
the tool is installed in. You cannot specify another home.
```

Because the `deinstall` tool is run from within an Oracle home, the `deinstall` tool cannot be run with the `-home` option. The `deinstall` tool can only be run as `./deinstall` from within an Oracle home.

1.2.1.3 Deinstall Upgraded 11.2 Oracle RAC and Oracle Grid Infrastructure for a Cluster Homes

After you deconfigure and deinstall an upgraded Oracle Database 11*g* Release 2 (11.2) Oracle RAC home and to deconfigure and deinstall an 11.2 Oracle Grid Infrastructure for a cluster home, you must detach any pre-11.2 Oracle RAC software homes from the central Inventory (reference Bug 8666509).

Detach the pre-11.2 Oracle RAC homes from the central inventory with the following command:

```
ORACLE_HOME/oui/bin/runInstaller -detachHome ORACLE_HOME_NAME=pre-11.2_ORACLE_
HOME_NAME ORACLE_HOME=pre-11.2_ORACLE_HOME
```

1.2.1.4 Delete /tmp/install Directory Before Running the Deinstallation Tool
If the `/tmp/install` directory exists prior to running the deinstallation and deconfiguration tool, remove the directory and rerun the tool again (reference Bug 8729651).

1.2.2 Time Zone File Version

Oracle Database 11*g* Release 2 (11.2) ships with time zone file versions 1 through 11. If your older database is using a time zone file version greater than 11, then you must obtain the corresponding time zone file patch for 11.2 prior to upgrading your database.

You can use `SELECT VERSION FROM V$TIMEZONE_FILE` to identify the time zone file version when upgrading your Oracle Database 10*g* or 11*g* databases. The Pre-Upgrade Information Tool (`utlu112i.sql` and its counterpart used by the Database Upgrade Assistant) provides a warning if the time zone file version is not version 11. The warning recommends running the `DBMS_DST` package to upgrade the database time zone version to the latest and to update `TIMESTAMP WITH TIME ZONE` data as well. The Pre-Upgrade Information Tool also populates three new database properties (`DST_PRIMARY_TT_VERSION`, `DST_SECONDARY_TT_VERSION`, and `DST_UPGRADE_STATE`) in `sys.database_properties` that are pertinent to the time zone version and its upgrade. `DST_PRIMARY_TT_VERSION` property records the actual time zone version in use. The other two database properties will be used when you upgrade your time zone version using `DBMS_DST` package.

Note that, in this release, you have the option to retain your current time zone version after migrating to 11.2. For example, whether your application uses any `TIMESTAMP WITH TIME ZONE` type, you do not need to run the `DBMS_DST` package to upgrade your time zone version to the latest one available.

1.2.3 Oracle ASM Rolling Upgrade

Oracle Automatic Storage Management (Oracle ASM) rolling upgrade check does not allow rolling upgrade to be done from 11.1.0.6 to any later release (reference Bug 6872001). The following message is reported in the alert log:

```
Rolling upgrade from 11.1.0.6 (instance instance-number) to 11.x.x.x is not
supported
```

`ORA-15156` is signalled by LMON which will then terminate the instance.

When trying to upgrade Oracle ASM from 11.1.0.6 to a later release of Oracle ASM, apply the patch for this bug to 11.1.0.6 instances before rolling upgrade starts. This patch can be applied to 11.1.0.6 instances in a rolling fashion.

After the patch has been installed, set the `ASMCA_ROLLING_UPGRADE` user environment variable to `true`. For example:

```
$ EXPORT ASMCA_ROLLING_UPGRADE=true
```

1.2.4 UTC Time Zone Error When Upgrading From 9.2 to 11.2

When running upgrade scripts from Oracle9*i* Database Release 2 (9.2) to Oracle Database 11*g* Release 2 (11.2), you may encounter the following error:

```
ORA-27465: invalid value UTC for attribute DEFAULT_TIMEZONE
```

This error occurs because the default time zone file for release 9.2 does not have the coordinated universal time (UTC) zone. This message is produced by Oracle Scheduler that, by default, selects the UTC time zone and checks it against the default time zone file of release 9.2. This error message is expected and you can ignore it.

1.2.5 Standard Edition Starter Database Upgrade

When the Standard Edition (SE) starter database is upgraded, the following components cannot be upgraded by the SE server because they require options that are not installed in the Standard Edition:

- OLAP Catalog

- OLAP Analytic Workspace

- Oracle OLAP API

- Oracle Spatial

After the upgrade, these components will have a `STATUS` value of `OPTION OFF` in the `DBA_REGISTRY` view, and there will be some invalid objects in the associated component schemas. The Database Upgrade Assistant (DBUA) will show unsuccessful upgrades for these components (reference Bug 8621666).

1.2.6 Tablespace and Fast Recovery Area Sizing

> **Note:** Fast Recovery was previously known as Flash Recovery.

The Oracle Database 11*g* Pre-Upgrade Information Utility (`utlu112i.sql`) estimates the additional space that is required in the `SYSTEM` tablespace and in any tablespaces associated with the components that are in the database (for example, `SYSAUX`, `DRSYS`). For a manual upgrade, be sure to run this utility on your existing database prior to upgrading.

The tablespace size estimates may be too small, especially if Oracle XML DB is installed in your database. However, to avoid potential space problems during either a manual upgrade or an upgrade using the Database Upgrade Assistant (DBUA), you can set one data file for each tablespace to `AUTOEXTEND ON MAXSIZE UNLIMITED` for the duration of the upgrade.

If you are using file systems for data file storage, then be sure there is adequate space in the file systems for tablespace growth during the upgrade.

If you are using a Fast Recovery Area, then check that the size available is sufficient for the redo generated during the upgrade. If the size is inadequate, then an `ORA-19815` error will be written to the alert log, and the upgrade will stop until additional space is made available.

1.2.7 Setting Memory Target at Instance Startup on Linux

Starting with Oracle Database 11*g* Release 1 (11.1), Oracle provides the option of automatically managing SGA and PGA with a combined `MEMORY_TARGET` parameter

without having to set SGA_TARGET and PGA_AGGREGATE_TARGET explicitly. This is supported on Linux, Windows, Solaris, HPUX, and AIX (reference Bug 7258378).

If you see the ORA-00845 error reported on Linux machines at Oracle instance startup when using the MEMORY_TARGET parameter, then check the size of /dev/shm. If /dev/shm is not configured, then mount it sized to be at least the value of MEMORY_TARGET. If /dev/shm is configured but the amount of available space reported (through df -k /dev/shm) is less than MEMORY_TARGET, then free the space or mount a larger /dev/shm to satisfy the MEMORY_TARGET size. Note that if you set the MEMORY_MAX_TARGET parameter greater than MEMORY_TARGET, then ensure that /dev/shm is sized to be at least the value of MEMORY_MAX_TARGET.

1.2.7.1 Memory Target for Oracle Database Instances Running Database Configuration Assistant (DBCA) defaults to this Automatic Memory Management option. In the case of upgrade or manual database creation, MEMORY_TARGET can be specified in the initialization parameter file.

1.2.8 Upgrading Databases with Oracle Multimedia, Oracle Spatial, or XDK for Java

Oracle Multimedia (formerly called Oracle *inter*Media), Oracle Spatial, and Oracle XDK for Java use Oracle XML DB. If any of these components are installed with the database, then Oracle XML DB is automatically installed to support them.

1.2.9 Upgrading a Database With Oracle Warehouse Builder (OWB)

If you are upgrading a database with OWB installed and configured, the OWB component will not be upgraded as part of the database upgrade process and therefore the version of OWB will remain the same after the database upgrade (reference Bug 9473944). The OWB component must be upgraded as a separate step as described in the *Oracle Warehouse Builder Installation and Administration Guide*.

1.2.10 Compatibility with Oracle9*i* Database Release 2

Connecting the Oracle Database 11*g* Release 2 (11.2) client or server to Oracle9*i* Database release (9.2.0.4 and above) is supported. Similarly, connecting Oracle9*i* client (release 9.2.0.4 and above) to Oracle Database 11*g* Release 2 (11.2) is supported.

1.2.11 Enabling Automated Backups

While installing Oracle Database, the Specify Backup and Recovery Options screen may appear truncated if your system does not have the required fonts installed. If your system has only fixed-width fonts, then you may not be able to fully specify the required information in the Backup Job Credentials area of the screen. To work around this issue, do not select **Enable Automated Backups** on this screen. After the installation is complete, use Oracle Enterprise Manager 11*g* Database Control to enable automated backups.

1.2.12 Upgrading a Database With SQL Access Advisor Tasks

Due to internal structural changes to the SQL Access Advisor repository, a database upgrade resets all of the existing SQL Access Advisor tasks to their initial state. This action effectively deletes all of the recommendation information for tasks that have been successfully completed prior to the upgrade.

After upgrading, the recommendation information can be restored by reexecuting the existing SQL Access Advisor tasks.

1.2.13 Downgrade to Release 11.1.0.6

If you anticipate downgrading back to release 11.1.0.6, then apply the patch for Bug 7634119. This action avoids the following DBMS_XS_DATA_SECURITY_EVENTS error:

```
PLS-00306: wrong number or types of arguments in call
to 'INVALIDATE_DSD_CACHE' DBMS_XS_DATA_SECURITY_EVENTS
PL/SQL: Statement ignored
```

Apply this patch prior to running catrelod.sql.

1.2.14 Downgrade to Release 10.2.0.4

If you anticipate downgrading back to release 10.2.0.4, then apply the patch for Bug 4309607 to the 10.2.0.4 Oracle home prior to running catrelod.sql. This patch is not needed for later 10.2.0.*x* patch releases. Applying this patch avoids the following error:

```
ORA-00600: internal error code, arguments: [koputilcvto2n], [15], [1035], [], [],
[], [], []
```

1.2.15 Downgrade from 11.2 to 10.2 Results in an Error

Downgrades from 11.2.0.1 to 10.2.0.4 can result in an ORA-600 [koputilcvto2n] error (reference BLR 8568714 and Bug 4309607).

To workaround this problem, apply patch 4309607 for 10.2.0.2 to the 10.2.0.2 Oracle home prior to downgrading from 11.2.0.1 to 10.2.0.2.

1.2.16 Rolling Upgrade for Oracle Clusterware

In Oracle Database 11*g* Release 2 (11.2), there is a new prerequisite check during the rolling Oracle Clusterware upgrade cycle. On failure of this prerequisite check, the *Oracle Database Upgrade Guide* documents the following:

> When upgrading Oracle Clusterware from Oracle Clusterware 10*g* Release 2 (10.2) to Oracle Clusterware 11*g* Release 2 (11.2), you must first apply the patchset version 10.2.0.3.0 or later to your existing Oracle Clusterware before it can be upgraded.

Follow the rolling upgrade procedure in Appendix F.5.1 of the *Oracle Grid Infrastructure Installation Guide for Linux*.

Apart from this proposed solution, there is an alternative. You can upgrade all of the nodes of the cluster simultaneously in a non-rolling manner. Choosing this solution voids the 10.2.0.3.0 patchset requirement.

1.2.17 Reusing the Oracle 9*i* Database Release 2 Listener Port on Release 11.2

If you reuse the same Oracle9*i* Database Release 2 (9.2) listener port for the SCAN VIP listener while installing Oracle Clusterware release 11.2 on nodes that have Oracle RAC release 9*i*, then you must ensure that your 9.2 listener is stopped (reference Bug 8688350).

Alternatively, you can stop the 9.2 listener, add IP=FIRST for the listener configuration in 9.2 listener.ora file, and restart the 9.2 listener before continuing with the Oracle Grid Infrastructure for a cluster installation.

1.2.18 Upgrading a Database With Oracle Label Security (OLS)

If you are upgrading a pre-11.2 database with OLS installed and configured, then you need to do one of the following for the upgrade of the database to succeed:

- Execute a custom installation of OLS in the 11.2 installation before starting the upgrade of the pre-11.2 database.

- If you do not want OLS in the upgraded database, then deinstall OLS before starting the upgrade.

If you do not execute one of the previously mentioned actions, then you will encounter an error (`ORA-01012: not logged on`) during the upgrade (reference Bug 8707506). In addition, after executing the `SHUTDOWN IMMEDIATE` command, the following error appears:

```
ORA-12432: LBAC error: zllesesinit:OCIStmtExecute
```

1.2.19 Upgrading With Objects in the Recycle Bin

The recycle bin must be empty during an upgrade to avoid possible `ORA-00600` deadlock errors, as well as to minimize the time required to perform the upgrade (reference Bug 8632581).

To avoid this deadlock, use the `PURGE DBA_RECYCLEBIN` statement to remove items and their associated objects from the recycle bin and release their storage space prior to upgrading your database.

1.2.20 Upgrading an 11.2 Database Where Oracle JVM Does Not Exist

During a database upgrade to 11.2, if Oracle JVM (which creates the PL/SQL package `DBMS_JAVA`) does not exist in the database, then the following error appears (reference Bug 8746395):

```
ERROR at line 1:
ORA-06550: line 1, column 7:
PLS-00201: identifier 'SYS.DBMS_JAVA' must be declared
ORA-06550: line 1, column 7:
PL/SQL: Statement ignored
```

This error can be safely ignored and the upgrade will continue.

1.2.21 Diagnosability Framework Errors After Downgrading from Release 11.2 to 11.1

After downgrading from Oracle Database release 11.2 to 11.1, the error `ORA-48318` may occur in the alert log, or when using the ADR Command Interpreter (ADRCI) utility (reference Bug 6976775). An Alert Log example follows:

```
Sweep Incident[8130]: failed, err=[48318]
```

The following is an ADRCI example:

```
adrci> show incident
DIA-48458: "show incident" failed due to the following errors
DIA-48318: ADR Relation [INCIDENT] of version=4 cannot be supported
```

As a workaround, perform the following steps:

1. Determine the location of the ADR home:

   ```
   SQL> select value from v$diag_info where name = 'ADR Home';

   VALUE
   --------------------------------------------------------------
   /ade/mfallen_g1/oracle/log/diag/rdbms/g1/g1
   ```

2. Stop the database instance.

3. Remove the ADR home directory using operating system utilities. (It is automatically re-created with the proper versions when the instance is restarted.)

1.2.22 Response File-Based Installation Issues

Note the following when doing a response file-based installation:

- While saving a response file, if a file with the specified response file name already exists at the destination location and the installation user does not have write permissions to overwrite the file contents, then the Installer does not prompt an error. Instead, the Installer silently fails as if it was successfully saved.

  As a workaround, always save the interview inputs (or response) to a new file. If the selected path already exists, then ensure that the installation user has sufficient privileges to overwrite the contents (reference Bug 8725384).

- The SELECTED_LANGUAGES property in the response file does not have a single value that stands for all of the languages. If you need to install the product in all of the languages, then put all of the language codes in a comma-separated list. An example is provided in the sample response file that is shipped with the product (reference Bug 8630967).

- The oracle.install.db.InstallEdition property in the response file must not be provided with a value of PE. This value is reserved for Windows operating systems (reference Bug 8631270).

- While saving a response file for client installation in custom mode, a few components, even though they are selected, are not recorded in the saved response file (reference Bug 8722858). Manually enter these components in the response file for these components. For example:

```
oracle.network.cman:11.2.0.1.0 -- "Oracle Connection Manager"
oracle.network.listener:11.2.0.1.0 -- "Oracle Net Listener"
```

1.3 Features Not Available or Restricted in 11.2.0.1

The following is a list of components that are not available or are restricted in Oracle Database 11*g* Release 2 (11.2.0.1):

- Database Smart Flash Cache is supported on Solaris and Oracle Linux only.

- Using Internet Protocol Version 6 (IPv6) is not supported with the following:
  - Oracle RAC and Oracle Clusterware
  - Oracle Fail Safe

- Oracle Ultra Search has been desupported and is not shipping with Oracle Database 11*g* Release 2 (11.2).

- Downgrading from Oracle Database 11*g* Release 2 (11.2) to release 10.2.0.3 or release 10.2.0.4 is not supported with Oracle Database Vault.

- Oracle Automatic Storage Management Cluster File System (Oracle ACFS) resources are not supported for Oracle Restart configurations on all platforms. Oracle ACFS drivers must be manually unloaded and loaded; Oracle ACFS file systems must be manually unmounted and mounted (after the Oracle ASM instance is running); Oracle ACFS database home file systems can be placed into the Oracle ACFS mount registry to be mounted along with other registered Oracle ACFS file systems.

- Refer to Section 1.10.1 for globalization restrictions within Oracle Secure Backup.

- Refer to Section 1.14 for features that are no longer available in Oracle Data Mining.

- Security-Enhanced Linux (SELinux) is not supported on Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems.

The following sections discuss additional restrictions.

1.3.1 Edition-Based Redefinition

The following restrictions exist for Oracle Database 11*g* Release 2 (11.2):

- If the AS CHILD OF clause is omitted when creating an edition, the new edition is created as the child of the database default edition. However, the *Oracle Database SQL Language Reference* manual specifies that the new edition is created as the child of the one edition that does not have a child (also known as the leaf edition). The *Oracle Database SQL Language Reference* manual is correct and the present behavior is incorrect (reference Bug 8681882).

- The CREATE EDITIONING VIEW command succeeds when the owner of the new editioning view is not editions-enabled. However the *Oracle Database SQL Language Reference* manual specifies that an editioning view must be owned by an editions-enabled user. The *Oracle Database SQL Language Reference* is correct and the present behavior is incorrect (reference Bug 8583698).

- If an updatable join view is defined on editioning views and the editioning views have triggers defined on them, then DML operations on the updatable join view may fail with various internal errors (reference Bug 8688904).

- DML on editioning views that are defined on tables that have an object type or nested table columns may result in external or internal (ORA-00600) errors (reference bug 7697126).

1.3.2 Database Object Names

Oracle does not recommend using quoted identifiers for database object names. These quoted identifiers are accepted by SQL*Plus, but they may not be valid when using other tools that manage database objects.

1.4 Default Behavior Changes

This section describes some of the differences in behavior between Oracle Database 11*g* Release 2 (11.2) and previous releases. The majority of the information about upgrading and downgrading is already included in the *Oracle Database Upgrade Guide*.

1.4.1 Direct Insert Behavior Change

Direct insert requires memory for every partition loaded. The memory usage is even greater if the partitions are compressed. In previous releases, a direct insert would continue to allocate memory as rows were encountered for previously untouched partitions until all of the partitions were loaded or until it ran out of memory. In this case, the insert fails.

Starting in Oracle Database 11*g* Release 2 (11.2), direct insert limits the memory that is allocated. If direct insert reaches the limit and it acquires rows for partitions that have not been loaded during the insert, then direct load stores those rows in the temporary tablespace. Once all rows have been passed in, direct load will load the rows it stored in the temporary tablespace.

1.4.2 Audit Default Behavior Changes

Default Audit behavior changes include:

- Audit filename is now prefixed with the instance name and ends with a sequence number. For example:

  instanceName_ora_pid_seqNumber.aud / *instanceName_ora_pid_seqNumber*.xml

- An existing audit file is never appended to. If an audit file already exists, the sequence number is incremented and written to *instanceName_ora_pid_ seqNumber+1*.aud.

- There is a preconfigured threshold for audit file growth. The audit file associated with an active session remains open until one of the following limits is reached:

  - After the audit record is written, the audit file size becomes 10 megabytes or more.

  - After the audit record is written, the audit file age becomes 5 days or more.

  Once one of these thresholds is reached, a new audit file with an incremented sequence number is opened for further audit records.

- There are no updates to AUD$ anymore.

  - All logoff (action# 101) audit records are written separately.

  - If an event is audited BY SESSION, then every occurrence of the event becomes a new audit record in AUD$.

1.4.3 FILE_ACCESS_ROLE Default Behavior Change

The default behavior of the CTX system parameter FILE_ACCESS_ROLE has changed (reference Bug 8360111). Customers with existing Oracle Text indexes that use the file or URL datastore must take action to continue to use the indexes without error. The changes are as follows:

- If FILE_ACCESS_ROLE is null (the default), then access is not allowed. By default, users who were previously able to create indexes of this type will not be able to create these indexes after the change.

- FILE_ACCESS_ROLE is now checked for index synchronization and document service operations. By default, users will not be able to synchronize indexes of this type or use document service calls such as ctx_doc.highlight who were allowed to prior to this change.

- Only SYS will be allowed to modify FILE_ACCESS_ROLE. Calling ctx_adm.set_ parameter (FILE_ACESS_ROLE, *role_name*) as a user other than SYS will now raise the new error:

  DRG-10764: only SYS can modify FILE_ACCESS_ROLE

- Users can set FILE_ACCESS_ROLE to PUBLIC to explicitly disable this check (which was the previous default behavior).

1.4.4 Support for Raw Devices

Oracle Universal Installer (OUI) and Database Configuration Assistant (DBCA) do not support raw devices (or block devices on Linux). However, command-line utilities such as SQL*Plus and CRSCTL do support raw or block devices.

1.4.5 Oracle Clusterware and Oracle ASM Installed Into Oracle Grid Infrastructure for a Cluster Home

Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) are installed into a single Oracle home called the Grid home. This installation is referred to as the Oracle Grid Infrastructure for a cluster installation.

When upgrading, you must upgrade both Oracle Clusterware and Oracle ASM at the same time.

1.4.6 Deprecated Parameters in the Oracle ASM Instance

The following initialization parameters are deprecated in the Oracle ASM instance:

- `CLUSTER_DATABASE` parameter

  If the `INSTANCE_TYPE` is Oracle ASM and the Oracle RAC option is turned on, then you do not have to specify the `CLUSTER_DATABASE` parameter. In this case, the `CLUSTER_DATABASE` parameter defaults to `TRUE`.

- `COMPATIBLE` parameter

  Do not set the `COMPATIBLE` parameter in an Oracle ASM instance. To advance the disk group compatibility, change the `COMPATIBLE.[RDBMS|ASM|ADVM]` attributes of the disk group.

1.4.7 Fixed Views

All parameter views (for example, `V$PARAMETER`) in the Oracle ASM instance only report parameters that are relevant to the Oracle ASM instance.

1.4.8 Offset for CLOB and NCLOB APIs

Starting with Oracle Database 10*g* Release 1 (10.1), APIs that write to a `CLOB` or `NCLOB` cause error `ORA-22831` when the offset specified for the beginning of the write is not on a character boundary of the existing `LOB` data.

`LOB` APIs use UCS2 codepoint semantics for the amount and offset parameters on `CLOB`s or `NCLOB`s when the database default or national character set is multibyte. The specified offset is not on a character boundary if it points to the low (second) surrogate of a surrogate pair. In such situations, error `ORA-22831` occurs and the data is not written. Thus, this prevents the corruption of the character in the target `LOB`.

Contact Oracle Support Services to configure the database so that it does not return `ORA-22831`.

1.4.9 Use of the Append Hint Runs Out of Memory When Loading Many Partitions

Use of direct-path `INSERT` to load a large number of partitions can exceed memory limits, especially when data compression is specified (reference Bug 6749894). Starting in 11.2, the number of partitions loaded at the same time will be limited, based on the `PGA_AGGREGATE_TARGET` initialization parameter, to preserve memory. Rows that are not stored in the partitions that are currently being loaded are saved in the temporary tablespace. After all rows are loaded for the current set of partitions, other partitions are loaded from rows that are saved in the temporary tablespace.

This behavior helps prevent the direct-path `INSERT` from terminating because of insufficient memory.

1.4.10 Non-Uniform Memory Access Optimizations and Support Disabled in 11.2

With Oracle Database 11*g* Release 2 (11.2), non-uniform memory access support is disabled by default. This restriction applies to all platforms and operating systems (reference Bug 8450932).

Non-uniform memory access optimizations and support in the Oracle Database are only available for specific combinations of Oracle version, operating systems, and platforms. Work with Oracle Support Services and your hardware vendor to enable non-uniform memory access support.

1.4.11 View Changes for Advanced Compression

The `COMPRESS_FOR` column in various table views, such as `user_tables` and `dba_tables`, returns different values in 11.2 as compared to 11.1. The new values returned from a `COMPRESS_FOR` column are `BASIC` or `OLTP`. In 11.1, the value was `DIRECT LOAD ONLY` and `FOR ALL OPERATIONS`, respectively.

1.5 Oracle Automatic Storage Management (Oracle ASM)

The following sections describe information pertinent to Oracle Automatic Storage Management (Oracle ASM) in Oracle Database 11*g* Release 2 (11.2).

1.5.1 Oracle Database Data Files and Oracle ACFS

Oracle ACFS is the preferred file manager for non-database files. It is optimized for general purpose files. Oracle ACFS does not support any file that can be directly stored in Oracle ASM.

1.6 Oracle Enterprise Manager Database Control

The following sections describe information for Database Control in Oracle Database 11*g* Release 2 (11.2).

1.6.1 Reorganize Objects

The following applies to Reorganize Objects:

- The database can be adversely affected if metadata changes occur between the time the reorganization script is generated and the time it is completed.

- The database can be adversely affected if an offline reorganization is attempted while DDL is in progress against the tables being reorganized.

- Before using the Reorganize Objects wizard or the Make Tablespace Locally Managed wizard in Oracle Enterprise Manager, Oracle recommends that you back up your database.

1.6.2 Editing Multiple Objects

Oracle Enterprise Manager only supports having one browser window open for editing an object. For example, Oracle Enterprise Manager only supports editing one tablespace at a time.

> **Note:**   If a window is read-only, then you may have multiple browser windows open.

1.6.3 Requirements for Proxy Settings for Oracle XML DB

The following are requirements for Oracle XML DB proxy settings:

- For an Oracle XML DB Resource to be created using a URL behind the firewall, the proxy in `emoms.properties` needs to be set.

- When registering an XML Schema based on a URL, the URL is interpreted by the database itself, in which case, the database proxy might need to be set.

> **Note:** If the proxy settings cannot be changed, then a valid workaround is to save these files locally in the client machine and then use the "Local File" option to create a resource or register a schema.

1.6.4 Large Number of Database Objects May Require Increase in Heap Size

Oracle Enterprise Manager Database Control is configured with 192 MB (32-Bit)/384 MB (64-Bit) of heap memory. However, certain Oracle Enterprise Manager Database Control functionality (for example, Change Manager) may require higher memory settings if the database contains a large number of objects.

Memory settings can be increased using the following `emctl` command:

```
emctl config dbconsole -heap_size MemorySizeValue M
```

Oracle Enterprise Manager Database Control must be restarted for the new settings to take effect.

1.6.5 New Oracle Database 11*g* Release 2 Features Supported by Oracle Enterprise Manager Database Control 11.2 Only

Management support for new features available with Oracle Database 11*g* Release 2 (11.2.0.1), such as policy-managed databases and Oracle ASM Cluster File System, is available in Oracle Enterprise Manager Database Control 11.2 only.

Oracle Database 11g Release 2 (11.2.0.1) features are supported in Oracle Enterprise Manager Grid Control 11.1.

Oracle Enterprise Manager Grid Control 10.2.0.5, which provides management support for previous Oracle Database versions, does not support new Oracle Database 11*g* Release 2 (11.2) features. Single-instance database provisioning is also not supported by Grid Control 10.2.0.5.

Note that administrator-managed clustered databases that are upgraded to Oracle Database 11*g* Release 2 (11.2) can continue to be monitored using Grid Control 10.2.0.5. For more information on Oracle Database 11*g* Release 2 (11.2) monitoring support, click the **Certify** tab on the My Oracle Support web site (formerly Oracle*MetaLink*) at the following URL:

https://metalink.oracle.com

1.6.6 Database Control Patching

The following items apply to database patching procedures.

1.6.6.1 "Patch Oracle Clusterware" Deployment Procedure May Fail if Cluster Nodes are at Different Patching Levels
If the nodes of a cluster are at different patching levels, the "Patch Oracle Clusterware - Rolling" deployment procedure may succeed on some nodes, but fail during the "Conflict Analysis" step for others (reference Bug 8661258).

To resolve this issue, manually roll back the conflicting patch and retry the "Conflict Analysis" step. As an alternative, you can obtain a superset patch from Oracle Support Services and try the deployment procedure again.

1.6.6.2 SQL Scripts Must be Run Manually on Oracle RAC Instances During Patching A SQL script that runs on Oracle RAC instances as part of a patching procedure will only execute on the Oracle RAC instance on which Oracle Enterprise Manager Database Control is running (reference Bug 8645179).

Specifically, if the Oracle home being patched has multiple Oracle RAC instances running with it, and the patch being applied requires that SQL scripts be run on the instances, then "Patch Oracle RAC - Rolling" and "Patch Oracle RAC - All Nodes" Deployment Procedures will execute the SQL only on the Oracle Enterprise Manager Database Control Oracle RAC instance.

The SQL scripts must be run manually on other Oracle RAC instances by following the instructions in the patch README.

1.6.6.3 OPatch Must be Manually Updated for Patches Requiring the Latest OPatch Version Some patches require that the latest version of the OPatch utility be installed in the Oracle homes where they will be applied (reference Bug 8581434). However, OPatch upgrade is not supported in any of the Patching Deployment procedures.

Instead, you must manually download the latest OPatch version for the required platform from My Oracle Support (formerly Oracle*MetaLink* at https://metalink.oracle.com/) using patch number 6880880. You can then update each Oracle home with the latest version.

1.6.6.4 Standalone Listeners Must be Stopped Before Patching If a standalone listener is running in an Oracle home that is being patched with either the "Patch Oracle RAC - Rolling" or the "Patch Oracle RAC - All Nodes" Deployment Procedures, then patching fails with the following error:

```
CheckActiveFilesAndExecutables failed as ORACLE_HOME/bin/tnslsnr file is active.
```

To avoid this issue, stop all of the standalone listeners that are running from within the Oracle homes being patched before you begin patching (reference Bug 8581327).

1.6.6.5 Stop Multiple Database Instances Running in an Oracle Home Before Patching If more than one database instance is running within an Oracle home while being patched, then the "Patch Oracle Database" Deployment Procedure may fail in the "CheckActiveExecutables" step (reference Bug 6278749).

To avoid this issue, stop all database instances other than the one in which Database Control is running before patching.

1.6.6.6 Clusterware Bundle Patches Must be Applied to Oracle RAC Using OPatch If a Clusterware bundle patch contains multiple patches to be applied to an Oracle RAC database, then Database Control cannot be used to apply the patch (reference Bug 8692833). If it is, then the database might not start after patching.

Check the Oracle Clusterware bundle patch Readme to see if the patch contains multiple Oracle RAC patches. If it does, then the patch must be applied manually using OPatch.

1.6.6.7 Patching From Database Control Not Supported If Oracle RAC Does Not Span All Cluster Nodes If there is no Oracle RAC instance within a cluster that spans across all of the nodes within the cluster, and a patch is applied through Database Control, then those nodes not spanned by the Oracle RAC instance will not be patched (reference Bug 8752959).

In this scenario, the cluster must be patched manually by following the instructions in the patch Readme.

> **Note:** Oracle Clusterware patching should only be done from Database Control if Oracle RAC spans across all the nodes in a cluster. Patching of single-cluster installations from Database Control is not supported.

1.7 Database Security

Note the following changes in Database Security.

1.7.1 Changes in Default Security Settings

Java package `oracle.security.rdbms.server.AppCtx` is deprecated. Classes `AppCtxManager`, `AppCtxPermit`, `AppCtxUtil`, and `AppCtxException` are part of this package (reference Bug 6736417).

1.7.2 UNLIMITED TABLESPACE Privilege Changes

The `UNLIMITED TABLESPACE` system privilege will be removed from the `RESOURCE` role in a future Oracle Database release (reference Bug 7614645).

1.8 Encryption and Integrity

Note the following changes in the areas of encryption and integrity.

1.8.1 Encrypted Tablespaces

- You cannot encrypt an existing tablespace with an `ALTER TABLESPACE` statement. However, you can use Data Pump or SQL statements such as `CREATE TABLE AS SELECT` or `ALTER TABLE MOVE` to move existing table data into an encrypted tablespace.

- When recovering a database with encrypted tablespaces (for example, after a `SHUTDOWN ABORT` or a catastrophic error that brings down the database instance), you must open the wallet after database mount and before database open so the recovery process can decrypt data blocks and redo.

- The master encryption key for Transparent Data Encryption (TDE) tablespace encryption can leverage Hardware Security Modules (HSM) for full key management (for example, create, store, rotate, and retire). In 11.1.0.7, it was only possible to create and store the TDE tablespace encryption master key in HSM, not rotate. Master key for the TDE tablespace encryption can also be migrated from Oracle Wallet to HSM.

1.9 Java and Web Services

Note the following items when working with Java.

1.9.1 Oracle JVM

Oracle Database 11*g* Release 2 (11.2) includes a fully functional Java Virtual Machine (JVM), as well as the Java class libraries for Sun's Java Development Kit (JDK) 5.0. When combined with Oracle's JDBC and SQLJ, this release provides an enterprise class platform, Oracle JVM, for developing and deploying server-based Java applications. Refer to the Oracle JVM Readme file located at:

ORACLE_HOME/relnotes/readmes/README_javavm.txt

1.9.2 Announcing Oracle Universal Connection Pool (UCP) and Deprecation of JDBC Implicit Connection Cache in Favor of UCP

Starting with Oracle Database 11*g* Release 1 (11.1.0.7), Oracle has released the new Universal Connection Pool for JDBC. For more details, see the *Oracle Universal Connection Pool for JDBC Developer's Guide* at the following Web page:

http://www.oracle.com/technology/tech/java/sqlj_jdbc/UCP_dev_guide.pdf

Consequently, Oracle is deprecating the existing JDBC connection pool (that is, Implicit Connection Cache) that was introduced in Oracle Database 10*g* Release 1. Your applications will continue to work until formal desupport in a future release at which time a desupport notice will be posted.

Oracle encourages you to plan to adopt UCP for new applications and plan to change existing applications as indicated in *Transitioning to Oracle Universal Connection Pool (UCP)* at the following Web page:

http://www.oracle.com/technology/tech/java/sqlj_jdbc/pdf/ucp_transition_guide.pdf

UCP download and code samples are located at the following Web page:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/htdocs/ucp.html

1.9.3 JDBC

The Oracle JDBC product supports the latest Java/JDBC standards. For more details, refer to the JDBC Readme file located at:

ORACLE_HOME/relnotes/readmes/README_jdbc.txt

1.9.4 Web Services

As an alternative to Oracle Net, Oracle Database Web services provides nonconnected access to the database through standard Web services mechanisms. These include XML, SOAP, and WSDL, and can turn the database into a Web services provider. Similarly, the database itself can act as a Web service consumer and run external Web services. Important features of Web services include:

- A JAX-RPC based SOAP Client library supports invocation of external Web services from within the database, and applies the power of SQL to the results.

- **Web Services Call-In:** Deploying a JPublisher-generated Java class against Oracle Application Server 10*g* enables you to run database operations such as Java and PL/SQL procedures and packages, SQL queries, and DML operations.

- **Web Services Call-Out:** Deploying a JPublisher-generated Web services client from a WSDL and its PL/SQL wrapper supports invocation of external Web services from within the database.

1.10 Media Management Software

For environments that consist of a single server, Oracle offers Oracle Secure Backup Express to back up your Oracle Database and other critical Oracle infrastructure to tape. Oracle Secure Backup is fully integrated with Recovery Manager (RMAN) to provide data protection services. For larger environments, Oracle Secure Backup is available as a separately licensable product to back up many database servers and file systems to tape. Oracle Secure Backup release 10.3 is shipping with this Oracle Database 11*g* Release 2 (11.2). For more information on Oracle Secure Backup, refer to:

http://www.oracle.com/goto/osb/

1.10.1 Globalization Restrictions Within Oracle Secure Backup

The following globalization restrictions apply to Oracle Secure Backup:

- The Oracle Secure Backup Web Tool and command line interface are available in English only, and are not globalized. All messages and documentation are in English.

- Oracle Secure Backup does not support file names or RMAN backup names that are encoded in character sets that do not support null byte termination, such as Unicode UTF-16. Note that this restriction affects file names, not backup contents. Oracle Secure Backup can back up Oracle databases in any character set.

1.11 Oracle Application Express

Note the following items when working with Oracle Application Express.

To learn more about Oracle Application Express, refer to the *Oracle Application Express Release Notes* and the *Oracle Application Express Installation Guide.*

1.11.1 Checking for the Most Current Release

Oracle Application Express is released more frequently than Oracle Database. To view additional information about the most current release, refer to:

http://www.oracle.com/technetwork/developer-tools/apex/overview/index.html/

1.11.2 Configuring Oracle HTTP Server

When configuring Oracle HTTP Server for Oracle Application Express in a new installation, the database user APEX_PUBLIC_USER must be an unlocked account. To unlock the account for database user APEX_PUBLIC_USER, complete the following:

1. Start SQL*Plus and connect as SYS to the database where Oracle Application Express is installed. For example:

   ```
   $ ORACLE_HOME/bin/sqlplus
   SQL> CONNECT SYS as SYSDBA
   Enter password: SYS_password
   ```

2. Run the following command:

   ```
   SQL> ALTER USER APEX_PUBLIC_USER ACCOUNT UNLOCK
   ```

1.11.3 Database Compatible Parameter and Enabling Network Services

To run the examples in the Oracle Application Express Post-Installation tasks in the section titled "Enabling Network Services in Oracle Database 11*g*" in the *Oracle Database Installation Guide*, the COMPATIBLE initialization parameter of the database

must be set to at least 11.1.0.0.0. By default, an Oracle Database 11*g* database will already have the parameter set properly, but a database upgraded to 11*g* from a prior version may not.

Refer to the section titled "Creating and Configuring an Oracle Database" in the *Oracle Database Administrator's Guide* for information about changing database initialization parameters.

1.12 Oracle Client Applications

Oracle Client 11*g* contains advanced features for diagnosing issues, including the ability to dump diagnostic information when important errors are detected. By default, these dumps are restricted to a small subset of available information to ensure that application data is not dumped. However, in many installations, secure locations for dump files may be configured, ensuring the privacy of such logs. In such cases, it is recommended that you turn on full dumps; this can greatly speed the resolution of issues. Enable full dumps by adding the following line to the `sqlnet.ora` file that is used by your Oracle client installation:

```
DIAG_RESTRICTED=false
```

To verify that diagnosability features are working correctly, take the following steps:

1. Upgrade your application to use Oracle Database 11*g* client libraries.

2. Start your application.

3. Check the file `sqlnet.log` in your application's `TNS_ADMIN` directory for error messages indicating that diagnosability could not be started (normally this is due to invalid directory names or permissions).

Refer to the *Oracle Call Interface Programmer's Guide* for details.

1.13 Oracle Configuration Manager

Note the following for Oracle Configuration Manager.

1.13.1 cron Configuration Issue

If you are denied access to `cron`, then the configuration of Oracle Configuration Manager fails with the following error:

```
ORACLE_HOME/ccr/bin/setupCCR
** Installing base package **
Deploying core - Version 10.2.5.0.0
Error encountered in package deployment.
```

After the installation is complete, set the environment variable `CCR_DISABLE_CRON_ENTRY` to `TRUE` and try the configuration of Oracle Configuration Manager again using the following command:

```
ORACLE_HOME/ccr/bin/setupCCR
```

1.14 Oracle Data Mining

Oracle Data Mining scoring functions in Oracle Database 11*g* Release 2 are also available in Oracle Exadata Storage Server Software. Scoring capabilities in the storage layer permit very large data sets to be mined quickly, thus further increasing the competitive advantage already gained from Oracle in-database analytics. For

information about Oracle Exadata Storage Server Software, see
http://www.oracle.com/technology/products/bi/db/exadata/index.html.

The Data Mining Option, as an embedded feature of the database, is automatically installed with the Oracle Enterprise Edition Database. When installing the database with the Data Mining Option, choose the Data Warehouse configuration type for the most appropriate default initialization parameters.

In Oracle Database 11*g*, Data Mining models are implemented as data dictionary objects in the SYS schema. The DMSYS schema no longer exists.

Data Mining users must have the CREATE MINING MODEL privilege to create mining models in their own schema. Additional privileges are required for other data mining activities, as described in the *Oracle Data Mining Administrator's Guide*.

New data dictionary views for Oracle Data Mining were introduced in Oracle Database 11*g* Release 1 (11.1):

- USER/ALL/DBA_MINING_MODELS

- USER/ALL/DBA_MINING_MODEL_ATTRIBUTES

- USER/ALL/DBA_MINING_MODEL_SETTINGS

Demo programs that illustrate the Data Mining APIs (PL/SQL and Java) are installed with Oracle Database Examples. Instructions are in the *Oracle Data Mining Administrator's Guide*.

The Oracle Data Mining Scoring Engine Option, a separately installed database option in Oracle Database 10*g*, is not available in Oracle Database 11*g*. All functionality of the Data Mining Scoring Engine Option is offered in the Data Mining Option.

The Basic Local Alignment Search Tool (BLAST), previously supported by Oracle Data Mining, is not available in Oracle 11*g*.

1.15 Oracle Internet Directory

The Oracle Internet Directory product ships only with Oracle Application Server, not the Oracle Database 11*g* Release 1 (11.1) product set. The following information is included because Oracle Network functionality may use Oracle Internet Directory. Many of the administrative activities for Oracle Internet Directory have been consolidated into a single tool, Oracle Internet Directory Configuration Assistant (OIDCA). OIDCA should be used with Enterprise User Security and Network Names features under these conditions:

1. Enterprise User Security

   - Enterprise User Security only works with Identity Management Realms in this release. You must convert Oracle Contexts used in prior releases to Identity Management Realms using the OIDCA tool.

   - Use OIDCA when creating or updating the ldap.ora configuration file for discovering the Oracle Internet Directory server in the environment. When created with OIDCA, ldap.ora is located in the *ORACLE_HOME*/ldap/admin directory on Linux and UNIX operating systems and in the *ORACLE_HOME*\ldap\admin directory on Windows operating systems.

2. Network Names

   - Use OIDCA when creating, upgrading, and deleting Oracle Contexts.

- Use OIDCA when converting an Oracle Context from an earlier release to an Identity Management Realm.

- Use OIDCA when setting up the `ldap.ora` configuration file for discovering the Oracle Internet Directory server in the environment.

Note the following items when working with Oracle Internet Directory.

1.15.1 Using the Oracle Internet Directory Configuration Assistant

The Oracle Internet Directory Configuration Assistant (OIDCA) enables you to create, upgrade, and delete an Oracle Context, configure the file `ldap.ora`, and convert an Oracle Context to an Identity Management Realm.

The OIDCA syntax is:

```
oidca oidhost=host
      nonsslport=port |
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      propfile=properties file
```

To see the usage of OIDCA, enter `oidca -help` at the command prompt.

1.15.2 Creating an Oracle Context

The following syntax is used to create an Oracle Context in OIDCA; the parameters are described in the subsequent table.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      mode=CREATECTX
      contextdn=OracleContext DN
```

Parameters	Description
oidhost	OID server host; if not specified, default is `localhost`
nonsslport	OID server port; if not specified, default is `389`
sslport	OID SSL port; if not specified, default is `636`
dn	OID user, such as `cn=orcladmin`
pwd	OID user password
mode	Mode of the OIDCA; set to `CREATECTX`
contextdn	DN under which `OracleContext` must be created, such as `dc=acme, dc=com`

Note the following points:

- The `contextdn` must exist for this operation to be successful.

- This valid DN should not exist in OID: `"cn=oraclecontext,dc=acme, dc=com"`.

- This valid DN must exist in OID: `"dc=acme,dc=com"`.

- The parameters `mode` and `contextdn` can also be passed as a properties file.

- Specify the parameter `nonsslport=port` if you want to perform the operation using non-SSL mode.

- Specify the parameter `sslport=sslport` if you want to perform the operation using SSL mode.

- Either the `nonsslport` or the `sslport` parameter must be specified, but not both.

Functionality

1. The OIDCA verifies that `contextdn` has a valid DN syntax and that the entry exists in Oracle Internet Directory. Note that the OIDCA cannot create a root `OracleContext` explicitly. If there is no root Oracle Context, then OIDCA exits with an error.

2. If DN exists, then OIDCA verifies that the Oracle Context already exists.

   - If the Oracle Context already exists and is up-to-date, then OIDCA exits with the message `Oracle Context already exists and is up to date`.

   - If the Oracle Context already exists, but it is an older version, then OIDCA exits with the message `Oracle Context already exists and is of an older version`.

   - If the Oracle Context does not exist, then OIDCA creates the Oracle Context under this DN.

1.15.3 Upgrading an Oracle Context

To upgrade an `OracleContext` instance, use the following syntax; the parameters are listed in the subsequent table.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      mode=UPGRADECTX
      contextdn=OracleContext DN
```

Parameters	Description
oidhost	OID server host; if not specified, default is `localhost`
nonsslport	OID server port; if not specified, default is `389`
sslport	OID SSL port; if not specified, default is `636`
dn	OID user, such as `cn=orcladmin`
pwd	OID user password
mode	Mode of the OIDCA; always set to `UPGRADECTX`
contextdn	DN under which `OracleContext` must be created, such as `dc=acme, dc=com`

Note the following points:

- The `contextdn` must contain an `OracleContext` for this operation to be successful.

- The DNs `"cn=oraclecontext,dc=acme,dc=com"` and `"dc=acme,dc=com"` are both valid.

- The parameters `mode` and `contextdn` can also be passed as a properties file.

- Specify the parameter `nonsslport=port` if you want to perform the operation using a non-SSL mode.

- Specify the parameter `sslport=sslport` if you want to perform the operation using SSL mode.

- Either the `nonsslport` or the `sslport` parameter must be specified, but not both.

Functionality

1. OIDCA verifies that the `contextdn` has valid DN syntax and that `OracleContext` exists in Oracle Internet Directory. OIDCA cannot upgrade a root `OracleContext` explicitly. If there is no root `OracleContext`, then OIDCA sends an error message.

2. If `OracleContext` exists under `contextdn`,

   - The OIDCA checks if the `OracleContext` belongs to a realm, in which case it exits with the appropriate message. Note that `OracleContext` instances that belong to a realm cannot be upgraded.

   - The OIDCA verifies that the `OracleContext` is up-to-date, then exits with the message `Oracle Context already exists and is up to date`.

   - If the `OracleContext` is not up-to-date, then the OIDCA upgrades the `OracleContext` under this DN.

1.15.4 Deleting an Oracle Context

To delete an `OracleContext`, use the following syntax; the parameters are listed in the subsequent table.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      mode=DELETECTX
      contextdn=OracleContext DN
```

Parameters	Description
oidhost	OID server host; if not specified, default is `localhost`
nonsslport	OID server port; if not specified, default is `389`
sslport	OID SSL port; if not specified, default is `636`
dn	OID user, such as `cn=orcladmin`
pwd	OID user password
mode	Mode of the OIDCA; always set to `DELETECTX`
contextdn	DN under which OracleContext must be created, such as `dc=acme,` `dc=com`

Note the following points:

- The `contextdn` must contain an `OracleContext` for this operation to be successful.

- The DNs `"cn=oraclecontext, dc=acme,dc=com"` and `"dc=acme,dc=com"` are both valid.

- The parameters mode and contextdn can also be passed as a properties file.

- Specify the parameter nonsslport=port if you want to perform the operation using a non-SSL mode.

- Specify the parameter sslport=sslport if you want to perform the operation using SSL mode.

- Either the nonsslport or the sslport parameter must be specified, but not both.

Functionality

1. OIDCA verifies that the contextdn has valid DN syntax and that OracleContext exists in Oracle Internet Directory.

2. If OracleContext exists under contextdn,

   - The OIDCA checks if the OracleContext belongs to a realm, in which case it exits with the appropriate message. Note that OracleContext instances that belong to a realm cannot be deleted.

   - If OracleContext does not belong to a realm, then OIDCA deletes it.

1.15.5 Configuring the File ldap.ora

To configure the file ldap.ora, use the following syntax; the parameters are listed in the subsequent table.

```
oidca oidhost=host
     nonsslport=port
     sslport=SSL Port
     adminctx=Administrative context
     mode=LDAPORA
     dirtype=OID or AD
     -update
```

Parameters	Description
oidhost	OID server host; if not specified, default is localhost.
nonsslport	OID server port; determined using discovery APIs.
sslport	OID SSL port; determined using discovery APIs.
mode	Mode of the OIDCA; always set to LDAPORA.
dirtype	Directory type; possible values are OID and AD; mandatory attribute.
adminctx	Default administrative context, such as dc=acme,dc=com. If not specified, then determined using discovery.
-update	If this flag is specified, then overwrite existing ldap.ora; if not, then create ldap.ora only if it does not already exist.

Note the following points:

- Either the non-SSL or the SSL port must be specified. The other port is discovered.

- The parameters mode, dirtype, and adminctx can also be passed in within a properties file.

Functionality

1. Using the Discovery API, the OIDCA determines all the parameters not specified on the command line.

2. The OIDCA checks for the `ldap.ora` location using Discovery APIs.

   - If `ldap.ora` exists and the `-update` parameter is not specified, then exit with message `ldap.ora exists`.

   - If `ldap.ora` exists and the `-update` parameter is not specified, then update the existing `ldap.ora` using Discovery API.

   - If `ldap.ora` does not exist, then create a new `ldap.ora` file in a location in the following order:

     ```
     LDAP_ADMIN
     ORACLE_HOME/ldap/admin
     ```

1.15.6 Converting an Oracle Context to an Identity Management Realm

Oracle Database 10*g* entries must be stored in Oracle Internet Directory release 9.0.4 server. An Identity Management Realm release 9.0.4 is also required for Enterprise User Security, a feature of the Oracle Database 10*g*.

To convert an existing `OracleContext` to an Identity Management Realm, use the following syntax. The parameters are listed in the subsequent table. Note that the root of the `OracleContext` object is not converted.

```
oidca oidhost=host
     nonsslport=port
     sslport=SSL Port
     dn=binddn
     pwd=bindpwd
     mode=CTXTOIMR
     contextdn=OracleContext DN
```

Parameters	Description
oidhost	OID server host; default is `localhost`
nonsslport	OID server port; default is `389`
sslport	OID SSL port; default is `636`
dn	OID user, such as `cn=orcladmin`
pwd	OID user password
mode	Mode of the OIDCA; always set to `CTXTOIMR`
contextdn	DN under which `OracleContext` must be created, such as `dc=acme,dc=com`

Note the following points:

- The `OracleContext` must exist under the specified `contextdn`.

- The DNs `"cn=oraclecontext,dc=acme,dc=com"` and `"dc=acme, dc=com"` are both valid.

- The parameters `mode` and `contextdn` can also be passed in a properties file.

- Specify the parameter `nonsslport=port` if you want to perform the operation using a non-SSL mode.

- Specify the parameter `sslport=sslport` if you want to perform the operation using SSL mode.

- Either the `nonsslport` or the `sslport` parameter must be specified, but not both.

Functionality

1. The OIDCA checks if `contextdn` has valid DN syntax, and if it contains a valid `OracleContext`.

2. If `OracleContext` exists under `contextdn`,

   - The OIDCA checks if the `OracleContext` belongs to a realm. If it does, then it exits with an appropriate error message.

   - If `OracleContext` does not belong to a realm, then OIDCA upgrades it to the latest version, and converts it to a realm.

Note also:

- If the nickname attribute is not `cn`, then configure it as a user configuration attribute using the Oracle Internet Directory Self-Service Console.

- If you want to use the Oracle Internet Directory Self-Service Console to manage the users and groups in the converted realm, then you must set up the administrative privileges appropriately. For details, refer to the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

1.16 Oracle Multimedia

The name Oracle *inter*Media was changed to Oracle Multimedia in Oracle Database 11*g* Release 1 (11.1). The feature remains the same, only the name has changed. References to Oracle *inter*Media were replaced with Oracle Multimedia, however some references to Oracle *inter*Media or *inter*Media may still appear in graphical user interfaces, code examples, and related documents in the Oracle Database documentation library for 11*g* Release 2 (11.2).

With Oracle Database 11*g* Release 2 (11.2), the Oracle Multimedia DICOM feature has been made available in Oracle Database Standard Edition. It was previously only available in Oracle Database Enterprise Edition.

For additional information, refer to the Oracle Multimedia Readme file located at:

```
ORACLE_HOME/ord/im/admin/README.txt
```

1.17 Oracle Net Services

Oracle Database provides limited support for Internet Protocol Version 6 (IPv6) addressing and connectivity. For details, see *Oracle Database Net Services Administrator's Guide*.

1.18 Oracle Real Application Clusters

Note the following items when working with Oracle RAC.

1.18.1 Moving ORADISM to Local Nodes from NFS on an Oracle Home

If you install an Oracle RAC database into a shared Oracle home on an NFS device, then you must copy the ORADISM binary (`oradism`) into a local directory on each node (reference Bug 6522385 and 6525377).

It is possible to install Oracle Clusterware and Oracle RAC on NFS without encountering this problem. Refer to the *Oracle Grid Infrastructure Installation Guide for Linux* for more information.

To move `oradism`, take the following steps:

1. Copy the `ORACLE_HOME`/bin/oradism binary to an identical directory path on all cluster nodes. The path (for example, `/u01/local/bin` in the example in Step 2) must be local and not NFS. For example:

   ```
   cp -a ORACLE_HOME/bin/oradism/u01/local/bin
   ```

2. Run the following commands, as the root user, to set ownership and permissions of the `oradism` executable:

   ```
   $ chown root /u01/local/bin/oradism
   $ chmod 4750 /u01/local/bin/oradism
   ```

3. Create a symbolic link from the NFS shared home to the local `oradism` directory path. This needs to be done from one node only. Each node can then reference its own `oradism` using the `symlink` from the shared Oracle home. For example:

   ```
   $ cd /nfs/app/oracle/product/11.2.0/db_1/bin
   $ rm -f oradism
   $ ln -s /u01/local/bin/oradism oradism
   ```

4. If the Oracle home is an Oracle Database home directory, then repeat steps 1-3 for other binaries such as `extjob`, `jssu`, `nmb`, `nmhs` and `nmo`. You do not need to perform this step if the Oracle home is an Oracle Grid Infrastructure home directory.

1.18.2 Creating Server Pools

> **Caution:** By default, any named user may create a server pool. To restrict the operating system users that have this privilege, Oracle strongly recommends that you add specific users to the CRS Administrators list.
>
> See also *Oracle Clusterware Administration and Deployment Guide* for more information about adding users to the CRS Administrators list.

Create an initial, user-defined server pool immediately after the Oracle Grid Infrastructure for a cluster installation if you plan to use a policy-managed Oracle RAC database. As the Grid software owner, use the following SRVCTL command located in the path `Grid_home`/bin:

```
srvctl add srvpool -g srvpool_name -u max
```

1.18.3 Adding a Node in a Cluster

When adding a node in a cluster running a policy-managed database, Oracle Clusterware tries to start the new instance before the cloning procedure completes. The following steps should be used to add the node:

1. Run `addNode` for the Oracle Grid Infrastructure for a cluster to add the new node. Do not run the root scripts when prompted; you will run them later.

2. Install the Oracle RAC database software using a software-only installation. Ensure Oracle is linked with the Oracle RAC option.

3. Complete the root script actions for the Database home.

4. Complete the root scripts action for the Oracle Clusterware home and then finish the installation.

1.19 Oracle Grid Infrastructure for a Cluster

Note the following items when working with Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM), which are installed with an Oracle Grid Infrastructure for a cluster installation.

1.19.1 Configure and Use SSL Certificates to Setup Authentication

> **Note:** This affects the security in the connection between the Oracle Clusterware and the mid-tier or JDBC client.

JDBC or Oracle Universal Connection Pool's (UCP) Oracle RAC features like Fast Connection Failover (FCF) subscribe to notifications from the Oracle Notification Service (ONS) running on the Oracle RAC nodes. The connections between the ONS server in the database tier and the notification client in the mid-tier are usually not authenticated. It is possible to configure and use SSL certificates to setup the authentication but the steps are not clearly documented.

The workaround is as follows:

1. Create an Oracle Wallet to store the SSL certificate using the `orapki` interface:

   a. `cd $ORA_CRS_HOME/opmn/conf`

   b. `mkdir sslwallet`

   c. `orapki wallet create -wallet sslwallet -auto_login`

      When prompted, provide `ONS_Wallet` as the password.

   d. `orapki wallet add -wallet sslwallet -dn "CN=ons_test,C=US" -keysize 1024 -self_signed -validity 9999 -pwd ONS_Wallet`

   e. `orapki wallet export -wallet sslwallet -dn "CN=ons_test,C=US" -cert sslwallet/cert.txt -pwd ONS_Wallet`

2. Update the ONS configuration file on all nodes in the database tier to specify the location of the wallet created in Step 1:

   a. Open the file *ORA_CRS_HOME*/opmn/conf/ons.config

   b. Add the `walletfile` parameter to the `ons.config` file:

      `walletfile=`*ORA_CRS_HOME*`/opmn/conf/sslwallet`

   c. Restart the ONS servers with the `srvctl`:

      ```
      srvctl stop nodeapps
      srvctl start nodeapps
      ```

3. If you are running a client-side ONS daemon on the mid-tier, there are two possible configurations:

   - ONS started from OPMN (like in OracleAS 10.1.3.*x*) which uses `opmn.xml` for its configuration.

- ONS started standalone (like using `onsctl`), which uses `ons.config` for its configuration.

For case (1), refer to the *OPMN Administrator's Guide* for the Oracle Application Server release. This involves modifying the `opmn.xml` file to specify the wallet location.

For case (2), refer to the section titled *Configuration of ONS* in *Appendix B* of the *Oracle Database JDBC Developer's Guide* to modify the ons.config file to specify the wallet location created in Step 1.

4. If you are running remote ONS configuration without a client-side ONS daemon, refer to the "Remote ONS Subscription" subsection of the "Configuring ONS for Fast Connection Failover" subsection of the "Using Fast Connection Failover" section of the "Fast Connection Failover" chapter in the *Oracle Database JDBC Developer's Guide*.

Alternatively, you can specify the following string as the `setONSConfiguration` argument:

```
propertiesfile=location_of_a_Java_properties_file
```

The Java properties file should contain one or more of the ONS Java properties listed below, but at least the `oracle.ons.nodes` property. The values for these Java properties would be similar to those specified in the "Remote ONS Subscription" subsection previously noted in this step:

```
oracle.ons.nodes
oracle.ons.walletfile
oracle.ons.walletpassword
```

1.19.2 Changing Host Names

Avoid changing host names after you complete the Oracle Grid Infrastructure for a cluster installation. Nodes with changed host names must be deleted from the cluster and added back with the new name.

1.19.3 Changing NLS_LANG to Use a Different Character Set

The `NLS_LANG` parameter is now read from the environment and not from the Oracle Cluster Registry (OCR) (reference Bug 8258489). The following steps describe how to change `NLS_LANG`:

1. Run the following command on all nodes:

```
crsctl stop crs
```

2. Edit *Grid_home*/crs/install/s_crsconfig_*hostname*_env.txt on all nodes and set the `LANGUAGE_ID` parameter to the appropriate value. For example:

```
LANGUAGE_ID='JAPANESE_JAPAN.JA16EUC'
```

3. Run the following command on all nodes:

```
crsctl start crs
```

1.19.4 Requirements for Rolling Upgrade of Oracle Clusterware

When upgrading Cluster Ready Services (CRS) from release 10.2.0.1 or 10.2.0.2 to release 11.2, a rolling upgrade of CRS requires additional steps (reference Bug 5198903). Complete the following to successfully complete the rolling upgrade:

1. Upgrade Cluster Ready Services (that is, the CRS home) to release 10.2.0.3, or 10.2.0.2 with CRS Bundle Patch #2 (reference Bug 5256865).

2. Upgrade CRS from release 10.2.0.3 to release 11.2.

Note that each of the CRS upgrades in Steps 1 and 2 are rolling upgrades. Therefore, a cluster outage is not required for these upgrades. Also note that it is not possible to do an in-place upgrade to Oracle Clusterware 11*g* Release 2.

There is a similar requirement for the rolling upgrade of 10.1.0.2 to 11.2 (reference Bug 5860479). To successfully perform a rolling upgrade of CRS from 10.1.0.2, complete the following:

1. Upgrade CRS (that is, the CRS home) to release 10.1.0.5.

2. Upgrade CRS from release 10.1.0.5 to release 11.2.

The CRS upgrades in Steps 1 and 2 are rolling upgrades, therefore no cluster outage is required. This is a problem for Oracle RAC only.

1.19.5 Unavailable OCR Location Error

The following command does not work if one or more of the Oracle Cluster Registry (OCR) locations has been marked as `unavailable` (reference Bug 8608734):

```
ocrconfig -add new_ocr_location
```

The workaround is to remove the location marked as `unavailable` using the following command before adding the new OCR location:

```
ocrconfig -delete unavailable_ocr_location
```

1.19.6 Incorrect Status of the Remote Node Given

In some cases, CRSCTL commands with the `-n nodename ora -all` do not give the correct status of the remote node (reference Bug 8655571). Instead, you receive a message that the nodes are `unknown`. You can log into the node to run the command. The problem clears after a node reboot.

1.19.7 Oracle ACFS and Oracle Clusterware Stack Shut Down

When attempting to shut down Oracle Clusterware, the Oracle Clusterware stack may report that it did not successfully stop on selected nodes (reference Bug 8703150 and Bug 8651848). If the database home is on Oracle ACFS, then you may receive the following error:

```
CRS-5014: Agent orarootagent.bin timed out starting process acfsmount for action
```

This error can be ignored.

Alternatively, the Oracle Clusterware stack may report that it did not successfully stop on selected nodes due to the inability to shut down the Oracle ACFS resources. If this occurs, take the following steps:

- Ensure that all file system activity to Oracle ACFS mount points is quiesced by shutting down programs or processes and retry the shutdown.

- If the `ora.registry.acfs` resource check function times out, or the resource exhibits a state of `UNKNOWN` or `INTERMEDIATE`, then this may indicate an inability to access the Oracle Cluster Registry (OCR). The most common cause of this is a network failure. The commands `acfsutil registry` and `ocrcheck` may give you a

better indicator of the specific error. Clear this error and attempt to stop Oracle Clusterware again.

1.19.8 sudo Commands and Cluster Administration with Oracle Enterprise Manager Database Control

If you want to add or delete nodes from the cluster using Oracle Enterprise Manager Database Control then, to do so without being prompted to run separate scripts, you must set up passwordless `sudo` access for the Oracle Grid Infrastructure for a cluster installation software owner account (reference Bug 8489482).

If you do not want to set up passwordless `sudo`, then the Oracle Enterprise Manager Database Control deployment procedure provides instructions for you to log onto a terminal as `root` and run the commands that require root privileges. After you run these commands, you can continue running the deployment procedure in Oracle Enterprise Manager Database Control.

1.19.9 Cluster Ready Service May Fail When Accessing OCR

The Cluster Ready Service may fail in accessing the OCR if the underlying storage becomes corrupted (reference Bug 8704534).

This problem can only happen if the OCR is configured in only one location. This problem will not occur if there are two or more OCR locations configured.

If the OCR is configured in a normal or high redundancy disk group, you can recover from the corruption by taking either of the following steps:

- Execute the `ALTER DISKGROUP CHECK` statement if the disk group is already mounted, or

- Remount the disk group with the `FORCE` option, which also takes the disk offline when it detects the disk header corruption.

If you are using an external redundancy disk group or third-party storage, then you must restore the OCR from backup to recover from a corruption.

The workaround is to configure an additional OCR location on a different storage location using the `ocrconfig -add` command.

1.20 Oracle Real Application Testing

The following sections describe Oracle Real Application Testing issues.

1.20.1 Reading SQL Trace Files to Construct a SQL Tuning Set

Oracle Database supports the added functionality to read SQL trace files from Oracle Database 9*i* to construct a SQL tuning set that can be used as an input source for SQL Performance Analyzer. After being constructed, you can use SQL Performance Analyzer to execute the SQL tuning set on Oracle Database 10*g* Release 2 remotely over a database link. This functionality is provided so that you can use the Oracle Real Application Testing option to test the effect on SQL response time of a database upgrade from Oracle Database 9*i* to Oracle Database 10*g* Release 2. For details, refer to the *Oracle Database Real Application Testing Addendum*.

1.21 Oracle ODBC Driver

The Oracle ODBC Driver Readme file is located at:

ORACLE_HOME/odbc/html/ODBCRelnotesUS.htm

1.22 Oracle OLAP

Consider the following when working with Oracle OLAP:

1. If you are upgrading from a 32-bit version to 64-bit version, and you are using the OLAP option, see Note 352306.1 on My Oracle Support (formerly Oracle*MetaLink* at https://metalink.oracle.com/) for additional steps (reference Bug 4966492).

2. SQL Aggregation Management is a group of PL/SQL subprograms in `DBMS_CUBE` that support the rapid deployment of cube materialized views from existing relational materialized views. Cube materialized views are cubes that have been enhanced to use the automatic refresh and query rewrite features of Oracle Database. A single cube materialized view can replace many of the relational materialized views of summaries on a fact table, providing uniform response time to all summary data.

   In the process of creating the cube materialized views, `DBMS_CUBE` also creates a fully functional analytic workspace including a cube and the cube dimensions. The cube stores the data for a cube materialized view, instead of the table that stores the data for a relational materialized view. A cube can also support a wide range of analytic functions that enhance the database with information-rich content.

3. Enhanced logging capabilities are provided through the new PL/SQL package `DBMS_CUBE_LOG` that offers four different logs. Both Analytic Workspace Manager users and database administrators can view the progress of an analytic workspace build in `CUBE_BUILD_LOG`, find load errors in `CUBE_REJECTED_RECORDS`, check the validity of dimension hierarchies in `CUBE_DIMENSION_COMPILE`, and follow OLAP engine events in `CUBE_OPERATIONS_LOG`.

4. The OLAP Option, as an embedded feature of the database, is automatically installed with the Oracle Enterprise Edition Database. When installing the database with the OLAP Option, choose the Data Warehouse configuration type for the most appropriate default initialization parameters.

5. Analytic Workspace Manager 11.2.0 should be used when running Oracle Database 11*g* Release 2 in either Oracle Database 10*g* or Oracle Database 11*g* compatibility modes. An Oracle Database 10*g* format analytic workspace may be created by choosing the Oracle Database 10*g* cube type when creating a new analytic workspace.

6. OLAP 10*g* clients connecting to 10*g* style analytic workspaces in an Oracle Database 11*g* Release 2 database instance should upgrade their Java to version 1.5, as well as upgrade their OLAP API version to 11.2.0.

7. Oracle Business Intelligence Beans 10*g* and Oracle Discoverer for OLAP 10*g* can be used with 10*g* format analytic workspaces in Oracle Database 11*g*. They cannot be used with 11*g* format analytic workspaces. To upgrade the OLAP API client JARs, include these new JARs in the Java class search list before the Intelligence Beans or Discoverer JAR files or both.

8. The OLAPSYS schema and the CWM metadata contained within the schema are deprecated in Oracle Database 11*g* Release 2. The OLAP API can still query relational tables (ROLAP mode) using data dictionary metadata. For more information, refer to My Oracle Support Tech Note 445311.1.

1.23 Oracle Spatial

The Oracle Spatial readme file supplements the information in the following manuals: *Oracle Spatial Developer's Guide, Oracle Spatial Topology and Network Data Models*

Developer's Guide, and *Oracle Spatial GeoRaster Developer's Guide*. The Oracle Spatial readme file is located at:

ORACLE_HOME/md/doc/README.txt

1.24 Oracle SQL Developer

The Oracle SQL Developer readme file is located at:

ORACLE_HOME/sqldeveloper/readme.html

1.25 Oracle Text

Note the following items when working with Oracle Text. You should also check entries for the *Oracle Text Application Developer's Guide* in the Documentation Addendum.

1.25.1 Query Compatible Syntax Values

Chapter 6, titled "CTX_CLS Package", in the *Oracle Text Reference* states under the section titled "Query Compatible Syntax" that the doc_id column is a NUMBER. This is correct; however, the values stored in this column must be in the range 0-4294967295. The values must be stored in an unsigned 32-bit value. This range also pertains to catid, catdocid, and rescatid.

1.25.2 Oracle Text Supplied Knowledge Bases

An Oracle Text knowledge base is a hierarchical tree of concepts used for theme indexing, ABOUT queries, and deriving themes for document services. The following Oracle Text services require that a knowledge base be installed:

- Index creation using a BASIC_LEXER preference where INDEX_THEMES=YES

- SYNCing of an index where INDEX_THEMES=YES

- CTX_DOC.THEMEs

- CTX_DOC.POLICY_THEMEs

- CTX_DOC.GIST

- CTX_DOC.POLICY_GIST

- CTX_QUERY.HFEEDBACK

- CTX_QUERY.EXPLAIN, if using ABOUT or THEMES with TRANSFORM

- CTX_DOC.SNIPPET (if using the ABOUT operator)

- CTX_DOC.POLICY_SNIPPET (if using the ABOUT operator)

- CONTAINS queries that use ABOUT or THEMES with TRANSFORM

- The Knowledge Base Extension Compiler, ctxkbtc

- Clustering and classification services, if themes are specified

If you plan to use any of these Oracle Text features, then you should install the supplied knowledge bases, English and French, from the Oracle Database Examples media, available for download on OTN.

Note that you can extend the supplied knowledge bases, or create your own knowledge bases, possibly in languages other than English and French. For more

information about creating and extending knowledge bases, refer to the *Oracle Text Reference*.

For information about how to install products from the Oracle Database Examples media, refer to the *Oracle Database Examples Installation Guide* that is specific to your platform.

Supplied Knowledge Bases and Upgrades

Because the supplied knowledge bases are contained on the Oracle Database Examples media, they are not immediately available after an upgrade to Oracle Database 11*g* Release 1 (11.1). Oracle Text features that depend on the supplied knowledge bases available before the upgrade will not function after the upgrade, so you must install the supplied knowledge bases from the Oracle Database Examples media.

After an upgrade, you must regenerate all user extensions to the supplied knowledge bases. These changes affect all databases installed in the given `ORACLE_HOME`.

For more information on upgrading Oracle Text and supplied knowledge bases, refer to the *Oracle Database Upgrade Guide*, Chapter 4, "After Upgrading a Database", section "Upgrading Oracle Text". The *Oracle Text Application Developer's Guide* contains both general instructions for upgrading from previous releases of Oracle Text and information on supplied knowledge bases.

1.25.3 Document Filtering Support

Starting with version 11.1.0.7, Oracle Text uses Oracle Outside In HTML Export for document filtering instead of filtering technology licensed from Autonomy, Inc. With the migration to Oracle Outside In HTML Export, several document formats are no longer supported. Refer to Appendix B of the *Oracle Text Reference* for a complete list of filter-supported document formats as well as for a list of formats that are no longer supported in 11.1.0.7.

1.25.4 URL Datastore and File Datastore

With the default behavior change for `FILE_ACCESS_ROLE` (see Section 1.4.3, "FILE_ACCESS_ROLE Default Behavior Change"), indexing operations on new or existing Oracle Text index using `FILE` or `URL` datastore fail if `FILE_ACCESS_ROLE` parameter is set to `NULL` (reference Bug 8360111).

1.26 Oracle Ultra Search

Oracle Ultra Search has been de-supported. No new features are implemented in Ultra Search, and Ultra Search will no longer be included as part of future Oracle product sets. In March of 2006, Oracle launched Oracle Secure Enterprise Search (SES). Oracle SES is a faster and more secure product built on the technologies originally developed for Ultra Search. Customers using Ultra Search are strongly recommended to migrate to Oracle SES.

See Also: The Oracle Secure Enterprise Search page on the Oracle Technology Network at:

http://www.oracle.com/technology/products/oses/index.html

1.27 Oracle Warehouse Builder

For additional information about Oracle Warehouse Builder (OWB) in Oracle Database 11*g* Release 2 (11.2), refer to the *Oracle Warehouse Builder Release Notes.*

1.27.1 Upgrading Existing Warehouse Builder ETL Designs for Use With OWB 11.2

Oracle Warehouse Builder Installation and Administration Guide, Chapter 6, "Migrating Oracle Warehouse Builder from Earlier Releases," states that "To upgrade from versions of Oracle Warehouse Builder earlier than 10*g* Release 2, you must first upgrade to Oracle Warehouse Builder 10*g* Release 2 or Oracle Warehouse Builder 11*g* Release 1. You may then upgrade your repository to Oracle Warehouse Builder 11*g* Release 2."

Note, however, that the complete repository upgrade process is only needed to preserve runtime metadata, such as when ETL mappings and process flows were deployed and execution results.

If the goal is to move an OWB design from a previous release to OWB 11*g*R2 and the deployment and execution history from the previous repository is not needed, then it is not necessary to upgrade the entire repository. Specifically, when upgrading from OWB 10.1 or earlier, it is simpler to migrate only the design-time metadata.

You can migrate an existing design to OWB 11*g*R2 using the following steps:

1. Export each project from previous releases as MDL.

2. Import the MDL into an OWB 11*g*R2 workspace.

3. Upgrade locations to match the new environment.

1.28 Oracle Workflow

Oracle Workflow is not shipping with Oracle Database 11*g* Release 2 (11.2).

1.29 Oracle XML DB

Consider the following when working with Oracle XML DB:

- Compression is not supported with Oracle XML DB.

- The transportable tablespace feature is not supported on hierarchy-enabled XMLType tables.

1.29.1 Change in Semantics of xdb:defaultTable Annotation

There is a change in behavior in the semantics of `xdb:defaultTable` annotation while registering Oracle XML DB schemas in 11.2 as compared to 11.1 (reference Bug 7646934). If you specify `xdb:defaultTable="MY_TAB"` without specifying `xdb:sqlInline="false"`, Oracle XML DB creates the table as requested and implicitly marks it as an out-of-line table. This behavior is different from 11.1 where the `defaultTable` annotation was ignored when the `sqlInline` setting was missing.

1.29.2 VARRAY Storage Default Change

In Oracle Database 11*g* Release 1 (11.1), the default value for `xdb:storeVarrayAsTable` changed from `FALSE` to `TRUE` for XMLType object-relational storage. This default applied to the default table, but not when creating XMLType object-relational tables and columns after the schema registration (reference Bug 6858659). In Oracle Database 11*g* Release 2 (11.2), all `VARRAY` data elements are created as tables by default. This

provides a significant performance increase at query time. In addition, note the following:

- Tables created prior to 11.2 are not affected by this. The upgrade process retains storage parameters. This only affects tables created in 11.2 or later.

- You can retain the pre-11.2 default of VARRAY storage as LOBs if you have small VARRAY data elements and you read and or write the full VARRAY all at once. You have two options to revert to the pre-11.2 behavior:

  - Re-register the schema with xdb:storeVarrayAsTable=FALSE. This affects the default and non-default tables.

  - Or, when creating the table (for non default tables), you can use the STORE ALL VARRAYS AS LOBS clause to override the default for all VARRAY data elements in the XMLType. This clause can only be used during table creation. It will return an error if used in the table_props at schema registration time.

- For schemas registered prior to 11.2 (when the default storage for VARRAY data elements was LOB), you can use STORE ALL VARRAYS AS TABLES clause to override the default for all VARRAY data elements in the XMLType.

1.30 PL/SQL

Note the following when working with PL/SQL.

1.30.1 Secure Database By Default Changes

The PUBLIC EXECUTE privilege for the DBMS_RANDOM PL/SQL package will be deprecated in future Oracle Database releases (reference Bug 7591837). Users who need to run this PL/SQL package should be given explicit EXECUTE privileges.

1.30.2 Stale Native Code Files Are Being Cached

Natively compiled PL/SQL and native code generated by the JIT compiler for Oracle JVM, may be cached in operating system files. The SHUTDOWN ABORT and SHUTDOWN IMMEDIATE commands do not clean these cached native code files (reference Bug 8527383).

To avoid problems caused by stale files in the cache, clean the files before restarting the database instance. The stale files that belong to the instance can be identified by name. The name patterns are as follows where *sid_name* is the system identifier name:

```
JOXSHM_EXT_*_sid_name_*
PESHM_EXT_*_sid_name_*
PESLD_sid_name_*
```

It is advisable to remove stale files as part of a startup script. For example, to clean stale files for the instance passed to the startup csh script by name on Linux, add the following three lines to the startup script:

```
rm JOXSHM_EXT_*_$1_*
rm PESHM_EXT_*_$1_*
rm PESLD_$1_*
```

The location of the native cache depends on the platform. It is defined by the location of files created by the operating system to implement shm_open requests on that platform. For example, Solaris stores shared memory segments in /var/tmp/.SHMD and /tmp/.SHMD. Linux stores them in /dev/shm.

1.30.3 DBMS_SCHEDULER PL/SQL Package

Consider the following when working with the DBMS_SCHEDULER PL/SQL package.

1.30.3.1 Event-Based Jobs The Oracle Scheduler supports event-based jobs. These are jobs that are started when a message is enqueued into a database queue. When the database queue used for this purpose is a secure queue, the QUEUE_SPEC attribute of event-based jobs and schedules is a pair of the form *queue_name, agent_name*. The reason an agent has to be specified is to help the scheduler determine which messages the user can see. In the case of nonsecure queues, if the user has dequeue privileges on a queue, then he or she can dequeue any agent that is subscribed to the queue. This means that he or she can see all of the messages in the queue. In the case of secure queues, the user has to be granted privileges to not only the queue but to agents as well.

In addition to granting dequeue privileges, the queue owner has to make a call to dbms_aqadm.enable_db_access() to enable the user to dequeue as a specific agent. The agent could be subscribed to the queue using a rule that filters messages that the user should not see.

The scheduler always subscribes to every queue (secure or otherwise) using its own agent SCHEDULER$_EVENT_AGENT. Because the scheduler runs as user SYS, it can see all the messages in the queue. In the case of secure queues, however, this should not be done because the purpose of secure queues is to enable the queue owner to provide some type of limited access to the queue by users. Therefore, the scheduler requires an agent name in the QUEUE_SPEC attribute. This agent name is used solely for the purpose of determining which messages the user can see. The scheduler does not dequeue messages as this agent. When the scheduler receives a notification that a message has arrived in a secure queue, it performs the following checks for each event-based job that depends on the queue.

1. It checks to determine if the job owner has dequeue privileges on the queue.

2. It checks to see if the agent name provided in the queue spec can be used by the job owner.

3. It checks to see if the agent is currently subscribed to the queue.

4. It checks if the incoming message is visible to the agent.

If all of these checks succeed, then the scheduler launches the event-based job. Regardless of whether the checks pass, the message is dequeued using the agent SCHEDULER$_EVENT_AGENT.

When providing an agent for the QUEUE_SPEC attribute, the user can provide an existing agent name or have a new one created for this purpose. In either case, it is the user's or queue owner's responsibility to ensure that the agent dequeues its messages.

1.30.3.2 Remote Jobs The Oracle Scheduler supports running jobs on remote hosts and transferring files to or from remote hosts. To achieve this goal, it is first necessary to go through the setup described in the "Enabling and Disabling Remote External Jobs" section of Chapter 28 of the *Oracle Database Administrator's Guide*. This feature also requires that Oracle XML DB HTTP Server be enabled on the database.

To check whether Oracle XML DB HTTP Server is enabled, you can use the following command:

```
SQL> SELECT DBMS_XDB.GETHTTPPORT() FROM DUAL;
```

If this statement returns `0`, then you need to enable Oracle XML DB HTTP Server on a non-zero port by logging in as `SYS` and issuing the following commands:

```
SQL> EXEC DBMS_XDB.SETHTTPPORT (port);
SQL> COMMIT;
```

Substitute `port` with Oracle XML DB HTTP Server port you wish to use.

1.30.4 UTL_TCP, UTL_HTTP, UTL_SMTP, UTL_MAIL, and UTL_INADDR PL/SQL Packages

A new security measure is introduced in this release for the following network-related PL/SQL packages: `UTL_TCP`, `UTL_HTTP`, `UTL_SMTP`, `UTL_MAIL`, and `UTL_INADDR`. The invoker of those packages needs additional privileges to connect to an external host or to resolve the name or the IP address of a host. The packages check the invoker for the necessary privileges only when the calls are made at runtime and raises an exception if the invoker lacks the privileges. This new security measure is implemented by Oracle XML DB access control list (ACL) mechanism and, therefore, requires Oracle XML DB to be installed in order to use those packages.

Each external host that a database user wants to connect to or to resolve the name or IP address for from the database is restricted by an access control list (ACL). To grant a user the privileges for the host, the database administrator should create an ACL, add the privileges to the ACL for the user, assign the ACL to the host, and commit the changes using the `DBMS_NETWORK_ACL_ADMIN` PL/SQL package. For example, to give the users `SCOTT` and `ADAMS` the permission to connect to `www.oracle.com` via HTTP (namely to connect to TCP/IP port 80), the database administrator should complete the following:

```
SQL> REM Creates a new ACL and adds SCOTT the privilege to the ACL to make TCP
connections
SQL> EXECUTE DBMS_NETWORK_ACL_ADMIN.CREATE_ACL('acl_for_oracle.xml', -
> 'ACL for www.oracle.com', 'SCOTT', TRUE, 'connect')

SQL> REM Adds ADAMS the privilege to the ACL to make TCP connections also
SQL> EXECUTE DBMS_NETWORK_ACL_ADMIN.ADD_PRIVILEGE('acl_for_oracle.xml', -
> 'ADAMS', TRUE, 'connect')

SQL> REM Assigns the new ACL to www.oracle.com for TCP/IP port 80 (HTTP)
SQL> EXECUTE DBMS_NETWORK_ACL_ADMIN.ASSIGN_ACL('acl_for_oracle.xml', -
> 'www.oracle.com', 80)

SQL> REM Commits to make the ACL take effect
SQL> COMMIT
```

The invoker of the `UTL_TCP`, `UTL_HTTP`, `UTL_SMTP`, and `UTL_MAIL` packages needs the `'connect'` privilege to make TCP/IP, HTTP, or SMTP connections to the external host. The invoker of the `UTL_INADDR` package needs the `'resolve'` privilege to resolve the name or the IP address of the external host. Note that those privileges are not granted through the `GRANT` SQL statement but through the `DBMS_NETWORK_ACL_ADMIN` package.

The current ACL assignment to external hosts and the privileges currently defined in the ACLs are shown through the system catalog views `DBA_NETWORK_ACLS` and `DBA_NETWORK_ACL_PRIVILEGES`.

The invokers of other database components that use those PL/SQL packages to perform network operations from the database, which are Oracle XML DB HttpUriType, Oracle Multimedia, and Oracle Spatial, are subject to the same network permission check and require the same privileges.

1.30.5 DBMS_SQL Package

In Oracle Database 11*g* Release 1 (11.1), Oracle introduces a number of enhancements to `DBMS_SQL` to improve the security of the package:

- Prevent guessing of open cursor numbers

  A new error, `ORA-29471`, will be raised when any `DBMS_SQL` subprogram is called with a cursor number that does not denote an open cursor. When the error is raised, an alert is issued to the alert log and `DBMS_SQL` becomes inoperable for the life of the session.

  If the actual value for the cursor number in a call to `IS_OPEN` does denote a cursor that is currently open in the session, then the return value is `TRUE`. If the actual is null, then the return value is `FALSE`. Otherwise, you get the `ORA-29471` error.

  Note that the `DBMS_SQL.OPEN_CURSOR` function is the only `DBMS_SQL` subprogram that has no formal parameter for the cursor number. Rather, it returns a cursor number. Therefore, it is not within the scope of the rules.

- Prevent inappropriate use of a cursor

  Cursors are now better protected from security breaches that subvert known, existing cursors.

  Checks are always made when binding and executing. Optionally, checks may be performed for every single `DBMS_SQL` subprogram call. The check is:

  - `current_user` is the same on calling the subprogram in question as it was on calling the most recent parse.

  - The enabled roles on calling the subprogram must be a superset of the enabled roles on calling the most recent parse.

  As is always the case, for definer's right subprograms, roles are irrelevant.

  If either check fails, then `ORA-29470` is raised.

  The mechanism for defining when checks are performed is a new overload for the `OPEN_CURSOR` subprogram which takes a formal parameter, `security_level`, with allowed values `NULL`, `1` and `2`.

  When `security_level = 1` (or is `NULL`), the checks are made only when binding and executing. When `security_level = 2`, the checks are always made.

  This security regime is stricter than in 10.2 and previous releases. As a consequence, users of `DBMS_SQL` may encounter runtime errors on upgrade. While the regime makes for more secure applications, users may want to relax the security checks temporarily as they migrate to 11.1. If so, consult with Oracle Support Services on steps to relax the security checks.

1.30.6 UTL_FILE Package

The `UTL_FILE` package no longer opens a file if that file is a symbolic link. This new restriction closes a known security hole in the `UTL_FILE` package.

If this security restriction poses an undue burden during migration, then contact Oracle Support Services for assistance.

1.30.7 Interoperability Between Releases

Interoperability between an 11.1 database or Forms client and a 10.1 or 10.2 database requires a minimum patchset level of 10.1.0.5 (for 10.1) or 10.2.0.2 (for 10.2).

An attempt to reference a 10.1 or 10.2 PL/SQL unit or view under the following circumstances fails with a `PLS-801[55916]` error unless the 10.1 or 10.2 environment has been patched to the correct level:

- A PL/SQL unit, anonymous block, trigger, call statement, or SQL statement on an 11.1 database runs a PL/SQL unit on a 10.1 or 10.2 database across a database link.

- A PL/SQL unit, anonymous block, trigger, or call statement on an 11.1 database references a view on a 10.1 or 10.2 database across a database link and the view directly or indirectly references a PL/SQL function or an object type.

- An 11.1 Forms client runs a PL/SQL unit in a 10.1 or 10.2 database using RPC.

To avoid the `PLS-801[55916]` error, a minimum patchset level of 10.1.0.5 (for 10.1) or 10.2.0.2 (for 10.2) is required.

There is no interoperability problem between 10.1 and 10.2 environments.

1.31 Pro*C

The Pro*C readme file is located at:

`ORACLE_HOME/precomp/doc/proc2/readme.doc`

1.32 Pro*COBOL

The Pro*COBOL readme file is located at:

`ORACLE_HOME/precomp/doc/procob2/readme.doc`

1.33 SQLJ

SQLJ in Oracle Database 11*g* Release 2 (11.2) is supported with JDK 5.0 and JDK 6.0.

1.34 SQL*Plus

The SQL*Plus readme file is located at:

`ORACLE_HOME/sqlplus/doc/README.htm`

1.35 Summary Management

Note the following items when working with Summary Management.

1.35.1 Query Rewrite

The following items apply to Query Rewrite.

1.35.1.1 Fine Grained Auditing (FGA) If Fine Grained Auditing (FGA) is enabled on a table in the query, then Query Rewrite will not occur for this query.

1.35.1.2 Use of PARTITION Clause Query rewrite does not occur for queries that use the `PARTITION` clause in the `FROM` clause to access table partitions. In order for query rewrite to rewrite such queries, the `PARTITION` clauses must first be converted into equivalent selection predicates and added to the `WHERE` clause.

1.35.2 Feature Availability

The following are available in Enterprise Edition:

- Creation and refresh features of materialized views

- Query rewrite and materialized view advice from the SQL Access Advisor

1.35.3 NLS Parameters

When using or refreshing certain materialized views, you must ensure that your NLS parameters are the same as when you created the materialized view. Materialized views that fall under this restriction contain the following constructs:

1. **Expressions that may return different values, depending on NLS parameter settings**

   It is recommended to write such expressions in the NLS-independent way. For example:

   ```
   (date > DATE '2003-01-02')
   ```

   Or:

   ```
   (rate <= 2.150)
   ```

2. **Equijoins where one side of the join is character data**

   The result of this equijoin depends on collation which can change on a session basis, giving an incorrect result in the case of query rewrite or an inconsistent materialized view after a refresh operation.

3. **Expressions that generate internal conversion to character data in the select list of a materialized view, or inside an aggregate of a materialized aggregate view**

   This restriction does not apply to expressions that involve only numeric data; for example, a+b where a and b are numeric values.

1.36 Oracle Streams

Note the following when working with Oracle Streams.

1.36.1 Propagation from Release 9.2 or Release 10.1 to Release 11.2 Returns Error ORA-25334

Propagation from Oracle9*i* Database Release 2 (9.2) or Oracle Database 10*g* Release 1 (10.1) to Oracle Database 11*g* Release 2 (11.2) results in error ORA-25334 unless the patch to Bug 4285404 has been applied to the 9.2 database.

1.36.2 Oracle Streams Apply Handler Modifications Necessary for Handling New Error Messages That Replace ORA-1403 Errors

Customized DML and error handlers for Oracle Streams require modification to catch the additional Oracle errors ORA-26786 and ORA-26787 in place of the ORA-01403 No data found message. An ORA-26787 error is raised if the row to be updated or deleted does not exist in the target table. An ORA-26786 error is raised when the row exists in the target table, but the values of some columns do not match those of the LCR.

1.36.3 Oracle Streams Advanced Queuing

Subscription names for Advanced Queuing notification will become canonical. Thus, registrations for scott.queue1:sub and SCOTT.QUEUE1:SUB are for the same entity and are represented in canonical form as "SCOTT"."QUEUE1":"SUB".

1.36.4 DBMS_RULE_ADM Operation Privileges

New in Oracle Database 11*g* Release 2 (11.2), the `DBMS_RULE_ADM` package might require that the caller be specifically granted the privilege to perform the `DBMS_RULE_ADM` operation (reference Bug 5523578).

For the short term, if needed, the previous behavior can be restored by setting event 25476 to any nonzero level. However, usage of this event may be deprecated in a future release, therefore it is recommended that you grant the relevant privileges should `DBMS_RULE_ADM` return a security-related error.

1.36.5 UTL_SPADV Load Privileges

An additional privilege is required to load the `UTL_SPADV` package into the Oracle Streams administrator schema. The Oracle Streams administrator schema must have `EXECUTE` privilege on `DBMS_LOCK` in order to load this package.

1.37 Documentation Addendum

This section contains corrections to the following Oracle Documentation for this release:

- Section 1.37.1, "Oracle Database Heterogeneous Connectivity User's Guide"
- Section 1.37.2, "Oracle Database PL/SQL Packages and Types Reference"
- Section 1.37.3, "Oracle Grid Infrastructure Installation Guide"

1.37.1 Oracle Database Heterogeneous Connectivity User's Guide

The following information should be removed from Chapter 4, "Using Heterogeneous Services Agents", section "Determining the Heterogeneous Services Parameters":

> The Distributed Access Manager has a refresh capability available through the menu and toolbar that allows users to rerun queries if necessary and update the data. When the data is refreshed, the tool verifies that the set of registered agents remains the same. If it is not, then the global view is updated. See *Oracle Enterprise Manager Administrator's Guide* and online help for more information about the Distributed Access Manager.

1.37.2 Oracle Database PL/SQL Packages and Types Reference

The `listno` parameter of the `DBMS_UTILITY` was inadvertently excluded from the *Oracle Database PL/SQL Packages and Types Reference*. For example, the `GET_PARAMETER_VALUE` function should read as follows:

GET_PARAMETER_VALUE Function

This function gets the value of specified initialization parameter.

Syntax

```
DBMS_UTILITY.GET_PARAMETER_VALUE (
    parnam      IN        VARCHAR2,
    intval      IN OUT    BINARY_INTEGER,
    strval      IN OUT    VARCHAR2,
    listno      IN        BINARY_INTEGER DEFAULT 1)
  RETURN BINARY_INTEGER;
```

Parameters

Parameter	Description
parnam	Parameter name.
intval	Value of an integer parameter or the value length of a string parameter.
strval	Value of a string parameter.
listno	List item number. If retrieving parameter values for a parameter that can be specified multiple times to accumulate values, use this parameter to get each individual parameter.

1.37.3 Oracle Grid Infrastructure Installation Guide

In *Oracle Grid Infrastructure Installation Guide*, in the section "Creating Disk Groups for Oracle Database Data Files," the guide indicates that logical volume managers are not supported with Oracle RAC. This is incorrect. The correct support information is as follows:

> Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. In addition, Oracle RAC requires a cluster logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

1.38 Open Bugs

This section lists known bugs for this release. A supplemental list of bugs may be found as part of the release documentation specific for your platform.

1.38.1 Database Upgrade Assistant (DBUA) Known Bugs

Bug 8724757

When upgrading a release 9.2.0.8 database with archived redo logs and the LOG_ARCHIVE_FORMAT initialization parameter is explicitly set to %t_%s.dbf, DBUA shows an error recommending that you change the format to add %r and then continue. But %r is not a supported format in 9.2.

Workaround: Remove LOG_ARCHIVE_FORMAT from the initialization parameter file or SPFILE and use the default format. In release 9.2.0.8 the default was %t_%s.dbf and in 11.2 the default is %t_%r_%s.dbf. Therefore, using the default will automatically update the format.

1.38.2 Deinstallation Tool Known Bugs

Bug 8737202

When using a standalone version of the deinstallation tool to deinstall Oracle Grid Infrastructure for a cluster, the deinstallation fails to remove Oracle Clusterware on the remote nodes if the you run the tool and do not have permissions to create the deinstallation home location on the remote nodes.

Workaround: If the standalone version of the deinstallation tool is used to remove Oracle Grid Infrastructure for a cluster, you need permissions to create the deinstallation home location on all the Oracle Clusterware nodes.

Bug 8726637

If the Oracle Clusterware home is created under top level `'/'` directory, deinstallation fails to remove the Oracle Clusterware software on remote nodes.

Workaround: At the end of the deinstallation, run the following command from the local node for each of the remote nodes:

`'ssh remote_node rm -rf crs_home/'`

Bug 8680498

If you try to deinstall Oracle Restart on a machine where there is a single-instance database Oracle home that is not managed by Oracle Restart, the deinstallation tool removes the `/etc/oratab` file.

Workaround: Make a backup of the `/etc/oratab` file before running the deinstallation tool to deconfigure Oracle Restart, and then restore the file after you deinstall Oracle Grid Infrastructure for a cluster.

Bug 8644344

When running the deinstallation tool to deinstall the database, you will be prompted to expand the Oracle home and to select a component. If you select the top level component, `Oracle Database Server`, and do not select the Oracle home, OUI does not show the message to run the deinstall utility and proceeds with the deinstallation of the database.

Workaround: Run the deinstallation tool to deinstall the Oracle home.

Bug 8635356

If you are running the deinstall tool from `ORACLE_HOME` that is installed on shared NFS storage, then you will see errors related to `.nfs` files during `ORACLE_HOME` clean up.

Workaround: To remove the `ORACLE_HOME`, run the `rm -rf ORACLE_HOME` command after the deinstall tool exits. Alternatively, you can use the standalone `deinstall.zip` and specify the location of the `ORACLE_HOME`.

Bug 8403941

Oracle Grid Infrastructure for a cluster deinstallation will fail when the Deinstall tool is run from a read-only location.

Workaround: If a standalone version of Deinstall tool is used to remove Oracle Grid Infrastructure for a cluster, the Deinstall tool should be downloaded and run from the location that is writable to the user running the tool. The user should also have permissions to create the deinstallation home location on all of the Oracle Clusterware nodes.

Bug 8666509

A deinstallation of Oracle Clusterware should ask you to detach any pre-11.2 Oracle RAC homes from the Oracle inventory.

Workaround: After you deconfigure and deinstall an upgraded 11.2 Oracle RAC home and want to continue with deconfiguration and deinstallation of the Oracle Grid Infrastructure for a cluster home, first detach any pre-11.2 Oracle RAC software homes from the central Inventory.

1.38.3 Oracle Application Express Known Bugs

Refer to "Open Bugs and Known Issues" in *Oracle Application Express Release Notes*.

1.38.4 Oracle Automatic Storage Management (Oracle ASM) Known Bugs

Bug 9327020

The ASMCMD `md_restore` command requires the value for the `COMPATIBLE.RDBMS` attribute to be set to 11.2.0.0.0.

Workaround: None.

Bug 8936616

The `ASM_DISKSTRING` specified in Oracle ASM initialization parameters does not take effect when trying to create a disk group.

Workaround: Go to the **Create Disk Group** dialog and specify **Disk Discovery Path** in the **Change Disk Discovery Path** dialog. ASMCA will list all the matching disks.

Bug 8783301

When the data files of a tablespace, created on a disk group, are dropped, a dependency between the database and disk group is not updated. This can prevent the Oracle RAC instance from starting after the disk group is removed.

Workaround: Execute the following command with the `-a` option that specifies a list of current disk groups after a tablespace is created or dropped:

```
srvctl modify database -d orcl -a DG1,DG2
```

Bug 8722063

The ASMCMD `lsdg` command does not show the correct voting disk locations in an Oracle ASM disk group.

Workaround: Use the `VOTING_FILE` field of the `V$ASM_DISK` view to find out if the voting file is located on that disk or not. Use the following SQL command:

```
SELECT VOTING_FILE FROM V$ASM_DISK
```

Bug 8721918

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) error numbers span from `501-12000`. Errors ranging from `501-9999` are reported with the format of `ACFS-00501-ACFS-09999`. The `oerr` tool does not recognize Oracle ACFS error messages in this form due to the extra zeros (0) that are prepended to the error number to create a five-digit number. For example, if `ACFS-00501` is returned, then running the command `oerr ACFS 00501` will not return data.

Workaround: When you run `oerr`, remove any preceding zeros from the error number that is returned in the error message. For example, if `ACFS-00501` is returned in an error, then run `oerr` as follows:

```
$ oerr ACFS 501
```

Bug 8623938

Querying a GV$ view while an Oracle ASM cluster is in a rolling migration state might cause the process that issued the query to fail and return an error similar to the following:

```
ORA-0600: internal error code, arguments: [ksxpcini:rm], [3], [218103808]
```

Workaround: Do not query GV$ views while an Oracle ASM cluster is in a rolling migration state.

When Automatic Storage Management Configuration Assistant (ASMCA) performs a rolling upgrade of an Oracle ASM cluster, it queries GV$ views before putting the cluster in rolling migration state. If ASMCA fails before completing the rolling migration, stop the rolling migration before relaunching ASMCA. Otherwise, ASMCA will fail when it performs its GV$ query. The SQL command to stop the rolling migration is:

```
ALTER SYSTEM STOP ROLLING MIGRATION;
```

1.38.5 Oracle Clusterware Known Bugs

Bug 9061246

If the Global Services Daemon (GSD) is enabled, it may fail to start on all nodes of the cluster when the Oracle Clusterware is started or when the nodes are rebooted. The GSD state can be verified by running the command `srvctl status nodeapps`.

Workaround: On any one of the node in Oracle Clusterware, issue the following command to start GSD on all nodes:

```
srvctl start nodeapps
```

Bug 8974576

REMOTE_LISTENER parameter was not changed after a SRVCTL ADD SCAN_LISTENER or SRVCTL REMOVE SCAN_LISTENER command.

In a future release, Oracle Clusterware may start managing REMOTE_LISTENER as part of a project to support multiple public networks each with their own SCAN. But in release 11.2, because of the need to include node VIPs in the REMOTE_LISTENER parameter of pre-11.2 databases that are upgraded, Oracle Clusterware does not have all the data needed to keep REMOTE_LISTENER properly updated. This is, instead, implemented by DBCA and DBUA, or by manual means.

Workaround: To manually change the REMOTE_LISTENER parameter, execute the following:

```
SQL> ALTER SYSTEM SET REMOTE_
LISTENER='t2000-cluster0-scan.t2000-18.oraclecorp.com' sid='*' scope=memory;

System altered.

SQL> ALTER SYSTEM REGISTER;

SQL> SHOW PARAMETER LISTENER;

NAME                                 TYPE        VALUE
------------------------------------ ----------- -----------------------------
remote_listener                      string      t2000-cluster0-scan.t2000-18.o
```

Bug 8791771

Executing the `srvctl modify database -d db_unique_name -y manual` command does not prevent Oracle Clusterware from automatically restarting the Oracle RAC database after rebooting the system. The Oracle RAC database is restarted automatically.

Workaround: None.

Bug 8786770

Silent Cluster Ready Services (CRS) installation fails when Oracle RAC release 9.2 is present.

Workaround: Pass `oracle_install_crs_AdvancedInstall=true` from the command line. For example, a command similar to the following:

```
./runInstaller -responseFile /scratch/rsmith/grid.rsp -silent oracle_install_crs_
AdvancedInstall=true
```

Bug 8760631

When an application resource is created using `crsctl add resource`, it returns `CRS-2518`.

Workaround: Execute the CRSCTL command after an action script is created in the local file system on all nodes in a cluster or shared file system or Oracle ACFS.

Bug 8739811

When you install Oracle Clusterware with a shared `ORA_CRS_HOME`, you may get the following error:

```
Failed to rmtcopy "/tmp/fileM3zR7m" to "ORA_CRS_HOME/gpnp/manifest.txt" for nodes
{node1,node2}, rc=256
```

Workaround: You can ignore this error.

Bug 8737732

If you upgrade Oracle Clusterware from release 11.1.0.6 to release 11.2 and you had stored the 11.1.0.6 OCR and voting disks on raw devices, then `cluvfy` will fail.

Workaround: You can ignore the `cluvfy` failure.

Bug 8736093

The Cluster Ready Service may fail trying to perform an OCR operation after customers have completed the upgrade to Oracle Clusterware 11.2.0.1.0. This problem occurs during the following circumstances:

- After a successful upgrade to Oracle Clusterware 11.2.0.1.0, and not on new installations.

- On clusters that have more than 2 nodes.

- The master Cluster Ready Service fails after completing the upgrade and the non-master Cluster Ready Service on other nodes started with the cluster active version that was less than 11.2.0.1.0. If the non-master Cluster Ready Service has been recycled between this time, this problem will not occur.

Workaround: Stop the Cluster Ready Service on all nodes using `crsctl stop res ora.crsd -init`, and restart it using `crsctl start res ora.crsd -init`.

Bug 8733944

Due to a problem in Oracle Clusterware starting with release 11.1.0.7, with the patches required for Exadata support or 11.1.0.7 CRS bundle Patch 1, in some cases the CSS daemon may fail when the clusterware is brought down on another node, either due to a shutdown command or a failure.

The symptom is an `ASSERT` in the CSSD log indicating that a maximum value has been exceeded. For example:

```
Group ID of xxxx exceeds max value for global groups
```

Workaround: Oracle recommends that customers running with the Exadata support patches or 11.1.0.7 CRS Bundle Patch 1 apply the patch for this bug to avoid this problem.

This problem may also be seen during an upgrade from 11.1.0.7 with patches as indicated above. To eliminate the potential of an 11.1.0.7 node failing during upgrade, the patch for this bug may be applied to the 11.1.0.7 nodes prior to upgrade.

When upgrading, it is recommended that the upgrade be completed on all nodes without restarting any non-upgraded nodes during the upgrade. If an 11.1.0.7 node does fail while performing the upgrade, it should be upgraded as opposed to restarted.

Bug 8727190

For logical corruption of a voting disk, CSSD reopens the voting disk and automatically sets it to `ONLINE` to avoid all node reboot, if it is possible. However, CSSD cannot recognize the corrupted voting disk and sets it to `OFFLINE` when Oracle Clusterware is started the next time.

Workaround: Re-create the corrupted voting disk.

Bug 8726128

The following command fails for RHEL4/OEL4 system:

```
cluvfy stage -post crsinst
```

This happens because the validation of Udev attributes fails on the local node. This also results in the failure of the Cluster Verification Configuration Assistant near the end of the Oracle Grid Infrastructure for a cluster installation session. The detail installation log shows `ERROR:PRVF-9802`.

Workaround: Ignore the failure in checking Udev attributes on RHEL4/OEL4. If Udev attribute validation is the only reported problem, then the installation should be considered successful.

Bug 8725983

When the Oracle ASM instance crashes or terminates abnormally, Oracle Clusterware may fail to restart database services.

Workaround: Start the service resources manually with SRVCTL.

Bug 8721204

The "Check sudo permission" step will fail in the Database Control deployment procedure due to a sudo setting.

Workaround: If the `grep requiretty /etc/sudoers` command returns `Defaults requiretty`, then you need to comment out the following line in `/etc/sudoers`:

`#Defaults requiretty`

Bug 8716580

Listener resource fails to start because `TNS_ADMIN` environment variable is set in the Oracle Clusterware daemon processes and the value is not suitable for the listener resource.

Workaround:

Restart Oracle Clusterware in a session that does not have the `TNS_ADMIN` environment variable set.

Retry the "Check sudo permission" step.

Bug 8715968

The prerequisite check for the Oracle Grid Infrastructure for a cluster installation does not report a problem if the Network Time Protocol (NTP) configuration file is missing but the NTP daemon is running on a cluster node. This causes Cluster Time Synchronization Service to enter into observer mode, as shown in the cluster `alert.log` file.

Workaround: Use Cluster Time Synchronization Service (`ctss`) instead of Network Time Protocol (NTP) to provide time synchronization services. To do this, deconfigure and deinstall NTP. To deactivate the NTP service, stop the existing `ntpd` service, disable it from the initialization sequences, and remove the `ntp.conf` file. For details on how to do this, refer to "Network Time Protocol Setting" in *Oracle Grid Infrastructure Installation Guide for Linux*.

Bug 8711200

When running the installer, if you choose "Fix and Run Again" after the prerequisite check, the Cluster Verification Utility script sometimes fails because `runfixup.sh` does not exist.

Workaround: Remove `/tmp/CVU_11.2.0.1.0_userID` from all of the cluster nodes and run the installer again.

Bug 8702731

When stopping Oracle Clusterware, shutting down the database may time out when the `SHUTDOWN IMMEDIATE` command is used.

Workaround: Before stopping Oracle Clusterware, shut down the database using the `SHUTDOWN ABORT` command.

Bug 8692998

The required nodes on an `Extend Cluster` or `Delete Node` may not be displayed correctly. The problem occurs if the host collection has not occurred.

Workaround: Trigger a `Refresh Host Configuration` and then retry the `Extend Cluster` or `Delete Node` deployment procedure. Take the following steps:

1. Cancel out of the interview.

2. Click on the **Database** tab.

3. Click on the **Deployments** link under the **Related Links** section.

4. Select all hosts and then click the **Refresh hosts** button.

After the job execution is complete, go back to the Software and Support subtab under the Database tab to launch the `Extend Cluster` or `Delete Node` deployment procedure.

Bug 8683759

When the CRS home is upgraded to version 11.2 and if Oracle RAC databases of earlier versions (10*g* or 11.1) are present, the cluster events pertaining to the 10*g* or 11.1 Oracle RAC databases and services are not posted to ONS clients.

Workaround: None.

Bug 8683517

After upgrading Oracle Clusterware and Oracle ASM to release 11.2, if the database is not upgraded, Oracle Clusterware will not enforce a proper shutdown order when you attempt to stop the Oracle Clusterware stack. This is because Oracle Database releases prior to release 11.2 do not share a dependency with release 11.2 Oracle ASM disk groups. If Oracle Clusterware attempts to stop a disk group resource before its dependent database shuts down, the disk group may fail to stop because it is still in use, thereby preventing the Oracle Clusterware stack from stopping.

Workaround: To stop the Oracle Clusterware stack, use the following command:

```
crsctl stop crs -f
```

Bug 3841387, 8373758

Oracle resources for 10.1 Oracle RAC Databases may not operate properly after upgrading Oracle Clusterware to 11.2.

Workaround: Apply the patches for Bug 3841387 and Bug 8373758 to the Oracle Database home.

Bug 8674080

During a rolling upgrade from Oracle Clusterware version 11.1 to 11.2, cluster events are not posted to Oracle Notification Service (ONS) clients for events occurring on the nodes whose software version is not 11.2. In addition, the events are not received by ONS clients running on the pre-11.2 nodes for the events occurring on the 11.2 nodes.

Workaround: Complete the upgrade of all the nodes in the cluster to 11.2. The events are received when the Active Version of the cluster is release 11.2.0.1.

Bug 8668083

In Oracle Restart environments, the following commands cannot perform a verification and return an error message if the user equivalence does not exist for the local node:

```
cluvfy stage -pre hacfg
cluvfy stage [-pre | -post] cfs
cluvfy stage [-pre | -post] acfscfg
cluvfy comp admprv
cluvfy stage -pre dbcfg
cluvfy stage -pre dbinst
```

Workaround: Configure SSH in Oracle Restart.

Contact Oracle Support Services to obtain the patches for these bugs and apply the patches to the Oracle Database home.

Bug 8667622

When SCAN listeners are running on different nodes, stopping and restarting `nodeapps` with the `force` option may result in the SCAN listeners being started on the same node.

Workaround: Relocate the SCAN VIPs and their associated listeners to different nodes with the following command:

```
srvctl relocate scan -i ordinal_number -n node_name
```

Bug 8667127

When the public interface is disabled or fails to function, Oracle Clusterware may end up stopping service resources and not restarting them on another server.

Workaround: Start the service resources manually with SRVCTL.

Bug 8666501

The length of time between Oracle Clusterware agent log file rotations is short. The log file can be overwritten in a few days.

Workaround: None.

Bug 8666150

The command `cluvfy stage -pre nodeadd` fails at the shared resources check when the Oracle Grid Infrastructure for a cluster home does not exist or is not shared.

Workaround: Manually create the Oracle Grid Infrastructure for a cluster home before executing the `cluvfy` command.

Bug 8659440

The command `srvctl modify scan_listener -p new_endpoints` only accepts a TCP port number as the `new_endpoints` argument.

Workaround: If it is necessary to update SCAN listeners to use endpoints that are not TCP port numbers, then run `srvctl stop listener -f` and `srvctl remove scan_listener -f` to stop and remove all SCAN listeners. Then, run `srvctl add scan_listener -p new_endpoints` to add SCAN listeners using the desired new endpoints, where `new_endpoints` has the following format:

```
[TCP:]port[/IPC:key][/NMP:pipe_name][/TCPS:s_port]
```

Bug 8657184

If two network interfaces are configured as public network interfaces in the cluster, the failure of one public interface on a node does not result in automatic VIP failover to the other public interface.

Workaround: If multiple public networks interfaces are present, then use interface bonding for high availability. At the Oracle Clusterware installer "Specify Network Interface Usage" screen, choose only one (bonded) interface as public. When configuring public networks with `srvctl add nodeapps` or `srvctl add vip`, specify only a single network interface name in the `-A` or `-S` argument.

Bug 8652158

After a VIP on a non-default network (created with command `srvctl add vip -k 2` ...) fails over to another node in the cluster due to a public network failure, it does not fall back automatically to its original node when the network is restored back to an operational state.

Workaround: Manually stop the concerned VIP and restart it using the following commands:

- srvctl stop vip -i *vipname*
- srvctl start vip -i *vipname*

This `start` command automatically starts the VIP on the correct node.

Bug 8648234

Executing the following commands may fail the check for membership in the OSDBA group if you use a different operating system user for each software installation and the Oracle Grid Infrastructure for a cluster installation user is not a member of the OSDBA group.

```
cluvfy comp sys -p crs
cluvfy comp sys -p ha
```

Workaround: You can ignore this failure if the installation user owns only Oracle Grid Infrastructure for a cluster. That is, if you use Job Role Separation, then the Oracle Grid Infrastructure for a cluster owner need not be a member of the OSDBA group. However, membership in the OSDBA group is a requirement for an Oracle Database owner.

Bug 8644639

When creating an Oracle ACFS mount point and adding it to the registry, the mount point is not mounted automatically if the following conditions are met:

1. The mount point directory was previously registered with the Oracle ACFS Registry.

2. The mount point directory had been previously mounted.

3. The mount point had then been unmounted and removed from the Oracle ACFS Registry.

4. The `ora.registry.acfs` resource has not been restarted since the mount point was deleted from the registry.

Workaround: Remove the mount point directory from the file */tmp*/.usm_state_file.

Bug 8262786, 8373758, 8441769, 8406545

Oracle resources for 10.2 and 11.1 Oracle RAC databases may not operate properly after upgrading Oracle Clusterware to 11.2.

Workaround: Apply the patches for Bugs 8262786, 8373758, 8441769, and 8406545 to the Oracle Database home.

Bug 8635825

If two network interfaces are configured as public network interfaces in the cluster, failure of one public interface on a node does not result in automatic virtual internet protocol address (VIP) failover to the other public interface.

Workaround: If multiple public networks interfaces are present, use interface bonding for high availability. At the Oracle Clusterware Installer's "Specify Network Interface Usage" screen, choose only one (bonded) interface as public. When configuring public networks with `srvctl add nodeapps` or `srvctl add vip`, specify only a single network interface name in the `-A` or `-S` argument.

Bug 8632280

After you complete an Oracle Clusterware upgrade to release 11.2, running the `oifcfg getif` command displays the error `PRIF-30: Network information in OCR profile needs to be synced with GPnPprofile`.

Workaround: Run the following command from the Oracle Clusterware home `bin` directory as the `root` user where *interface_name* is the name of the interface and *subnet* is the name of the subnet:

```
oifcfg setif -global interface_name/subnet public
```

Bug 8612566

Some public network failures result in OCI FAN clients no longer receiving any notifications, and hanging for the duration of their TCP timeout.

Workaround: None.

Bug 8607080

Oracle Clusterware ignores the value of `NLS_LANG` that is set in the environment. This means that the setting put into Oracle Clusterware at installation time must be used by all administrators.

Commands such as `crsctl` and `srvctl` are affected by this because CRSD and OHASD error messages are translated into the default value before being passed back to the CLI.

Workaround: If the language ID has been incorrectly set at installation time and you want to change the `NLS_LANG` value, execute the following commands on all nodes:

1. `crsctl stop crs`

2. `edit s_crsconfig_hostname_env.txt`

3. `crsctl start crs`

This workaround is to be used to change the value of `NLS_LANG` after Oracle Clusterware has been installed. To ensure that proper `NLS_LANG` is used during installation of Oracle Clusterware, execute the following command:

```
edit grid_home/crs/install/crsconfig_params
```

For example:

```
SILENT=false
ORACLE_OWNER=grid
ORA_DBA_GROUP=oinstall
ORA_ASM_GROUP=asmadmin
LANGUAGE_ID='JAPANESE_JAPAN.JA16EUC
```

Bug 8600842

When performing a rolling upgrade of Oracle Clusterware to release 11.2, some or all of the resources (such as Oracle RAC database instances, listener processes, and database services) may not be automatically started on the last node (upgrade is usually done on a node-at-a-time basis).

Workaround: Resources should be manually automatically started using normal management interfaces (for example, SRVCTL for Oracle database instances and services).

Bug 8568386

When adding a voting disk to a shared file system, its group permission is automatically set to the root user. However, the group permission should be set to the Oracle Grid Infrastructure for a cluster installation user.

Workaround: Manually change the group permission of the voting disk by using the `chown` command.

Bug 8520511

If you install Oracle Database release 10.2 or release 11.1 with data files on Oracle ASM on a release 11.2 Oracle Clusterware with OCR voting disk on Oracle ASM, the DBCA installation fails with the error `ORA-00119`.

Workaround: Apply the one-off patches that are mentioned in Bug 8288940 or create a symbolic link using the following command:

```
Symlink Pre-11.2_Oracle_Database_home/network/admin/listener.ora --> 11.2_Oracle_
Clusterware_home/network/admin/endpoints_listener.ora
```

Where *Pre-11.2_Oracle_Database_home* is the Oracle Database release 10.2 or release 11.1 database home.

Bug 8493252

In rare cases, Oracle Clusterware installation may fail during the `AttachHome` operation when the remote node closes the SSH connection.

Workaround: Set the `LoginGraceTime` parameter to 0 in the SSH daemon configuration file `/etc/ssh/sshd_config` on all cluster nodes.

Bug 8448079

When Oracle Clusterware is stopped and restarted, Oracle ASM disk groups do not retain their pre-CRS shutdown states. Release 11.2 database resources will have a start dependency on the Oracle ASM disk group. In this case, Oracle Clusterware will start the Oracle ASM disk group before starting the database resource. However, Oracle

Database releases prior to release 11.2 do not have a dependency on the Oracle ASM disk group and will fail to start.

Workaround: Prevent the Oracle ASM parameter `ASM_DISKGROUPS` from being set to `NULL` when stopping Oracle Clusterware by disabling Oracle ASM disk group resources and manually setting the Oracle ASM parameter `ASM_DISKGROUPS`. The following example is disabling the disk group resources for disk groups `data` and `fra`:

```
srvctl disable diskgroup -g data
srvctl disable diskgroup -g fra
sqlplus / as sysasm
SQL> ALTER SYSTEM SET asm_diskgroups=data,fra scope=both;
```

> **Note:** If you are upgrading the database to release 11.2 using Database Upgrade Assistant, ensure that the disk groups used by the database are enabled. Otherwise, Database Upgrade Assistant may fail during the upgrade process. When Database Upgrade Assistant attempts to start a database resource after an upgrade, it first tries to start the disk group resources on which the database datafiles reside. If the disk group resources are disabled, then these disk groups will fail to start. To enable disk group resources, run the following commands before using Database Upgrade Assistant:
>
> ```
> srvctl enable diskgroup -g data
> srvctl enable diskgroup -g fra
> ```

Bug 8434421

The first Cluster Synchronization Services Daemon (CSSD) log file (`ocssd.log`) is not overwritten by log rotation because it is owned by the root user.

Workaround: Either remove `ocssd.log` or change its ownership.

Bug 8309620

After 11.2 Oracle Clusterware is installed and configured on a cluster with OCR and voting disk on a file system rather than Oracle ASM, and later a pre-11.2 Oracle Database is installed, the pre-11.2 DBCA will attempt to bring up 11.2 Oracle ASM as a single-instance Oracle ASM rather than a cluster Oracle ASM instance.

Workaround: Execute the following commands after the 11.2 Oracle Clusterware configuration is complete, but before installing a pre-11.2 Oracle Database, to start and then stop Oracle ASM on any node:

```
srvctl start asm -n node
srvctl stop asm -n node
```

Bug 8276914

If `ORACLE_BASE` is set in the environment of the root user before starting Oracle Clusterware, then the location of the Oracle Clusterware log directory of user-defined listeners and SCAN listeners changes from `ORACLE_BASE/diag/tnslsnr` to *Grid_ home*`/log/diag/tnslsnr` after the node is rebooted.

Workaround: Do not set `ORACLE_BASE` in the root user's environment before starting Oracle Clusterware.

1.38.6 Oracle Configuration Manager Known Bugs

Bug 8709678

The error `[: -eq unary operator expected :failed` is returned when configuring the Oracle Configuration Manager either using `setupCCR` or `configCCR`.

Workaround: Disable the creation of a `crontab` entry by defining the environmental variable `CCR_DISABLE_CRON_ENTRY` to `true` and rerun the Oracle Configuration Manager command to configure the OCM instance.

1.38.7 Oracle Data Guard Known Bugs

Bug 8716064

The `ANALYZE TABLE VALIDATE` statement on an Oracle Active Data Guard standby database may lead to `ORA-600 [4555]/ORA-600[25027]`.

Workaround: Issue the command on the primary database.

Bug 8428523

The `ALTER TABLE RENAME` statement on the primary database may lead to incorrect results on the Active Data Guard standby database after the `RENAME` DDL redo has been applied.

Workaround: Flush the shared pool on the Active Data Guard standby database after the `RENAME` DDL redo has been applied.

Bug 8363109

Flashback queries of a recent SCN on an Active Data Guard standby database may return `ORA-8181`. One example is that queries using asynchronous XMLIndex may get `ORA-8181` when executed on an Active Data Guard standby database. Another example is that queries using XMLType stored in Object Relational may get `ORA-8181` when executed on an Active Data Guard standby database.

Workaround: Issue the queries on the primary database.

1.38.8 Oracle Database Enterprise Edition Known Bugs

Bug 9181675

The `asmgidwrap` script needs to be called if you are creating a database manually on Oracle ASM to avoid a permission error.

Workaround: For a role-separated installation (that is, there is a different user and group for grid and RDBMS), use DBCA to create the database that automatically calls `asmgidwrap` script while creating a database on Oracle ASM. If you choose to create a database manually, the script needs to be called explicitly so the proper group can be set to avoid a permission error.

Bug 8867819

The execution plan may change in release 11.2 since the default behavior for optimizer feedback is now on. Therefore, when no change is made, two SPA trials may report regression depending on the query.

Workaround: For SPA to be consistent, turn off optimizer feedback. Use `_optimizer_use_feedback=false`.

Bug 8747281

In 11.2, OCR has been replaced by the Grid Plug and Play profile as a source for cluster interconnects for Oracle RAC database instances. However, the fixed view `V$CLUSTER_INTERCONNECTS`, still displays the `SOURCE` of these interconnects as `OCR`, rather than `Profile`.

Workaround: None.

Bug 8729627

When using 11.1 DBCA to remove a database on a cluster running 11.2 Oracle Clusterware, a `PRKP-1061/CRS-2524` error may be displayed because the database resource is locked.

Workaround: You can ignore the message. Click **OK** to continue.

Bug 8722270

When the database is configured to run on a subset of cluster nodes, and Oracle Clusterware is started on a node where the database had not previously run, the database may not start properly.

Workaround: Start the database on the affected node using the `srvctl start instance` command.

Bug 8714383

Oracle Call Interface (OCI) FAN clients do not receive `DATABASE DOWN` events when the last instance of an Oracle RAC database fails or is shut down.

Workaround: None.

Bug 8682141

Creating the encrypted tablespace using the `CREATE LIKE` option results in an internal error.

Workaround: create the encrypted tablespace using the `CREATE` option instead of the `CREATE LIKE` option.

Bug 8679182

When configuring a database on a cluster that has multiple public subnets defined for its VIPs (for example, using a command similar to `srvctl add vip -k 2 -A ...`), the database agent automatically sets `LOCAL_LISTENER` to the listener on the default network. This may duplicate a listener set in `LISTENER_NETWORKS`.

Workaround: Do not specify listeners in `LISTENER_NETWORKS` that are on the default public subnet.

Bug 8671032

In Oracle Enterprise Manager, the "TIP" obtained from the "Run Analysis" on the Automatic Undo Management page is not correct. For example, the Recommended

Undo Tablespace Size should be three times the Minimum Size, to allow for workload fluctuations, but Oracle Enterprise Manager shows different values.

Workaround: The "TIP" should read as follows:

```
Oracle advises that you configure the undo tablespace to be three times the
Recommended Undo Tablespace Size to allow for workload fluctuations
```

Bug 8660571

Issuing a DML statement against a partitioned table with the parallel option enabled and at least one XMLIndex index in asynchronous (deferred) maintenance mode may fail.

Workaround: Disable the parallel option for the affected table as shown in the following example:

```
ALTER TABLE partitioned_table NOPARALLEL;
```

Bug 8658472

When a database is created using DBCA on a cluster, SPFILE is used by default and the location is stored in Oracle Cluster Registry (OCR). But, the SPFILE location can not be cleared to use PFILE.

Workaround: If PFILE is used after database creation, then the database resource needs to be re-created by running the following commands without the -p option:

```
srvctl remove database -d db_unique_name
srvctl add database -d db_unique_name -o ORACLE_HOME -a diskgroup
```

Bug 8639114

Continuous Query Notification (CQN) may not properly support tables that are OLTP compressed. The errors that you may see include ORA-7445 or ORA-600.

Workaround: Do not use CQN + table compression together.

Bug 8601114

The Database Control operation "Import from Export Files" or "Import from Database" is erroneously reported as failed if the "Action to Take if Table Exists" is either Skip, Append, or Truncate.

Workaround: None. The operation succeeds.

Bug 8595729

A large number of audit files may be created in the $GRID_HOME/rdbms/audit/ directory by the Oracle Clusterware monitoring of Oracle ASM disk group resources.

Workaround: Periodically remove the audit files with the prefix +asm stored in the $GRID_HOME/rdbms/audit/ directory.

Bug 8595653

The endpoints_listener.ora file is used to get endpoints of the default listener when data files of Oracle Database 11*g* Release 1 or Oracle Database 10*g* Release 2 are created on a release 11.2 Oracle ASM disk group. However, when the listener is modified (such

as changing a port number using Network Configuration Assistant), the `endpoints_listener.ora` file is not updated.

Workaround: If a default listener is modified, edit the `endpoints_listener.ora` file manually. The file is stored in the `$GRID_HOME/network/admin/` directory.

Bug 8304261

If compression attributes are specified at the table level and in the sub-partition template, newly created interval partitions do not use the compression attributes defined in the sub-partition template, but use the table level compression instead.

Workaround: Redefine compression attributes in sub-partitions.

1.38.9 Oracle Database Vault Known Bugs

Bug 8728716

`DVSYS.DBMS_MACADM.UPDATE_RULE_SET` procedure may not update some attributes of the target rule set correctly.

Workaround: Delete the rule set and re-create a new rule set with desired attributes.

Bug 8706788

Oracle Enterprise Manager may generate a realm configuration alert.

Workaround: Remove the participant membership of `WKSYS` and `WKUSER` from `Oracle Data Dictionary` realm.

Bug 8686258

Database Vault policy cannot be managed in Oracle Enterprise Manager Database Control because the following message is displayed in Database Vault Administration page:

```
"OPERATOR TARGET" privilege does not exist.
"You must have OPERATOR TARGET privilege to perform this operation."
```

Workaround: To manage Database Vault policy using Oracle Enterprise Manager, the Database Vault administrator must have the EM Administrator privilege. If you do not want to grant the EM Administrator privilege to the Database Vault administrator, then use the Database Vault Administrator page directly. For additional information, see *Oracle Database Vault Administrator's Guide*.

Bug 8474817

If a user is granted Database Vault scheduler job authorization by `DBMS_MACADM.AUTHORIZE_SCHEDULER_USER` API, dropping this user does not remove the user from the list of authorized users.

Workaround: Invoke the `DBMS_MACADM.UNAUTHORIZE_SCHEDULER_USER` API call to remove the user's authorization manually.

Bug 8341283

The `ACTION_NAME` entry in the `DVSYS.AUDIT_TRAIL$` table displays `Realm Authorization Audit` for a failed realm enforcement if the audit option is set to audit

on success and failure. The RETURNCODE will show the correct error code that was triggered.

Workaround: Use the RETURNCODE value to determine whether a violation has occurred and the ACTION_NAME column to identify whether the audit was generated by a realm enforcement or command rule enforcement.

Bug 8247647

Dropping a user does not automatically remove the user from the list of authorized users for Oracle Data Pump.

Workaround: Invoke the DBMS_MACADM.UNAUTHORIZE_DATAPUMP_USER API call to manually remove the user's authorization for Oracle Data Pump.

Bug 7033772

Database Vault Administrator (DVA) does not work after an Enterprise Manager DBControl upgrade.

Workaround: Manually redeploy DVA after DBControl has been upgraded. You can follow the steps described in Appendix C, Section "Deploying Database Vault Administrator to the Database Console OC4J Container" of the *Oracle Database Vault Administrator's Guide*.

1.38.10 Oracle Enterprise Manager Database Control Known Bugs

Bug 8702939

If you attempt to add a new node, then "Grid Infrastructure User Permission Checks" and or the "Checks if Staging Area Path is writable" may fail.

Workaround: Manually create the directory for the Oracle RAC database home as Home/EMStage/PAF with 775 permissions and retry the failed step.

Bug 8757994

If you explicitly set a value in any of the following text boxes in the **Placement** section of the **Add Resource** dialog box and then change focus to any other text box, you will not be able to set a value in the text box to which you changed your focus and the page appears frozen for several seconds:

- Favor the placement to some Server Pools and or Servers:
  - Server Pools text box
  - Servers text box
- Restrict the placement to some Server Pools or Servers:
  - Server Pools text box
  - Servers text box

Workaround: When the page appears frozen, wait 5 to 10 seconds and try again. Or, set the value of the Server Pools or Servers text box by selecting from the List of Values instead of explicitly setting the values in the text box.

Bug 8724077

Due to a mass cluster file transfer problem in large clusters, typically in clusters with more than eight nodes, DBCA may intermittently report a failure when configuring Database Control and display the error message NULL. DBCA then and asks you to run the EMCA command line tool to configure Database Control.

Workaround: Run the command `emca -config dbcontrol db -repos recreate -cluster` after the database creation process completes.

Bug 8716161

The following error is returned when attempting to stop DBConsole:

```
$ emctl stop dbconsole
Environment variable ORACLE_UNQNAME not defined. Please set ORACLE_UNQNAME to
database unique name.
```

Workaround: Set `ORACLE_UNQNAME` environment variable to match the `db_unique_name` initialization parameter.

Bug 8692858

Oracle RAC Add Node with Oracle ASM storage fails intermittently due to updates in the `chartCache` files in Oracle Enterprise Manager Database Control.

Workaround: Before you use Oracle Enterprise Manager Database Control to extend the cluster, or before you run the `addNode.sh` script, add the following directory to *ORACLE_HOME*/rdbms/install/install.excl, where *ORACLE_HOME* is the Oracle RAC home:

```
oc4j/j2ee/oc4j_applications/applications/em/em/images/chartCache/*
```

Bug 8741598

When adding nodes with Database Control, the `cluvfy stage -pre nodeadd` command fails without providing an error.

Workaround:

If the `cluvfy pre` or `cluvfy post add node` checks fail, then you can ignore the step and run the following commands to verify the validity of the new nodes and proceed with the rest of the deployment steps. Use the following syntax where *grid_infrastructure_ORACLE_HOME* is the name of the home of the Oracle Grid Infrastructure for a cluster installation home and *comma_separated_list_of_new_nodes* is the list of new nodes:

```
grid_infrastructure_ORACLE_HOME/bin/cluvfy -pre nodeadd -n comma_separated_list_
of_new_nodes -verbose -fixup
grid_infrastructure_ORACLE_HOME/bin/cluvfy -post nodeadd -n comma_separated_list_
of_new_nodes -verbose
```

Bug 8674920

If the installation owners for the Oracle Grid Infrastructure for a cluster and Oracle Database are different, then the owners of Oracle ASM binaries and Oracle Enterprise Manager Agent binaries are also different. When you start Support Workbench, the error message `Error Operation failed - Operation failed` might appear, because the Oracle Enterprise Manager Agent is running as a different user, and Support Workbench does not have permissions for the Oracle ASM target.

Workaround: None.

Bug 8655988

On the "Hang Analysis" page of Oracle Enterprise Manager Database Control for Cluster Database, the text in the "Wait" graph is shown as garbled characters when you use the Japanese locale.

Workaround: None.

Bug 8444875

You cannot create an application VIP from Oracle Enterprise Manager Database Control.

Workaround: Use the following command as root with the appropriate name for the VIP and the appropriate operating system user ID:

```
appvipcfg create -network=1 -ip=10.185.148.111 -vipname=appsvip -user=oracle
```

Bug 8350597

If Database Control is running in an IPv6 environment, then you cannot use it to monitor Exadata cells and you should not add Exadata cells as targets.

Workaround: None.

Bug 8638267

If you select the "Desktop Class" style database configuration in the Installer, and after completing the installation you attempt to create a database using DBCA or any Database Control setup using Oracle Enterprise Manager Configuration Assistant (EMCA), then you must set the ORACLE_HOSTNAME environment variable to 'localhost'. If you do not set ORACLE_HOSTNAME, then DBCA fails while configuring Oracle Enterprise Manager with the following error:

```
Listener is not up or database service is not registered with it. Start the
Listener and register database service and run EM Configuration Assistant again.
```

Workaround: Set the ORACLE_HOSTNAME environment variable to 'localhost' and retry creating the database.

1.38.11 Oracle Grid Control Known Bugs

Bug 8609658

In release 11.2, only the SYSASM user has the privilege to start up or shut down Oracle ASM. Preferred credentials stored in Oracle ASM from prior releases using the SYSDBA role can no longer be used due to Oracle ASM File Access Control.

Workaround: Specify startup and shutdown credentials for Oracle ASM using the SYSASM role.

1.38.12 Oracle Grid Infrastructure for a Cluster Known Bugs

Bug 8741415

In the "One Click Extend Cluster Database" deployment procedure, the "Checks for sudo permission" step fails because there is no /usr/local/bin/perl.

Workaround: You can ignore this step.

Bug 8739647

When installing the Oracle Grid Infrastructure for a cluster home as a separate user (someone other than the user who owns the database *ORACLE_HOME*), if `/opt/oracle` already exists on the host system, the installation will not change the group ownership of the directory if it is wrong. If the user does not have the appropriate group ownership set, one or more of the following errors would be seen in the alert log or the DBCA log:

ORA-15081
ORA-15025
ORA-15186

Workaround: Before installing the Oracle Grid Infrastructure for a cluster home, change the group ownership of `/opt/oracle` to the group that owns the Oracle installation inventory—which typically is `oinstall`—as shown in the following example:

```
chgrp oinstall /opt/oracle
```

> **Note:**   You need to run this command as the root user.

Bug 8679335

If you install Oracle Grid Infrastructure for a cluster and Oracle RAC as different users, you may get an error due to insufficient permissions on the *ORACLE_HOME*/EMStage directory, where *ORACLE_HOME* is the Oracle RAC home.

Workaround: Manually create the directory *ORACLE_HOME*/EMStage/PAF with 775 permissions and retry the failed step.

1.38.13 Oracle Multimedia Known Bugs

Bug 8725198

The File interface of the `DicomLocatorService` and `DicomTagService` classes in the Oracle Multimedia Mid-Tier Java API may raise the `java.io.IOException: "Too many open files"` exception if executed on large numbers of files.

Workaround: Create an `InputStream` from a File and call the `InputStream` interface of the `DicomLocatorService` and `DicomTagService` classes of the Oracle Multimedia Mid-Tier Java API instead of the File interface.

1.38.14 Oracle OLAP Known Bugs

Bug 8474636

An error occurs when building a partitioned cube under the following circumstances:

- `NLS_DATE_LANGUAGE` is set to a non-English language.

- The cube is partitioned by a Time dimension.

- The column that maps to the Time dimension has a Date data type.

Workaround: Set `NLS_DATE_LANGUAGE` to `ENGLISH`. You can set `NLS_LANGUAGE` to whatever language your applications support.

1.38.15 Oracle Net Services Known Bugs

Bug 8682102

Once a wallet is created and a master key generated, the wallet must not be deleted. The database needs to be able to locate this master key to perform some encryption-related operations.

Workaround: If you deleted the wallet and if a backup copy of the wallet exists, you may be able to restore the backup copy to the wallet location specified in the `sqlnet.ora` file.

Bug 7602834

On the "Network Administration" home page, if you select "Directory Naming" and provide Oracle Internet Directory credentials, you will get an internal server error.

Workaround 1:   Use Oracle Net Manager instead.

Workaround 2:   Run the following commands:

```
cp ORACLE_HOME/network/lib/libnldapj11.so ORACLE_HOME/lib/.
cd ORACLE_HOME/bin
./emctl stop dbconsole
./emctl start dbconsole
```

1.38.16 Oracle Real Application Clusters Known Bugs

Bug 8716518

If using a separate user to install Oracle Restart, before DBCA, add the user `grid` to the group `dba`. For example:

```
# usermod -a -G dba grid
```

Workaround: None.

Bug 8581327

If standalone listeners are running from the Oracle RAC `ORACLE_HOME` that is being patched, then the procedures "Patch Oracle RAC - Rolling" and "Patch Oracle RAC - All Nodes" fail at the step `CheckActiveFilesAndExecutables` because the file `ORACLE_HOME`/bin/tnslsnr is active.

Workaround: Stop the standalone listeners that are running from the Oracle RAC `ORACLE_HOME` that is being patched and retry the procedure that failed.

1.38.17 Oracle Streams Known Bugs

Bug 8739024

In statement DML handlers, binding `NCLOB` columns or `CLOB` columns with varying width character sets using an LCR field expression such as `:new.nclob_column`, can result in an incorrect NCLOB or CLOB in the destination table.

Workaround: Use procedure DML handlers for customizing the apply for tables with these columns.

Bug 8735201

The Streams Apply process may abort with an `ORA-26786` error when applying changes from OLTP compressed tables. An `ORA-26786` error typically indicates that a data conflict has occurred while applying the source row.

Workaround: On the source database, check if the table involved in the error has OLTP compression enabled. Substitute the appropriate table name for *tablename* in the appropriate query below. For a non-partitioned table, use the following query:

```
SELECT table_name, compression, compress_for FROM user_tables WHERE
table_name = 'tablename';
```

For a partitioned table, use the following query:

```
SELECT table_name, partition_name, tablespace_name, compression, compress_for
 FROM user_tab_partitions WHERE table_name = 'tablename';
```

If a conflict error occurs, use either a Streams conflict handler or error handler to allow the DML to be executed. To ensure that the apply process continues processing following an occurrence of the `ORA-26786` error, set the `disable_on_error` apply parameter to `'N'`:

```
exec DBMS_APPLY_ADM.SET_PARAMETER('apply_name', 'disable_on_error', 'N');
```

1.38.18 Oracle Universal Installer Known Bugs

Bug 9593297

During an upgrade of 10.2.0.5 Cluster Ready Services to 11.2.0.1 Oracle Clusterware, Oracle Universal Installer (OUI) hangs at the first screen.

Workaround: Before starting 11.2.0.1 Oracle Clusterware upgrade, shutdown Oracle Clusterware using the following command. The upgrade needs to be performed in non-rolling approach.

Bug 8979500

There is an error when installing the Oracle Grid Infrastructure for a cluster with the Grid Naming Service (GNS) and the Oracle Cluster Registry (OCR) and voting disks stored on Oracle ASM.

When installing the Oracle Grid Infrastructure for a cluster for Oracle Database 11*g* Release 2 with GNS and the OCR and voting disks on Oracle ASM, `root.sh` may fail with error `PROT-1: Failed to initialize ocrconfig`. The failing command is `ocrconfig -upgrade oracle dba`, and `TNS-12560` can be found in the SQL*Net log file.

Workaround: Check the `/etc/group` entry for the group being used by OCR, temporarily reduce the number of users listed for this group, and restart the installation.

Bug 8947155

When installing a single-instance Oracle Database with the Desktop class option, and if the domain name in the global database name exceeds 128 characters, Oracle Universal Installer does not display a warning. The installation fails during the Database Configuration Assistant phase with the following error:

```
ERROR at line 1:
```

```
ORA-02086: database (link) name is too long
```

Workaround: Ensure that the domain portion of the Global Database Name does not exceed 128 characters including punctuation.

Bug 8756488

The **Installer Summary** dialog displays a size requirement that does *not* include log and other configuration data sizes.

Workaround: Refer to your respective installation guide for information about the correct amount of free disk space that is required for a successful installation.

Bug 8752235

If you select Oracle Automatic Storage Management (Oracle ASM) as the installation location for Oracle Cluster Registry (OCR) and then move to the next Oracle Universal Installer (OUI) screen and change the installation location for OCR to `File System`, OUI executes `root.sh` before starting the Automatic Storage Management Configuration Assistant (ASMCA), and ASMCA fails.

Workaround: Either do not execute ASMCA or ignore its failure and finish the installation.

Bug 8744509

During silent installation, Oracle Universal Installer does not accept a fully qualified host name for the `responseFile` property `oracle.crs.config.clusterNodes. responseFile`.

Workaround: In the response file to be used for silent installation, specify the host name without the domain. For example:

```
oracle.install.crs.config.clusterNodes=node1,node2
```

Not this:

```
oracle.install.crs.config.clusterNodes=node1.mydomain.com,node2.mydomain.com
```

Bug 8729326

When upgrading to 11.2 Clusterware, the Installer invokes ASMCA in silent mode to upgrade Oracle ASM into Oracle Grid Infrastructure for a cluster home. Oracle ASM upgrade is handled in rolling fashion when upgrading from 11.1.0.7. Prior versions of Oracle ASM instances are upgraded in non-rolling fashion and Oracle ASM-based databases are bounced without any prior warning.

Workaround: You can plan your database outage to be the point where you acknowledge the Installer prompt after executing `root.sh` on all nodes. At this point, CRS is upgraded in rolling fashion and the Installer will be calling ASMCA to upgrade Oracle ASM, which will bounce databases as part of Oracle ASM upgrade.

Bug 8714456

Choose the "Upgrade Grid Infrastructure" option to upgrade older versions of Oracle Automatic Storage Management (Oracle ASM) to Oracle Database 11g Release 2 (11.2).

Workaround: None.

Bug 8675426

The Enterprise user security utility, UMU (User Migration Utility), fails with the following error when migrating database users to the directory:

```
Exception in thread "main" java.lang.NoClassDefFoundError:
oracle/security/pki/OraclePKIProvider
        at java.lang.Class.forName0(Native Method)
        at java.lang.Class.forName(Class.java:242)
```

Workaround: The *ORACLE_HOME*/umu script expects a jar under the ORACLE_HOME/ldap/jlib directory. However, the jar is located under ORACLE_HOME/jlib instead. Create a symbolic link for ORACLE_HOME/jlib under ORACLE_HOME/ldap/jlib. For example, on a Linux, use the following command in the ORACLE_HOME/ldap directory:

```
ln -s ORACLE_HOME/jlib jlib
```

Bug 8670579

When running the Linux 32-bit operating system on 64-bit Advanced Micro Devices (AMD) Athlon hardware, there is a bug in third party code where the CPU type is not correctly identified. This may result in an application crash when using DBCA or NETCA startup.

Workaround: Apply the patch for this bug to avoid this problem.

Bug 8666656

The Oracle Universal Installer (OUI) runInstaller script that resides in the Oracle home (*ORACLE_HOME*/oui/bin/runInstaller) cannot be used to install the 11.2.0.1 releases of Oracle Database, Oracle Grid Infrastructure for a cluster, and Oracle Database Client.

Workaround: Use Oracle Universal Installer on the respective 11.2.0.1.0 product media to install each product.

Bug 8638708

If you select the database configuration Desktop Class in Oracle Universal Installer (OUI), listener and database control are configured with 'localhost' as the host name. The Oracle Enterprise Manager Database Control start and stop operations using emctl may fail.

Workaround: For Database Control start and stop operations that use emctl in that home, set the ORACLE_HOSTNAME environment variable to 'localhost'.

Bug 8407818

After adding a new node to a shared Oracle database home using addNode.sh, the /etc/oratab on the newly added node gets an entry of the source database name that exists on the source node from where addNode.sh was run. The /etc/oratab file on the new node is supposed to get the database entry after the database instance is added for the new node using DBCA.

Workaround: Before invoking DBCA from the source node to add a new database instance for the new node, open the /etc/oratab file on the new node using an editor and remove the entry made for the source database name.

Bug 7388999

The `addNode.sh` script is not supported in GUI mode.

Workaround: Only use `addNode.sh` in silent mode with the required parameters.

1.38.19 Database Control Patching Known Bugs

Bug 6489313

The deployment procedures for patching Oracle Standalone Databases and Oracle RAC Databases do not run SQL scripts if operating system authentication is disabled for those database installations. The procedures fail after applying the patch.

Workaround:

As described by the instructions in the README of the patch, run the SQL scripts on the database instances that are running out of the patched Oracle homes.

1.38.20 Oracle Warehouse Builder Known Bugs

Refer to "Open Bugs and Known Issues" in *Oracle Warehouse Builder Release Notes*.

1.38.21 RMAN Known Bugs

Bug 8652023

When `CATALOG BACKUPPIECE` is used near the end of the month, the completion date for the backup piece might have an invalid date.

Workaround: Do not use `CATALOG BACKUPPIECE` near the end of the month.

Bug 8631293

The Recovery Manager (RMAN) command `DUPLICATE` might fail with `RMAN-06551` if the parameters `DB_RECOVERY_FILE_DEST` and `DB_RECOVERY_FILE_DEST_SIZE` are set in the `init.ora` file of the new database in which the `DUPLICATE` command is executing.

Workaround: Do not set the parameters `DB_RECOVERY_FILE_DEST` and `DB_RECOVERY_FILE_DEST_SIZE` in the `init.ora` file of the new database until the `DUPLICATE` command has finished executing.

Bug 8572507

When you run the RMAN `list backupset` command on a backup set that contains data file information about dropped files, an error is returned.

Workaround: Use the RMAN `list backup` command instead of `list backupset`.

Bug 8559328

Attempting a Recovery Manager (RMAN) Tablespace Point-in-Time Recovery (TSPITR) on a tablespace that has objects that depend on user-defined object types causes the Data Pump Import utility (`impdp`) to fail with `ORA-31684`.

Workaround: Ignore the failure and alter the tablespace state to `READ WRITE`. Although the tablespace state appears to be `READ ONLY` (instead of `OFFLINE`, as expected), when you alter the tablespace state to `READ WRITE`, all recovered objects are available.

2 Readme Information for Oracle Database 11*g* Release 2 (11.2.0.2)

> **Note:** If you are on Oracle Database 11*g* Release 2 (11.2.0.2), this is the Readme section that you need to read.

This section of the Readme contains the following sub-sections:

Section 2.1, "Nomenclature Changes"

Section 2.2, "Compatibility, Upgrading, Downgrading, and Installation"

Section 2.3, "Features Not Available or Restricted in 11.2.0.2"

Section 2.4, "Default Behavior Changes"

Section 2.5, "Oracle Automatic Storage Management (Oracle ASM)"

Section 2.6, "Oracle Enterprise Manager Database Control"

Section 2.7, "Database Security"

Section 2.8, "Encryption and Integrity"

Section 2.9, "Java and Web Services"

Section 2.10, "Media Management Software"

Section 2.11, "Oracle Application Express"

Section 2.12, "Oracle Client Applications"

Section 2.13, "Oracle Configuration Manager"

Section 2.14, "Oracle Data Mining"

Section 2.15, "Oracle Internet Directory"

Section 2.16, "Oracle Multimedia"

Section 2.17, "Oracle Net Services"

Section 2.18, "Oracle Real Application Clusters"

Section 2.19, "Oracle Grid Infrastructure for a Cluster"

Section 2.20, "Oracle Real Application Testing"

Section 2.21, "Oracle ODBC Driver"

Section 2.22, "Oracle OLAP"

Section 2.23, "Oracle Spatial"

Section 2.24, "Oracle SQL Developer"

Section 2.25, "Oracle Text"

Section 2.26, "Oracle Ultra Search"

Section 2.27, "Oracle Warehouse Builder"

Section 2.28, "Oracle Workflow"

Section 2.29, "Oracle XML DB"

Section 2.30, "PL/SQL"

2.1 Nomenclature Changes

Note the following nomenclature changes:

- Flash Recovery Area has been renamed to Fast Recovery Area.

- The name Oracle *inter*Media was changed to Oracle Multimedia in Oracle Database 11*g* Release 1 (11.1). The feature remains the same, only the name has changed.

2.2 Compatibility, Upgrading, Downgrading, and Installation

For late-breaking updates and best practices about preupgrade, post-upgrade, compatibility, and interoperability discussions, see Note 785351.1 on My Oracle Support (formerly Oracle*MetaLink* at `https://metalink.oracle.com/`) that links to the "Upgrade Companion" web site for Oracle Database 11*g* Release 2.

Caution: After installation is complete, do not manually remove or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running. If you remove these files, then Oracle software can encounter intermittent hangs. Oracle Grid Infrastructure for a cluster and Oracle Restart installations fail with the following error:

```
CRS-0184: Cannot communicate with the CRS daemon.
```

2.2.1 Multiple Interconnects and Oracle ACFS

If you have Oracle ACFS file systems on Oracle Grid Infrastructure for a cluster 11g release 2 (11.2.0.1), you upgrade Oracle Grid Infrastructure to 11g release 2 (11.2.0.2), and you take advantage of Redundant Interconnect Usage and add one or more additional private interfaces to the private network, then you must restart the Oracle ASM instance on each upgraded cluster member node (reference Bug 9969133).

2.2.2 Deinstallation Restrictions

The following sections describe deinstallation and deconfiguration restrictions. See Section 2.38.2, "Deinstallation Tool Known Bugs" for additional information.

2.2.2.1 Deinstall Upgraded 11.2 Oracle RAC and Oracle Grid Infrastructure for a Cluster Homes

After you deconfigure and deinstall an upgraded Oracle Database 11*g* Release 2 (11.2) Oracle RAC home and to deconfigure and deinstall an 11.2 Oracle Grid Infrastructure

for a cluster home, you must detach any pre-11.2 Oracle RAC software homes from the central Inventory (reference Bug 8666509).

Detach the pre-11.2 Oracle RAC homes from the central inventory with the following command:

```
ORACLE_HOME/oui/bin/runInstaller -detachHome ORACLE_HOME_NAME=pre-11.2_ORACLE_
HOME_NAME ORACLE_HOME=pre-11.2_ORACLE_HOME
```

2.2.2.2 Error When Running Deinstallation from an Existing Oracle Home With the -home Option
If you try to run the `deinstall` tool from the product home with the `-home` option, then the deinstallation fails with the following error message:

```
$ ./deinstall -home /scratch/user_dir/oracle/product/11.2.0/dbhome_1
Error: invalid argument -home.
Since the tool is run from inside an Oracle Home it will deinstall the home
the tool is installed in. You cannot specify another home.
```

Because the `deinstall` tool is run from within an Oracle home, the `deinstall` tool cannot be run with the `-home` option. The `deinstall` tool can only be run as `./deinstall` from within an Oracle home.

2.2.3 Time Zone File Version
Oracle Database 11*g* Release 2 (11.2) ships with time zone file versions 1 through 14. If your older database is using a time zone file version greater than 14, then you must obtain the corresponding time zone file patch for 11.2 prior to upgrading your database.

You can use `SELECT VERSION FROM V$TIMEZONE_FILE` to identify the time zone file version when upgrading your Oracle Database 10*g* or 11*g* databases. The Pre-Upgrade Information Tool (`utlu112i.sql` and its counterpart used by the Database Upgrade Assistant) provides a warning if the time zone file version is not version 11. The warning recommends running the `DBMS_DST` package to upgrade the database time zone version to the latest and to update `TIMESTAMP WITH TIME ZONE` data as well. The Pre-Upgrade Information Tool also populates three new database properties (`DST_PRIMARY_TT_VERSION`, `DST_SECONDARY_TT_VERSION`, and `DST_UPGRADE_STATE`) in `sys.database_properties` that are pertinent to the time zone version and its upgrade. `DST_PRIMARY_TT_VERSION` property records the actual time zone version in use. The other two database properties will be used when you upgrade your time zone version using `DBMS_DST` package.

Note that, in this release, you have the option to retain your current time zone version after migrating to 11.2. For example, whether your application uses any `TIMESTAMP WITH TIME ZONE` type, you do not need to run the `DBMS_DST` package to upgrade your time zone version to the latest one available.

2.2.4 Oracle ASM Rolling Upgrade
Oracle Automatic Storage Management (Oracle ASM) rolling upgrade check does not allow rolling upgrade to be done from 11.1.0.6 to any later release (reference Bug 6872001). The following message is reported in the alert log:

```
Rolling upgrade from 11.1.0.6 (instance instance-number) to 11.x.x.x is not
supported
```

`ORA-15156` is signalled by LMON which will then terminate the instance.

When trying to upgrade Oracle ASM from 11.1.0.6 to a later release of Oracle ASM, apply the patch for this bug to 11.1.0.6 instances before rolling upgrade starts. This patch can be applied to 11.1.0.6 instances in a rolling fashion.

2.2.5 UTC Time Zone Error When Upgrading From 9.2 to 11.2

When running upgrade scripts from Oracle9*i* Database Release 2 (9.2) to Oracle Database 11*g* Release 2 (11.2), you may encounter the following error:

```
ORA-27465: invalid value UTC for attribute DEFAULT_TIMEZONE
```

This error occurs because the default time zone file for release 9.2 does not have the coordinated universal time (UTC) zone. This message is produced by Oracle Scheduler that, by default, selects the UTC time zone and checks it against the default time zone file of release 9.2. This error message is expected and you can ignore it.

2.2.6 Tablespace and Fast Recovery Area Sizing

> **Note:** Fast Recovery was previously known as Flash Recovery.

The Oracle Database 11*g* Pre-Upgrade Information Utility (`utlu112i.sql`) estimates the additional space that is required in the SYSTEM tablespace and in any tablespaces associated with the components that are in the database (for example, SYSAUX, DRSYS). For a manual upgrade, be sure to run this utility on your existing database prior to upgrading.

The tablespace size estimates may be too small, especially if Oracle XML DB is installed in your database. However, to avoid potential space problems during either a manual upgrade or an upgrade using the Database Upgrade Assistant (DBUA), you can set one data file for each tablespace to AUTOEXTEND ON MAXSIZE UNLIMITED for the duration of the upgrade.

If you are using file systems for data file storage, then be sure there is adequate space in the file systems for tablespace growth during the upgrade.

If you are using a Fast Recovery Area, then check that the size available is sufficient for the redo generated during the upgrade. If the size is inadequate, then an ORA-19815 error will be written to the alert log, and the upgrade will stop until additional space is made available.

2.2.7 Memory Target for Oracle Database Instances

Running Database Configuration Assistant (DBCA) defaults to this Automatic Memory Management option. In the case of upgrade or manual database creation, MEMORY_TARGET can be specified in the initialization parameter file.

2.2.8 Upgrading a Database With Oracle Data Mining (ODM)

If you upgrade a database with the Data Mining option from 11.2.0.1 to 11.2.0.2, make sure that the DMSYS schema does not exist in your 11.2.0.1 database. If it does, you should drop the DMSYS schema and its associated objects from the database as follows:

```
SQL> CONNECT / AS SYSDBA;
SQL> DROP USER DMSYS CASCADE;
SQL> DELETE FROM SYS.EXPPKGACT$ WHERE SCHEMA = 'DMSYS';
SQL> SELECT COUNT(*) FROM DBA_SYNONYMS WHERE TABLE_OWNER = 'DMSYS';
```

If the above SQL returns non-zero rows, create and run a SQL script as shown in the following example:

```
SQL> SET HEAD OFF
SQL> SPOOL dir_path/DROP_DMSYS_SYNONYMS.SQL
SQL> SELECT 'Drop public synonym ' ||'"'||SYNONYM_NAME||'";' FROM DBA_SYNONYMS
WHERE TABLE_OWNER = 'DMSYS';
SQL> SPOOL OFF
SQL> @dir_path/DROP_DMSYS_SYNONYMS.SQL
SQL> EXIT;
```

If you upgrade a database from 10*g* to 11.2, all Data Mining metadata objects are migrated from DMSYS to SYS. After the upgrade, when you determine that there is no need to perform a downgrade, set the initialization parameter COMPATIBLE to 11.2 and drop the DMSYS schema and its associated objects as described above.

2.2.9 Compatibility with Oracle9*i* Database Release 2

Connecting the Oracle Database 11*g* Release 2 (11.2) client or server to Oracle9*i* Database release (9.2.0.4 and above) is supported. Similarly, connecting Oracle9*i* client (release 9.2.0.4 and above) to Oracle Database 11*g* Release 2 (11.2) is supported.

2.2.10 Enabling Automated Backups

While installing Oracle Database, the Specify Backup and Recovery Options screen may appear truncated if your system does not have the required fonts installed. If your system has only fixed-width fonts, then you may not be able to fully specify the required information in the Backup Job Credentials area of the screen. To work around this issue, do not select **Enable Automated Backups** on this screen. After the installation is complete, use Oracle Enterprise Manager 11*g* Database Control to enable automated backups.

2.2.11 Upgrading a Database With SQL Access Advisor Tasks

Due to internal structural changes to the SQL Access Advisor repository, a database upgrade resets all of the existing SQL Access Advisor tasks to their initial state. This action effectively deletes all of the recommendation information for tasks that have been successfully completed prior to the upgrade.

After upgrading, the recommendation information can be restored by reexecuting the existing SQL Access Advisor tasks.

2.2.12 Downgrade to Release 11.1.0.6

If you anticipate downgrading back to release 11.1.0.6, then apply the patch for Bug 7634119. This action avoids the following DBMS_XS_DATA_SECURITY_EVENTS error:

```
PLS-00306: wrong number or types of arguments in call
to 'INVALIDATE_DSD_CACHE' DBMS_XS_DATA_SECURITY_EVENTS
PL/SQL: Statement ignored
```

Apply this patch prior to running catrelod.sql.

2.2.13 Rolling Upgrade for Oracle Clusterware

In Oracle Database 11*g* Release 2 (11.2), there is a new prerequisite check during the rolling Oracle Clusterware upgrade cycle. On failure of this prerequisite check, the *Oracle Database Upgrade Guide* documents the following:

When upgrading Oracle Clusterware from Oracle Clusterware 10*g* Release 2 (10.2) to Oracle Clusterware 11*g* Release 2 (11.2), you must first apply the patchset version 10.2.0.3.0 or later to your existing Oracle Clusterware before it can be upgraded.

Follow the rolling upgrade procedure in Appendix F.5.1 of the *Oracle Grid Infrastructure Installation Guide*.

Apart from this proposed solution, there is an alternative. You can upgrade all of the nodes of the cluster simultaneously in a non-rolling manner. Choosing this solution voids the 10.2.0.3.0 patchset requirement.

2.2.14 Upgrading With Objects in the Recycle Bin

The recycle bin must be empty during an upgrade to avoid possible `ORA-00600` deadlock errors, as well as to minimize the time required to perform the upgrade (reference Bug 8632581).

To avoid this deadlock, use the `PURGE DBA_RECYCLEBIN` statement to remove items and their associated objects from the recycle bin and release their storage space prior to upgrading your database.

2.2.15 Oracle ACFS Registry May Be in an Inconsistent State After Installing or Upgrading to 11.2.0.2.0 or After An Oracle Clusterware Restart

The `ora.acfs.registry` resource may be in an inconsistent state after installing or upgrading to 11.2.0.2.0 or after an Oracle Clusterware restart (reference Bug 9876173 and Bug 9864447):

- After completing a new 11.2.0.2.0 installation, if the voting disk or the Oracle Cluster Registry (OCR) or both are on Oracle ASM, then `ora.registry.acfs` will have the following values:

  ```
  ./crsctl status resource ora.registry.acfs
   NAME=ora.registry.acfs
   TYPE=ora.registry.acfs.type
   TARGET=ONLINE, ONLINE
   STATE=ONLINE, ONLINE
  ```

  (Note that `ora.asm` will be in `STATE ONLINE`, `TARGET ONLINE`.)

  However, if the voting disk or the Oracle Cluster Registry (OCR) or both are not on Oracle ASM, then `ora.registry.acfs` will have the following values:

  ```
  ./crsctl status resource ora.registry.acfs
   NAME=ora.registry.acfs
   TYPE=ora.registry.acfs.type
   TARGET=ONLINE , ONLINE
   STATE=OFFLINE, OFFLINE
  ```

  (Note that `ora.asm` will be in `STATE OFFLINE`, `TARGET OFFLINE`.)

- After completing an upgrade from 11.2.0.1.0 to 11.2.0.2.0, `TARGET` and `STATE` for `ora.registry.acfs` will be set to the value they were in 11.2.0.1.0, prior to upgrading to 11.2.0.2.0.

- At any time, if Oracle Clusterware is restarted, `ora.registry.acfs` will have the following values:

  ```
  ./crsctl status resource ora.registry.acfs
   NAME=ora.registry.acfs
  ```

```
      TYPE=ora.registry.acfs.type
      TARGET=ONLINE , ONLINE
      STATE=ONLINE, ONLINE
```

(Note that `ora.asm` will be in `STATE ONLINE`, `TARGET ONLINE`.)

It is not necessary to take any action, however, if desired, to disable Oracle ACFS, enter the command `acfsroot unregister`, which will remove the `ora.registry` (and `ora.drivers`) resources completely.

2.2.16 Oracle ACFS File Systems Must be Manually Dismounted Prior to Upgrade, Deinstallation, or Direct Shutdown of Oracle Clusterware or Oracle ASM

Oracle ACFS file systems must be manually dismounted prior to upgrade, deinstallation, or direct shutdown of Oracle Clusterware or Oracle ASM. Also, manually dismount non-Oracle ACFS file systems which are running over Oracle ASM Dynamic Volume Manager (Oracle ADVM). This is also recommended when rebooting the system (reference Bug 8594128, Bug 9547527, and Bug 9777741).

When doing an upgrade, deinstallation, or direct shutdown of Oracle Clusterware or Oracle ASM, to ensure that Oracle Clusterware (or Oracle ASM) comes down cleanly, you must dismount all Oracle ACFS file systems manually prior to Oracle Cluserware shutdown. Use the `lsof` and `fuser` commands (Linux and UNIX) or the `handle` and `wmic` commands (Windows) to identify processes which are active on the Oracle ACFS file systems. To ensure that these processes are no longer active, dismount all Oracle ACFS file systems and issue Oracle Clusterware shutdown. Otherwise errors may be issued during Oracle Clusterware shutdown relating to activity on Oracle ACFS file systems which will stop the successful shutdown of Oracle Clusterware.

2.2.17 Data Pump Export Operations Do Not Work if the DMSYS Schema is Not Removed as Part of the Upgrade to 11.2.0.2

Data Pump Export operations do not work if the `DMSYS` schema is not removed as part of the upgrade to 11.2.0.2 (reference Bug 10007411). The reported error is similar to the following:

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
ORA-39126: Worker unexpected fatal error in
KUPW$WORKER.GET_TABLE_DATA_OBJECTS []
ORA-31642: the following SQL statement fails:
BEGIN "DMSYS"."DBMS_DM_MODEL_EXP".SCHEMA_CALLOUT(:1,0,1,'10.01.00.05.00');
END;
ORA-06512: at "SYS.DBMS_SYS_ERROR", line 86
ORA-06512: at "SYS.DBMS_METADATA", line 1245
ORA-04063: package body "DMSYS.DBMS_DM_MODEL_EXP" has errors
ORA-06508: PL/SQL: could not find program unit being called:
"DMSYS.DBMS_DM_MODEL_EXP"
```

The pre-upgrade checks for 11.2.0.2 report the action that should be taken before the upgrade:

```
The DMSYS schema exists in the database. Prior to performing an upgrade Oracle
recommends that the DMSYS schema, and its associated objects be removed from the
database. Refer to the Oracle Data Mining Administration Guide for the
instructions on how to perform this task.
```

Until this step is taken, Data Pump Export will not work.

2.2.18 catrelod.sql Fails if the Time Zone File Version Used by the Database Does Not Exist in Oracle Home

The following error is returned when `catrelod.sql` is run as part of the downgrade process if you previously installed a recent version of the time zone file and used the `DBMS_DST` PL/SQL package to upgrade `TIMESTAMP WITH TIME ZONE` data to that version (reference Bug 9803834):

```
ORA-00600: internal error code, arguments: [qcisSetPlsqlCtx:tzi init],
 [], [], [], [], [], [], [], [], [], [], []
```

See Step 2 of 'Downgrade the Database' in Chapter 6 of the *Oracle Database Upgrade Guide* for more details.

If you previously installed a recent version of the time zone file and used the `DBMS_DST` PL/SQL package to upgrade `TIMESTAMP WITH TIME ZONE` data to that version, then you must install the same version of the time zone file in the release to which you are downgrading. For example, the latest time zone files that are supplied with Oracle Database 11*g* Release 2 (11.2) are version 14. If, after the database upgrade, you had used `DBMS_DST` to upgrade the `TIMESTAMP WITH TIME ZONE` data to version 14, then install the version 14 time zone file in the release to which you are downgrading. This ensures that your `TIMESTAMP WITH TIME ZONE` data is not logically corrupted during retrieval. To find which version your database is using, query `V$TIMEZONE_FILE`.

Also see the *Oracle Database Globalization Support Guide* for more information on installing time zone files.

2.2.19 Error Downgrading from 11.2.0.2 to 10.2.0.5

During 10.2.0.5 downgrade from 11.2.0.2, `catrelod.sql` fails with the following internal error (reference Bug 9975586):

```
ORA-00600: internal error code, arguments: [17069]
```

To workaround this problem, add the following to the initialization parameter file:

```
event='logon trace name kra_options level 7'
```

If using `spfile`, issue the following command:

```
alter system set event='logon trace name kra_options level 7' scope=spfile;
```

Restart the instance and continue running the downgrade script. After downgrade, remove the event from the initialization parameter file.

2.2.20 Downgrading a Database Having Database Control Configured

Consider the following when downgrading the database while having Database Control configured (reference Bug 9922349):

1. If you are upgrading from 11.2.0.1 to 11.2.0.2 and then plan to downgrade to 11.2.0.1, you need to apply the following patches in order to downgrade Database Control as part of the database downgrade:

   - 11.2.0.1 PSU2 bundle

   - One-off patch for Bug 8795792

   - One-off patch for Bug 10013976

   Without these patches, the `emdwgrd` utility would fail with `IMPORT` (`impdp`) errors when restoring Database Control data.

2. When running `emdwgrd` on 11.2.0.1 Oracle RAC databases, you may need to pass an additional parameter, `-serviceAlias`, if you do not have system identifier (SID) aliases defined in `tnsnames.ora`. This is also needed for single instance if SID and database names are different. For example:

```
emdwgrd -save [-cluster] -sid SID [-serviceAlias tns_alias] -path save_
directory
emdwgrd -restore -tempTablespace TEMP [-cluster] -sid SID [-serviceAlias tns_
alias] -path save_directory
```

3. In the case of in-place downgrade from 11.2.0.2 to 11.2.0.1 using the same Oracle home, you do not need to run `emca -restore` before running `emdwngrd -restore`.

2.3 Features Not Available or Restricted in 11.2.0.2

The following is a list of components that are not available or are restricted in this release of Oracle Database 11*g* Release 2 (11.2.0.2):

- Patching from Database Control is deprecated and no longer available.

- Oracle Database release 11.2.0.1 upgrade to 11.2.0.2 clusterware is not supported if the 11.2.0.1 Oracle Grid Infrastructure for a cluster is installed in a non-shared Oracle home and 11.2.0.2 Oracle Grid Infrastructure for a cluster is installed in a shared Oracle home (reference Bug 10074804). Both 11.2.0.1 and 11.2.0.2 should be installed in either a shared or non-shared Oracle home.

- All Oracle Grid Infrastructure patch set upgrades must be out-of-place upgrades, in which case you install the patch set into a new Oracle Grid home (reference Bug 10210246). In-place patch set upgrades are not supported.

- Database Smart Flash Cache is supported on Solaris and Oracle Linux only.

- Using Internet Protocol Version 6 (IPv6) is not supported with the following:

  - Oracle RAC and Oracle Clusterware

  - Oracle Fail Safe

- Oracle Ultra Search has been desupported and is not shipping with Oracle Database 11*g* Release 2 (11.2).

- Downgrading from Oracle Database 11*g* Release 2 (11.2) to release 10.2.0.3 or release 10.2.0.4 is not supported with Oracle Database Vault.

- Oracle Automatic Storage Management Cluster File System (Oracle ACFS) resources are not supported for Oracle Restart configurations on all platforms. Oracle ACFS drivers must be manually unloaded and loaded; Oracle ACFS file systems must be manually unmounted and mounted (after the Oracle ASM instance is running); Oracle ACFS database home file systems can be placed into the Oracle ACFS mount registry to be mounted along with other registered Oracle ACFS file systems.

- Oracle ACFS Replication is not supported in Oracle Restart.

- Refer to Section 2.10.1 for globalization restrictions within Oracle Secure Backup.

- Refer to Section 2.14 for features that are no longer available in Oracle Data Mining.

- Security-Enhanced Linux (SELinux) is not supported on Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems.

The following sections discuss additional restrictions.

2.3.1 Database Object Names

Oracle does not recommend using quoted identifiers for database object names. These quoted identifiers are accepted by SQL*Plus, but they may not be valid when using other tools that manage database objects.

2.4 Default Behavior Changes

This section describes some of the differences in behavior between Oracle Database 11*g* Release 2 (11.2) and previous releases. The majority of the information about upgrading and downgrading is already included in the *Oracle Database Upgrade Guide*.

2.4.1 Direct Insert Behavior Change

Direct insert requires memory for every partition loaded. The memory usage is even greater if the partitions are compressed. In previous releases, a direct insert would continue to allocate memory as rows were encountered for previously untouched partitions until all of the partitions were loaded or until it ran out of memory. In this case, the insert fails.

Starting in Oracle Database 11*g* Release 2 (11.2), direct insert limits the memory that is allocated. If direct insert reaches the limit and it acquires rows for partitions that have not been loaded during the insert, then direct load stores those rows in the temporary tablespace. Once all rows have been passed in, direct load will load the rows it stored in the temporary tablespace.

2.4.2 Audit Default Behavior Changes

Default Audit behavior changes include:

- Audit filename is now prefixed with the instance name and ends with a sequence number. For example:

  instanceName_ora_pid_seqNumber.aud / *instanceName_ora_pid_seqNumber*.xml

- An existing audit file is never appended to. If an audit file already exists, the sequence number is incremented and written to *instanceName_ora_pid_ seqNumber+1*.aud.

- There is a preconfigured threshold for audit file growth. The audit file associated with an active session remains open until one of the following limits is reached:

  - After the audit record is written, the audit file size becomes 10 megabytes or more.

  - After the audit record is written, the audit file age becomes 5 days or more.

  Once one of these thresholds is reached, a new audit file with an incremented sequence number is opened for further audit records.

- There are no updates to AUD$ anymore.

  - All logoff (action# 101) audit records are written separately.

  - If an event is audited BY SESSION, then every occurrence of the event becomes a new audit record in AUD$.

2.4.3 FILE_ACCESS_ROLE Default Behavior Change

The default behavior of the CTX system parameter FILE_ACCESS_ROLE has changed (reference Bug 8360111). Customers with existing Oracle Text indexes that use the file

or URL datastore must take action to continue to use the indexes without error. The changes are as follows:

- If `FILE_ACCESS_ROLE` is null (the default), then access is not allowed. By default, users who were previously able to create indexes of this type will not be able to create these indexes after the change.

- `FILE_ACCESS_ROLE` is now checked for index synchronization and document service operations. By default, users will not be able to synchronize indexes of this type or use document service calls such as `ctx_doc.highlight` who were allowed to prior to this change.

- Only SYS will be allowed to modify `FILE_ACCESS_ROLE`. Calling `ctx_adm.set_parameter (FILE_ACESS_ROLE, role_name)` as a user other than SYS will now raise the new error:

  ```
  DRG-10764: only SYS can modify FILE_ACCESS_ROLE
  ```

- Users can set `FILE_ACCESS_ROLE` to `PUBLIC` to explicitly disable this check (which was the previous default behavior).

2.4.4 Support for Raw Devices

Oracle Universal Installer (OUI) and Database Configuration Assistant (DBCA) do not support raw devices (or block devices on Linux). However, command-line utilities such as SQL*Plus and CRSCTL do support raw or block devices.

2.4.5 Oracle Clusterware and Oracle ASM Installed Into Oracle Grid Infrastructure for a Cluster Home

Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) are installed into a single Oracle home called the *Grid_home*. This installation is referred to as the Oracle Grid Infrastructure for a cluster installation.

When upgrading, you must upgrade both Oracle Clusterware and Oracle ASM at the same time.

2.4.6 Deprecated Parameters in the Oracle ASM Instance

The following initialization parameters are deprecated in the Oracle ASM instance:

- `CLUSTER_DATABASE` parameter

  If the `INSTANCE_TYPE` is Oracle ASM and the Oracle RAC option is turned on, then you do not have to specify the `CLUSTER_DATABASE` parameter. In this case, the `CLUSTER_DATABASE` parameter defaults to `TRUE`.

- `COMPATIBLE` parameter

  Do not set the `COMPATIBLE` parameter in an Oracle ASM instance. To advance the disk group compatibility, change the `COMPATIBLE.[RDBMS|ASM|ADVM]` attributes of the disk group.

2.4.7 Fixed Views

All parameter views (for example, `V$PARAMETER`) in the Oracle ASM instance only report parameters that are relevant to the Oracle ASM instance.

2.4.8 Offset for CLOB and NCLOB APIs

Starting with Oracle Database 10*g* Release 1 (10.1), APIs that write to a `CLOB` or `NCLOB` cause error `ORA-22831` when the offset specified for the beginning of the write is not on a character boundary of the existing `LOB` data.

`LOB` APIs use UCS2 codepoint semantics for the amount and offset parameters on `CLOB`s or `NCLOB`s when the database default or national character set is multibyte. The specified offset is not on a character boundary if it points to the low (second) surrogate of a surrogate pair. In such situations, error `ORA-22831` occurs and the data is not written. Thus, this prevents the corruption of the character in the target `LOB`.

Contact Oracle Support Services to configure the database so that it does not return `ORA-22831`.

2.4.9 Use of the Append Hint Runs Out of Memory When Loading Many Partitions

Use of direct-path `INSERT` to load a large number of partitions can exceed memory limits, especially when data compression is specified (reference Bug 6749894). Starting in 11.2, the number of partitions loaded at the same time will be limited, based on the `PGA_AGGREGATE_TARGET` initialization parameter, to preserve memory. Rows that are not stored in the partitions that are currently being loaded are saved in the temporary tablespace. After all rows are loaded for the current set of partitions, other partitions are loaded from rows that are saved in the temporary tablespace.

This behavior helps prevent the direct-path `INSERT` from terminating because of insufficient memory.

2.4.10 Non-Uniform Memory Access Optimizations and Support Disabled in 11.2

With Oracle Database 11*g* Release 2 (11.2), non-uniform memory access support is disabled by default. This restriction applies to all platforms and operating systems (reference Bug 8450932).

Non-uniform memory access optimizations and support in the Oracle Database are only available for specific combinations of Oracle version, operating systems, and platforms. Work with Oracle Support Services and your hardware vendor to enable non-uniform memory access support.

2.4.11 View Changes for Advanced Compression

The `COMPRESS_FOR` column in various table views, such as `user_tables` and `dba_tables`, returns different values in 11.2 as compared to 11.1. The new values returned from a `COMPRESS_FOR` column are `BASIC` or `OLTP`. In 11.1, the value was `DIRECT LOAD ONLY` and `FOR ALL OPERATIONS`, respectively.

2.5 Oracle Automatic Storage Management (Oracle ASM)

The following sections describe information pertinent to Oracle Automatic Storage Management (Oracle ASM) in Oracle Database 11*g* Release 2 (11.2).

2.5.1 Oracle Database Data Files and Oracle ACFS

Oracle ACFS is the preferred file manager for non-database files. It is optimized for general purpose files. Oracle ACFS does not support any file that can be directly stored in Oracle ASM.

2.5.2 Oracle Homes on Oracle ACFS Supported Starting With Release 11.2

Placing Oracle homes on Oracle ACFS is supported starting with Oracle Database release 11.2 (reference Bug 10144982). Oracle ACFS can result in unexpected and inconsistent behavior if you attempt to place Oracle homes on Oracle ACFS on database versions prior to 11.2.

2.6 Oracle Enterprise Manager Database Control

The following sections describe information for Database Control in Oracle Database 11*g* Release 2 (11.2).

2.6.1 Reorganize Objects

The following applies to Reorganize Objects:

■ The database can be adversely affected if metadata changes occur between the time the reorganization script is generated and the time it is completed.

■ The database can be adversely affected if an offline reorganization is attempted while DDL is in progress against the tables being reorganized.

■ Before using the Reorganize Objects wizard or the Make Tablespace Locally Managed wizard in Oracle Enterprise Manager, Oracle recommends that you back up your database.

2.6.2 Editing Multiple Objects

Oracle Enterprise Manager only supports having one browser window open for editing an object. For example, Oracle Enterprise Manager only supports editing one tablespace at a time.

> **Note:** If a window is read-only, then you may have multiple browser windows open.

2.6.3 Large Number of Database Objects May Require Increase in Heap Size

Oracle Enterprise Manager Database Control is configured with 192 MB (32-Bit)/384 MB (64-Bit) of heap memory. However, certain Oracle Enterprise Manager Database Control functionality (for example, Change Manager) may require higher memory settings if the database contains a large number of objects.

Memory settings can be increased using the following `emctl` command:

```
emctl config dbconsole -heap_size MemorySizeValue M
```

Oracle Enterprise Manager Database Control must be restarted for the new settings to take effect.

2.6.4 New Oracle Database 11*g* Release 2 Features Supported by Oracle Enterprise Manager Database Control 11.2 Only

Management support for new features available with Oracle Database 11g Release 2 (11.2.0.2), such as Oracle RAC One Node Database and Oracle Database QoS Management, is available in Oracle Enterprise Manager Database Control 11.2.0.2 only.

2.7 Database Security

Note the following changes in Database Security.

2.7.1 Configure and Use SSL Certificates to Setup Authentication

> **Note:** This affects the security in the connection between the Oracle Clusterware and the mid-tier or JDBC client.

JDBC or Oracle Universal Connection Pool's (UCP) Oracle RAC features like Fast Connection Failover (FCF) subscribe to notifications from the Oracle Notification Service (ONS) running on the Oracle RAC nodes. The connections between the ONS server in the database tier and the notification client in the mid-tier are usually not authenticated. It is possible to configure and use SSL certificates to setup the authentication but the steps are not clearly documented.

The workaround is as follows:

1. Create an Oracle Wallet to store the SSL certificate using the `orapki` interface:

    a. `cd $ORA_CRS_HOME/opmn/conf`

    b. `mkdir sslwallet`

    c. `orapki wallet create -wallet sslwallet -auto_login`

    When prompted, provide `ONS_Wallet` as the password.

    d. `orapki wallet add -wallet sslwallet -dn "CN=ons_test,C=US" -keysize 1024 -self_signed -validity 9999 -pwd ONS_Wallet`

    e. `orapki wallet export -wallet sslwallet -dn "CN=ons_test,C=US" -cert sslwallet/cert.txt -pwd ONS_Wallet`

2. Update the ONS configuration file on all nodes in the database tier to specify the location of the wallet created in Step 1:

    a. Open the file *ORA_CRS_HOME*/opmn/conf/ons.config

    b. Add the `walletfile` parameter to the `ons.config` file:

    `walletfile=`*ORA_CRS_HOME*`/opmn/conf/sslwallet`

    c. Restart the ONS servers with the `srvctl`:

    ```
    srvctl stop nodeapps
    srvctl start nodeapps
    ```

3. If you are running a client-side ONS daemon on the mid-tier, there are two possible configurations:

    -   ONS started from OPMN (like in OracleAS 10.1.3.*x*) which uses `opmn.xml` for its configuration.

    -   ONS started standalone (like using `onsctl`), which uses `ons.config` for its configuration.

    For case (1), refer to the *OPMN Administrator's Guide* for the Oracle Application Server release. This involves modifying the `opmn.xml` file to specify the wallet location.

    For case (2), refer to the section titled *Configuration of ONS* in *Appendix B* of the *Oracle Database JDBC Developer's Guide* to modify the ons.config file to specify the wallet location created in Step 1.

4. If you are running remote ONS configuration without a client-side ONS daemon, refer to the "Remote ONS Subscription" subsection of the "Configuring ONS for Fast Connection Failover" subsection of the "Using Fast Connection Failover" section of the "Fast Connection Failover" chapter in the *Oracle Database JDBC Developer's Guide*.

Alternatively, you can specify the following string as the `setONSConfiguration` argument:

```
propertiesfile=location_of_a_Java_properties_file
```

The Java properties file should contain one or more of the ONS Java properties listed below, but at least the `oracle.ons.nodes` property. The values for these Java properties would be similar to those specified in the "Remote ONS Subscription" subsection previously noted in this step:

```
oracle.ons.nodes
oracle.ons.walletfile
oracle.ons.walletpassword
```

2.8 Encryption and Integrity

Note the following changes in the areas of encryption and integrity.

2.8.1 Encrypted Tablespaces

- You cannot encrypt an existing tablespace with an `ALTER TABLESPACE` statement. However, you can use Data Pump or SQL statements such as `CREATE TABLE AS SELECT` or `ALTER TABLE MOVE` to move existing table data into an encrypted tablespace.

- When recovering a database with encrypted tablespaces (for example, after a `SHUTDOWN ABORT` or a catastrophic error that brings down the database instance), you must open the wallet after database mount and before database open so the recovery process can decrypt data blocks and redo.

- The master encryption key for Transparent Data Encryption (TDE) tablespace encryption can leverage Hardware Security Modules (HSM) for full key management (for example, create, store, rotate, and retire). In 11.1.0.7, it was only possible to create and store the TDE tablespace encryption master key in HSM, not rotate. Master key for the TDE tablespace encryption can also be migrated from Oracle Wallet to HSM.

2.9 Java and Web Services

Note the following items when working with Java.

2.9.1 Oracle JVM

Oracle Database 11*g* Release 2 (11.2) includes a fully functional Java Virtual Machine (JVM), as well as the Java class libraries for Sun's Java Development Kit (JDK) 5.0. When combined with Oracle's JDBC and SQLJ, this release provides an enterprise class platform, Oracle JVM, for developing and deploying server-based Java applications. Refer to the Oracle JVM Readme file located at:

```
ORACLE_HOME/relnotes/readmes/README_javavm.txt
```

2.9.2 Announcing Oracle Universal Connection Pool (UCP) and Deprecation of JDBC Implicit Connection Cache in Favor of UCP

Starting with Oracle Database 11*g* Release 1 (11.1.0.7), Oracle has released the new Universal Connection Pool for JDBC. For more details, see the *Oracle Universal Connection Pool for JDBC Developer's Guide* at the following Web page:

http://www.oracle.com/technology/tech/java/sqlj_jdbc/UCP_dev_guide.pdf

Consequently, Oracle is deprecating the existing JDBC connection pool (that is, Implicit Connection Cache) that was introduced in Oracle Database 10*g* Release 1. Your applications will continue to work until formal desupport in a future release at which time a desupport notice will be posted.

Oracle encourages you to plan to adopt UCP for new applications and plan to change existing applications as indicated in *Transitioning to Oracle Universal Connection Pool (UCP)* at the following Web page:

http://www.oracle.com/technology/tech/java/sqlj_jdbc/pdf/ucp_transition_guide.pdf

UCP download and code samples are located at the following Web page:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/htdocs/ucp.html

2.9.3 JDBC

The Oracle JDBC product supports the latest Java/JDBC standards. For more details, refer to the JDBC Readme file located at:

ORACLE_HOME/relnotes/readmes/README_jdbc.txt

2.9.4 Web Services

As an alternative to Oracle Net, Oracle Database Web services provides nonconnected access to the database through standard Web services mechanisms. These include XML, SOAP, and WSDL, and can turn the database into a Web services provider. Similarly, the database itself can act as a Web service consumer and run external Web services. Important features of Web services include:

- A JAX-RPC based SOAP Client library supports invocation of external Web services from within the database, and applies the power of SQL to the results.

- **Web Services Call-In:** Deploying a JPublisher-generated Java class against Oracle Application Server 10*g* enables you to run database operations such as Java and PL/SQL procedures and packages, SQL queries, and DML operations.

- **Web Services Call-Out:** Deploying a JPublisher-generated Web services client from a WSDL and its PL/SQL wrapper supports invocation of external Web services from within the database.

2.10 Media Management Software

For environments that consist of a single server, Oracle offers Oracle Secure Backup Express to back up your Oracle Database and other critical Oracle infrastructure to tape. Oracle Secure Backup is fully integrated with Recovery Manager (RMAN) to provide data protection services. For larger environments, Oracle Secure Backup is available as a separately licensable product to back up many database servers and file systems to tape. Oracle Secure Backup release 10.3 is shipping with this Oracle Database 11*g* Release 2 (11.2). For more information on Oracle Secure Backup, refer to

http://www.oracle.com/goto/osb/

2.10.1 Globalization Restrictions Within Oracle Secure Backup

The following globalization restrictions apply to Oracle Secure Backup:

- The Oracle Secure Backup Web Tool and command line interface are available in English only, and are not globalized. All messages and documentation are in English.

- Oracle Secure Backup does not support file names or RMAN backup names that are encoded in character sets that do not support null byte termination, such as Unicode UTF-16. Note that this restriction affects file names, not backup contents. Oracle Secure Backup can back up Oracle databases in any character set.

2.11 Oracle Application Express

Note the following items when working with Oracle Application Express.

To learn more about Oracle Application Express, refer to the *Oracle Application Express Release Notes* and the *Oracle Application Express Installation Guide.*

2.11.1 Checking for the Most Current Release

Oracle Application Express is released more frequently than Oracle Database. To view additional information about the most current release, refer to:

http://www.oracle.com/technetwork/developer-tools/apex/overview/index.html

2.11.2 Configuring Oracle HTTP Server

When configuring Oracle HTTP Server for Oracle Application Express in a new installation, the database user APEX_PUBLIC_USER must be an unlocked account. To unlock the account for database user APEX_PUBLIC_USER, complete the following:

1. Start SQL*Plus and connect as SYS to the database where Oracle Application Express is installed. For example:

   ```
   $ ORACLE_HOME/bin/sqlplus
   SQL> CONNECT SYS as SYSDBA
   Enter password: SYS_password
   ```

2. Run the following command:

   ```
   SQL> ALTER USER APEX_PUBLIC_USER ACCOUNT UNLOCK
   ```

2.11.3 Database Compatible Parameter and Enabling Network Services

To run the examples in the Oracle Application Express Post-Installation tasks in the section titled "Enabling Network Services in Oracle Database 11*g*" in the *Oracle Database Installation Guide*, the COMPATIBLE initialization parameter of the database must be set to at least 11.1.0.0.0. By default, an Oracle Database 11*g* database will already have the parameter set properly, but a database upgraded to 11*g* from a prior version may not.

Refer to the section titled "Creating and Configuring an Oracle Database" in the *Oracle Database Administrator's Guide* for information about changing database initialization parameters.

2.12 Oracle Client Applications

Oracle Client 11*g* contains advanced features for diagnosing issues, including the ability to dump diagnostic information when important errors are detected. By default, these dumps are restricted to a small subset of available information to ensure that application data is not dumped. However, in many installations, secure locations for dump files may be configured, ensuring the privacy of such logs. In such cases, it is recommended that you turn on full dumps; this can greatly speed the resolution of issues. Enable full dumps by adding the following line to the sqlnet.ora file that is used by your Oracle client installation:

```
DIAG_RESTRICTED=false
```

To verify that diagnosability features are working correctly, take the following steps:

1. Upgrade your application to use Oracle Database 11*g* client libraries.

2. Start your application.

3. Check the file sqlnet.log in your application's TNS_ADMIN directory for error messages indicating that diagnosability could not be started (normally this is due to invalid directory names or permissions).

Refer to the *Oracle Call Interface Programmer's Guide* for details.

2.13 Oracle Configuration Manager

Note the following for Oracle Configuration Manager.

2.13.1 cron Configuration Issue

If you are denied access to cron, then the configuration of Oracle Configuration Manager fails with the following error:

```
ORACLE_HOME/ccr/bin/setupCCR
** Installing base package **
Deploying core - Version 10.2.5.0.0
Error encountered in package deployment.
```

After the installation is complete, set the environment variable CCR_DISABLE_CRON_ENTRY to TRUE and try the configuration of Oracle Configuration Manager again using the following command:

```
ORACLE_HOME/ccr/bin/setupCCR
```


2.14 Oracle Data Mining

Oracle Data Mining scoring functions in Oracle Database 11*g* Release 2 are also available in Oracle Exadata Storage Server Software. Scoring capabilities in the storage layer permit very large data sets to be mined quickly, thus further increasing the competitive advantage already gained from Oracle in-database analytics. For information about Oracle Exadata Storage Server Software, see http://www.oracle.com/technology/products/bi/db/exadata/index.html.

The Data Mining Option, as an embedded feature of the database, is automatically installed with the Oracle Enterprise Edition Database. When installing the database with the Data Mining Option, choose the Data Warehouse configuration type for the most appropriate default initialization parameters.

In Oracle Database 11*g*, Data Mining models are implemented as data dictionary objects in the `SYS` schema. The `DMSYS` schema no longer exists.

Data Mining users must have the `CREATE MINING MODEL` privilege to create mining models in their own schema. Additional privileges are required for other data mining activities, as described in the *Oracle Data Mining Administrator's Guide*.

New data dictionary views for Oracle Data Mining were introduced in Oracle Database 11*g* Release 1 (11.1):

- `USER/ALL/DBA_MINING_MODELS`

- `USER/ALL/DBA_MINING_MODEL_ATTRIBUTES`

- `USER/ALL/DBA_MINING_MODEL_SETTINGS`

Demo programs that illustrate the Data Mining APIs (PL/SQL and Java) are installed with Oracle Database Examples. Instructions are in the *Oracle Data Mining Administrator's Guide*.

The Oracle Data Mining Scoring Engine Option, a separately installed database option in Oracle Database 10*g*, is not available in Oracle Database 11*g*. All functionality of the Data Mining Scoring Engine Option is offered in the Data Mining Option.

The Basic Local Alignment Search Tool (BLAST), previously supported by Oracle Data Mining, is not available in Oracle 11*g*.

2.15 Oracle Internet Directory

The Oracle Internet Directory product ships only with Oracle Application Server, not the Oracle Database 11*g* Release 1 (11.1) product set. The following information is included because Oracle Network functionality may use Oracle Internet Directory. Many of the administrative activities for Oracle Internet Directory have been consolidated into a single tool, Oracle Internet Directory Configuration Assistant (OIDCA). OIDCA should be used with Enterprise User Security and Network Names features under these conditions:

1. Enterprise User Security

   - Enterprise User Security only works with Identity Management Realms in this release. You must convert Oracle Contexts used in prior releases to Identity Management Realms using the OIDCA tool.

   - Use OIDCA when creating or updating the `ldap.ora` configuration file for discovering the Oracle Internet Directory server in the environment. When created with OIDCA, `ldap.ora` is located in the `ORACLE_HOME`/ldap/admin directory on Linux and UNIX operating systems and in the `ORACLE_HOME`\ldap\admin directory on Windows operating systems.

2. Network Names

   - Use OIDCA when creating, upgrading, and deleting Oracle Contexts.

   - Use OIDCA when converting an Oracle Context from an earlier release to an Identity Management Realm.

   - Use OIDCA when setting up the `ldap.ora` configuration file for discovering the Oracle Internet Directory server in the environment.

Note the following items when working with Oracle Internet Directory.

2.15.1 Using the Oracle Internet Directory Configuration Assistant

The Oracle Internet Directory Configuration Assistant (OIDCA) enables you to create, upgrade, and delete an Oracle Context, configure the file ldap.ora, and convert an Oracle Context to an Identity Management Realm.

The OIDCA syntax is:

```
oidca oidhost=host
      nonsslport=port |
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      propfile=properties file
```

To see the usage of OIDCA, enter oidca -help at the command prompt.

2.15.2 Creating an Oracle Context

The following syntax is used to create an Oracle Context in OIDCA; the parameters are described in the subsequent table.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      mode=CREATECTX
      contextdn=OracleContext DN
```

Parameters	Description
oidhost	OID server host; if not specified, default is localhost
nonsslport	OID server port; if not specified, default is 389
sslport	OID SSL port; if not specified, default is 636
dn	OID user, such as cn=orcladmin
pwd	OID user password
mode	Mode of the OIDCA; set to CREATECTX
contextdn	DN under which OracleContext must be created, such as dc=acme, dc=com

Note the following points:

- The contextdn must exist for this operation to be successful.

- This valid DN should not exist in OID: "cn=oraclecontext,dc=acme, dc=com".

- This valid DN must exist in OID: "dc=acme,dc=com".

- The parameters mode and contextdn can also be passed as a properties file.

- Specify the parameter nonsslport=port if you want to perform the operation using non-SSL mode.

- Specify the parameter sslport=sslport if you want to perform the operation using SSL mode.

- Either the nonsslport or the sslport parameter must be specified, but not both.

Functionality

1. The OIDCA verifies that `contextdn` has a valid DN syntax and that the entry exists in Oracle Internet Directory. Note that the OIDCA cannot create a root `OracleContext` explicitly. If there is no root Oracle Context, then OIDCA exits with an error.

2. If DN exists, then OIDCA verifies that the Oracle Context already exists.

   - If the Oracle Context already exists and is up-to-date, then OIDCA exits with the message `Oracle Context already exists and is up to date`.

   - If the Oracle Context already exists, but it is an older version, then OIDCA exits with the message `Oracle Context already exists and is of an older version`.

   - If the Oracle Context does not exist, then OIDCA creates the Oracle Context under this DN.

2.15.3 Upgrading an Oracle Context

To upgrade an `OracleContext` instance, use the following syntax; the parameters are listed in the subsequent table.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      mode=UPGRADECTX
      contextdn=OracleContext DN
```

Parameters	Description
oidhost	OID server host; if not specified, default is `localhost`
nonsslport	OID server port; if not specified, default is `389`
sslport	OID SSL port; if not specified, default is `636`
dn	OID user, such as `cn=orcladmin`
pwd	OID user password
mode	Mode of the OIDCA; always set to `UPGRADECTX`
contextdn	DN under which `OracleContext` must be created, such as `dc=acme, dc=com`

Note the following points:

- The `contextdn` must contain an `OracleContext` for this operation to be successful.

- The DNs `"cn=oraclecontext,dc=acme,dc=com"` and `"dc=acme,dc=com"` are both valid.

- The parameters `mode` and `contextdn` can also be passed as a properties file.

- Specify the parameter `nonsslport=port` if you want to perform the operation using a non-SSL mode.

- Specify the parameter `sslport=sslport` if you want to perform the operation using SSL mode.

- Either the `nonsslport` or the `sslport` parameter must be specified, but not both.

Functionality

1. OIDCA verifies that the `contextdn` has valid DN syntax and that `OracleContext` exists in Oracle Internet Directory. OIDCA cannot upgrade a root `OracleContext` explicitly. If there is no root `OracleContext`, then OIDCA sends an error message.

2. If `OracleContext` exists under `contextdn`,

   - The OIDCA checks if the `OracleContext` belongs to a realm, in which case it exits with the appropriate message. Note that `OracleContext` instances that belong to a realm cannot be upgraded.

   - The OIDCA verifies that the `OracleContext` is up-to-date, then exits with the message `Oracle Context already exists and is up to date`.

   - If the `OracleContext` is not up-to-date, then the OIDCA upgrades the `OracleContext` under this DN.

2.15.4 Deleting an Oracle Context

To delete an `OracleContext`, use the following syntax; the parameters are listed in the subsequent table.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      mode=DELETECTX
      contextdn=OracleContext DN
```

Parameters	Description
oidhost	OID server host; if not specified, default is `localhost`
nonsslport	OID server port; if not specified, default is `389`
sslport	OID SSL port; if not specified, default is `636`
dn	OID user, such as `cn=orcladmin`
pwd	OID user password
mode	Mode of the OIDCA; always set to `DELETECTX`
contextdn	DN under which OracleContext must be created, such as `dc=acme, dc=com`

Note the following points:

- The `contextdn` must contain an `OracleContext` for this operation to be successful.

- The DNs `"cn=oraclecontext, dc=acme,dc=com"` and `"dc=acme,dc=com"` are both valid.

- The parameters `mode` and `contextdn` can also be passed as a properties file.

- Specify the parameter `nonsslport=port` if you want to perform the operation using a non-SSL mode.

- Specify the parameter `sslport=sslport` if you want to perform the operation using SSL mode.

- Either the `nonsslport` or the `sslport` parameter must be specified, but not both.

Functionality

1. OIDCA verifies that the `contextdn` has valid DN syntax and that `OracleContext` exists in Oracle Internet Directory.

2. If `OracleContext` exists under `contextdn`,

    - The OIDCA checks if the `OracleContext` belongs to a realm, in which case it exits with the appropriate message. Note that `OracleContext` instances that belong to a realm cannot be deleted.

    - If `OracleContext` does not belong to a realm, then OIDCA deletes it.

2.15.5 Configuring the File ldap.ora

To configure the file `ldap.ora`, use the following syntax; the parameters are listed in the subsequent table.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      adminctx=Administrative context
      mode=LDAPORA
      dirtype=OID or AD
      -update
```

Parameters	Description
oidhost	OID server host; if not specified, default is `localhost`.
nonsslport	OID server port; determined using discovery APIs.
sslport	OID SSL port; determined using discovery APIs.
mode	Mode of the OIDCA; always set to `LDAPORA`.
dirtype	Directory type; possible values are `OID` and `AD`; mandatory attribute.
adminctx	Default administrative context, such as `dc=acme,dc=com`. If not specified, then determined using discovery.
-update	If this flag is specified, then overwrite existing `ldap.ora`; if not, then create `ldap.ora` only if it does not already exist.

Note the following points:

- Either the non-SSL or the SSL port must be specified. The other port is discovered.

- The parameters `mode`, `dirtype`, and `adminctx` can also be passed in within a properties file.

Functionality

1. Using the Discovery API, the OIDCA determines all the parameters not specified on the command line.

2. The OIDCA checks for the `ldap.ora` location using Discovery APIs.

    - If `ldap.ora` exists and the `-update` parameter is not specified, then exit with message `ldap.ora exists`.

    - If `ldap.ora` exists and the `-update` parameter is not specified, then update the existing `ldap.ora` using Discovery API.

- If `ldap.ora` does not exist, then create a new `ldap.ora` file in a location in the following order:

  ```
  LDAP_ADMIN
  ORACLE_HOME/ldap/admin
  ```

2.15.6 Converting an Oracle Context to an Identity Management Realm

Oracle Database 10*g* entries must be stored in Oracle Internet Directory release 9.0.4 server. An Identity Management Realm release 9.0.4 is also required for Enterprise User Security, a feature of the Oracle Database 10*g*.

To convert an existing `OracleContext` to an Identity Management Realm, use the following syntax. The parameters are listed in the subsequent table. Note that the root of the `OracleContext` object is not converted.

```
oidca oidhost=host
      nonsslport=port
      sslport=SSL Port
      dn=binddn
      pwd=bindpwd
      mode=CTXTOIMR
      contextdn=OracleContext DN
```

Parameters	Description
`oidhost`	OID server host; default is `localhost`
`nonsslport`	OID server port; default is `389`
`sslport`	OID SSL port; default is `636`
`dn`	OID user, such as `cn=orcladmin`
`pwd`	OID user password
`mode`	Mode of the OIDCA; always set to `CTXTOIMR`
`contextdn`	DN under which `OracleContext` must be created, such as `dc=acme`, `dc=com`

Note the following points:

- The `OracleContext` must exist under the specified `contextdn`.

- The DNs "`cn=oraclecontext,dc=acme,dc=com`" and "`dc=acme, dc=com`" are both valid.

- The parameters `mode` and `contextdn` can also be passed in a properties file.

- Specify the parameter `nonsslport=port` if you want to perform the operation using a non-SSL mode.

- Specify the parameter `sslport=sslport` if you want to perform the operation using SSL mode.

- Either the `nonsslport` or the `sslport` parameter must be specified, but not both.

Functionality

1. The OIDCA checks if `contextdn` has valid DN syntax, and if it contains a valid `OracleContext`.

2. If `OracleContext` exists under `contextdn`,

- The OIDCA checks if the `OracleContext` belongs to a realm. If it does, then it exits with an appropriate error message.

- If `OracleContext` does not belong to a realm, then OIDCA upgrades it to the latest version, and converts it to a realm.

Note also:

- If the nickname attribute is not `cn`, then configure it as a user configuration attribute using the Oracle Internet Directory Self-Service Console.

- If you want to use the Oracle Internet Directory Self-Service Console to manage the users and groups in the converted realm, then you must set up the administrative privileges appropriately. For details, refer to the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

2.16 Oracle Multimedia

The name Oracle *inter*Media was changed to Oracle Multimedia in Oracle Database 11*g* Release 1 (11.1). The feature remains the same, only the name has changed. References to Oracle *inter*Media were replaced with Oracle Multimedia, however some references to Oracle *inter*Media or *inter*Media may still appear in graphical user interfaces, code examples, and related documents in the Oracle Database documentation library for 11*g* Release 2 (11.2).

With Oracle Database 11*g* Release 2 (11.2), the Oracle Multimedia DICOM feature has been made available in Oracle Database Standard Edition. It was previously only available in Oracle Database Enterprise Edition.

For additional information, refer to the Oracle Multimedia Readme file located at:

```
ORACLE_HOME/ord/im/admin/README.txt
```

2.17 Oracle Net Services

Oracle Database provides limited support for Internet Protocol Version 6 (IPv6) addressing and connectivity. For details, see *Oracle Database Net Services Administrator's Guide*.

2.18 Oracle Real Application Clusters

Note the following items when working with Oracle RAC.

2.18.1 Moving root Owned Binaries that need setuid to Local Nodes from NFS

If you install an Oracle RAC database into a shared Oracle home on an NFS device, then you must copy the ORADISM binary (`oradism`) into a local directory on each node (reference Bug 7210614).

To move `oradism`, take the following steps:

1. Copy the `ORACLE_HOME`/bin/oradism binary to an identical directory path on all cluster nodes. The path (for example, `/u01/local/bin` in the example in Step 2) must be local and not NFS. For example:

   ```
   cp -a ORACLE_HOME/bin/oradism /u01/local/bin
   ```

2. Run the following commands, as the root user, to set ownership and permissions of the `oradism` executable:

```
$ chown root /u01/local/bin/oradism
$ chmod 4750 /u01/local/bin/oradism
```

3. Create a symbolic link from the NFS shared home to the local `oradism` directory path. This needs to be done from one node only. Each node can then reference its own `oradism` using the `symlink` from the shared Oracle home. For example:

```
$ cd /nfs/app/oracle/product/11.2.0/db_1/bin
$ rm -f oradism
$ ln -s /u01/local/bin/oradism oradism
```

4. If the Oracle home is an Oracle Database home directory, then repeat steps 1-3 for other binaries such as `extjob`, `jssu`, `nmb`, `nmhs` and `nmo`. You do not need to perform this step if the Oracle home is an Oracle Grid Infrastructure home directory.

2.18.2 Creating Server Pools

> **Caution:** By default, any named user may create a server pool. To restrict the operating system users that have this privilege, Oracle strongly recommends that you add specific users to the CRS Administrators list.
>
> See also *Oracle Clusterware Administration and Deployment Guide* for more information about adding users to the CRS Administrators list.

Create an initial, user-defined server pool immediately after the Oracle Grid Infrastructure for a cluster installation if you plan to use a policy-managed Oracle RAC database. As the Grid software owner, use the following SRVCTL command located in the path *Grid_home*/bin:

```
srvctl add srvpool -g srvpool_name -u max
```

2.18.3 Using the Same Cluster for Administrator-Managed and Policy-Managed Databases

If you are creating an administrator-managed database on a cluster that already hosts policy-managed databases, then you must carefully select the nodes for the administrator-managed database (reference Bug 10027250). This is because the nodes that you select for an administrator-managed database that are in policy-managed server pools will be moved into the Generic server pool as part of this process. If you select nodes that already run other policy-managed database instances, then DBCA prompts you with a message that lists the instances and services that will be shut down when DBCA creates the administrator-managed database. If you select the `Yes` button on the dialog box when DBCA asks "`Do you want to continue?`," then your policy-managed database instances and services will be shut down as a result of the administrator-managed database creation process.

Note: This is also true if you use the `srvctl add instance` command, which gives a similar error message indicating that the databases would be shut down. If you also use the force option (`-f`) with the `srvctl add instance` command, then this is the same as choosing `Yes` on the DBCA dialog. Doing this shuts down any policy-managed databases that are running on the node before moving the node into the Generic server pool.

2.18.4 Adding a Node in a Cluster

When adding a node in a cluster running a policy-managed database, Oracle Clusterware tries to start the new instance before the cloning procedure completes. The following steps should be used to add the node:

1. Run `addNode` for the Oracle Grid Infrastructure for a cluster to add the new node. Do not run the root scripts when prompted; you will run them later.

2. Install the Oracle RAC database software using a software-only installation. Ensure Oracle is linked with the Oracle RAC option.

3. Complete the root script actions for the Database home.

4. Complete the root scripts action for the Oracle Clusterware home and then finish the installation.

2.19 Oracle Grid Infrastructure for a Cluster

Note the following items when working with Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM), which are installed with an Oracle Grid Infrastructure for a cluster installation.

2.19.1 In-Place Patch Set Upgrades Are Not Supported

All Oracle Grid Infrastructure patch set upgrades must be out-of-place upgrades, in which case you install the patch set into a new Oracle Grid home (reference Bug 10210246). In-place patch set upgrades are not supported.

2.19.2 Changing Host Names

Avoid changing host names after you complete the Oracle Grid Infrastructure for a cluster installation. Nodes with changed host names must be deleted from the cluster and added back with the new name.

2.19.3 Oracle ACFS and Oracle Clusterware Stack Shut Down

When attempting to shut down Oracle Clusterware, the Oracle Clusterware stack may report that it did not successfully stop on selected nodes (reference Bug 8651848). If the database home is on Oracle ACFS, then you may receive the following error:

```
CRS-5014: Agent orarootagent.bin timed out starting process acfsmount for action
```

This error can be ignored.

Alternatively, the Oracle Clusterware stack may report that it did not successfully stop on selected nodes due to the inability to shut down the Oracle ACFS resources. If this occurs, take the following steps:

- Ensure that all file system activity to Oracle ACFS mount points is quiesced by shutting down programs or processes and retry the shutdown.

- If the `ora.registry.acfs` resource check function times out, or the resource exhibits a state of `UNKNOWN` or `INTERMEDIATE`, then this may indicate an inability to access the Oracle Cluster Registry (OCR). The most common cause of this is a network failure. The commands `acfsutil registry` and `ocrcheck` may give you a better indicator of the specific error. Clear this error and attempt to stop Oracle Clusterware again.

2.20 Oracle Real Application Testing

The following sections describe Oracle Real Application Testing issues.

2.20.1 Reading SQL Trace Files to Construct a SQL Tuning Set

Oracle Database supports the added functionality to read SQL trace files from Oracle
Database 9*i* to construct a SQL tuning set that can be used as an input source for SQL
Performance Analyzer. After being constructed, you can use SQL Performance
Analyzer to execute the SQL tuning set on Oracle Database 10*g* Release 2 remotely
over a database link. This functionality is provided so that you can use the Oracle Real
Application Testing option to test the effect on SQL response time of a database
upgrade from Oracle Database 9*i* to Oracle Database 10*g* Release 2. For details, refer to
the *Oracle Database Real Application Testing Addendum*.

2.21 Oracle ODBC Driver

The Oracle ODBC Driver Readme file is located at:

ORACLE_HOME/odbc/html/ODBCRelnotesUS.htm

2.22 Oracle OLAP

Consider the following when working with Oracle OLAP:

1. SQL Aggregation Management is a group of PL/SQL subprograms in DBMS_CUBE
   that support the rapid deployment of cube materialized views from existing
   relational materialized views. Cube materialized views are cubes that have been
   enhanced to use the automatic refresh and query rewrite features of Oracle
   Database. A single cube materialized view can replace many of the relational
   materialized views of summaries on a fact table, providing uniform response time
   to all summary data.

   In the process of creating the cube materialized views, DBMS_CUBE also creates a
   fully functional analytic workspace including a cube and the cube dimensions. The
   cube stores the data for a cube materialized view, instead of the table that stores
   the data for a relational materialized view. A cube can also support a wide range
   of analytic functions that enhance the database with information-rich content.

2. Enhanced logging capabilities are provided through the new PL/SQL package
   DBMS_CUBE_LOG that offers four different logs. Both Analytic Workspace Manager
   users and database administrators can view the progress of an analytic workspace
   build in CUBE_BUILD_LOG, find load errors in CUBE_REJECTED_RECORDS, check the
   validity of dimension hierarchies in CUBE_DIMENSION_COMPILE, and follow OLAP
   engine events in CUBE_OPERATIONS_LOG.

3. The OLAP Option, as an embedded feature of the database, is automatically
   installed with the Oracle Enterprise Edition Database. When installing the
   database with the OLAP Option, choose the Data Warehouse configuration type
   for the most appropriate default initialization parameters.

4. Analytic Workspace Manager 11.2.0 should be used when running Oracle
   Database 11*g* Release 2 in either Oracle Database 10*g* or Oracle Database 11*g*
   compatibility modes. An Oracle Database 10*g* format analytic workspace may be
   created by choosing the Oracle Database 10*g* cube type when creating a new
   analytic workspace.

5. OLAP 10*g* clients connecting to 10*g* style analytic workspaces in an Oracle Database 11*g* Release 2 database instance should upgrade their Java to version 1.5, as well as upgrade their OLAP API version to 11.2.0.

6. Oracle Business Intelligence Beans 10*g* and Oracle Discoverer for OLAP 10*g* can be used with 10*g* format analytic workspaces in Oracle Database 11*g*. They cannot be used with 11*g* format analytic workspaces. To upgrade the OLAP API client JARs, include these new JARs in the Java class search list before the Intelligence Beans or Discoverer JAR files or both.

7. The OLAPSYS schema and the CWM metadata contained within the schema are deprecated in Oracle Database 11*g* Release 2. The OLAP API can still query relational tables (ROLAP mode) using data dictionary metadata. For more information, refer to My Oracle Support Tech Note 445311.1.

2.23 Oracle Spatial

The Oracle Spatial readme file supplements the information in the following manuals: *Oracle Spatial Developer's Guide*, *Oracle Spatial Topology and Network Data Models Developer's Guide*, and *Oracle Spatial GeoRaster Developer's Guide*. The Oracle Spatial readme file is located at:

ORACLE_HOME/md/doc/README.txt

2.24 Oracle SQL Developer

The Oracle SQL Developer readme file is located at:

ORACLE_HOME/sqldeveloper/readme.html

2.25 Oracle Text

Note the following items when working with Oracle Text. You should also check entries for the *Oracle Text Application Developer's Guide* in the Documentation Addendum.

2.25.1 Query Compatible Syntax Values

Chapter 6, titled "CTX_CLS Package", in the *Oracle Text Reference* states under the section titled "Query Compatible Syntax" that the *doc_id* column is a number. This is correct; however, the values stored in this column must be in the range `0-4294967295`. The values must be stored in an unsigned 32-bit value. This range also pertains to `catid`, `catdocid`, and `rescatid`.

2.25.2 Oracle Text Supplied Knowledge Bases

An Oracle Text knowledge base is a hierarchical tree of concepts used for theme indexing, `ABOUT` queries, and deriving themes for document services. The following Oracle Text services require that a knowledge base be installed:

- Index creation using a `BASIC_LEXER` preference where `INDEX_THEMES=YES`

- `SYNC`ing of an index where `INDEX_THEMES=YES`

- `CTX_DOC.THEMES`

- `CTX_DOC.POLICY_THEMES`

- `CTX_DOC.GIST`

- `CTX_DOC.POLICY_GIST`

- `CTX_QUERY.HFEEDBACK`

- `CTX_QUERY.EXPLAIN`, if using `ABOUT` or `THEMES` with `TRANSFORM`

- `CTX_DOC.SNIPPET` (if using the `ABOUT` operator)

- `CTX_DOC.POLICY_SNIPPET` (if using the `ABOUT` operator)

- `CONTAINS` queries that use `ABOUT` or `THEMES` with `TRANSFORM`

- The Knowledge Base Extension Compiler, `ctxkbtc`

- Clustering and classification services, if themes are specified

If you plan to use any of these Oracle Text features, then you should install the supplied knowledge bases, English and French, from the Oracle Database Examples media, available for download on OTN.

Note that you can extend the supplied knowledge bases, or create your own knowledge bases, possibly in languages other than English and French. For more information about creating and extending knowledge bases, refer to the *Oracle Text Reference*.

For information about how to install products from the Oracle Database Examples media, refer to the *Oracle Database Examples Installation Guide* that is specific to your platform.

Supplied Knowledge Bases and Upgrades

Because the supplied knowledge bases are contained on the Oracle Database Examples media, they are not immediately available after an upgrade to Oracle Database 11*g* Release 1 (11.1). Oracle Text features that depend on the supplied knowledge bases available before the upgrade will not function after the upgrade, so you must install the supplied knowledge bases from the Oracle Database Examples media.

After an upgrade, you must regenerate all user extensions to the supplied knowledge bases. These changes affect all databases installed in the given `ORACLE_HOME`.

For more information on upgrading Oracle Text and supplied knowledge bases, refer to the *Oracle Database Upgrade Guide*, Chapter 4, "After Upgrading a Database", section "Upgrading Oracle Text". The *Oracle Text Application Developer's Guide* contains both general instructions for upgrading from previous releases of Oracle Text and information on supplied knowledge bases.

2.25.3 Document Filtering Support

Starting with version 11.1.0.7, Oracle Text uses Oracle Outside In HTML Export for document filtering instead of filtering technology licensed from Autonomy, Inc. With the migration to Oracle Outside In HTML Export, several document formats are no longer supported. Refer to Appendix B of the *Oracle Text Reference* for a complete list of filter-supported document formats as well as for a list of formats that are no longer supported in 11.1.0.7.

2.26 Oracle Ultra Search

Oracle Ultra Search has been desupported and is not shipping with Oracle Database 11*g* Release 2 (11.2). In March of 2006, Oracle launched Oracle Secure Enterprise Search (SES). Oracle SES is a faster and more secure product built on the technologies

originally developed for Ultra Search. Customers using Ultra Search are strongly recommended to migrate to Oracle SES.

> **See Also:** The Oracle Secure Enterprise Search page on the Oracle Technology Network at:
>
> http://www.oracle.com/technology/products/oses/index.html

2.27 Oracle Warehouse Builder

For additional information about Oracle Warehouse Builder (OWB) in Oracle Database 11*g* Release 2 (11.2), refer to the *Oracle Warehouse Builder Release Notes.*

2.28 Oracle Workflow

Oracle Workflow is not shipping with Oracle Database 11*g* Release 2 (11.2).

2.29 Oracle XML DB

The following features are not supported with Oracle XML DB:

- Flashback Archive

- Editioning Views

- Exadata Hybrid Columnar Compression (EHCC)

- SecureFiles LOB Encryption

- Oracle Label Security (OLS) with a hybrid structured and unstructured XMLIndex on the same XML document.

2.29.1 Change in Semantics of xdb:defaultTable Annotation

There is a change in behavior in the semantics of `xdb:defaultTable` annotation while registering Oracle XML DB schemas in 11.2 as compared to 11.1 (reference Bug 7646934). If you specify `xdb:defaultTable="MY_TAB"` without specifying `xdb:sqlInline="false"`, Oracle XML DB creates the table as requested and implicitly marks it as an out-of-line table. This behavior is different from 11.1 where the `defaultTable` annotation was ignored when the `sqlInline` setting was missing.

2.29.2 VARRAY Storage Default Change

In Oracle Database 11*g* Release 1 (11.1), the default value for `xdb:storeVarrayAsTable` changed from `FALSE` to `TRUE` for XMLType object-relational storage. This default applied to the default table, but not when creating XMLType object-relational tables and columns after the schema registration (reference Bug 6858659). In Oracle Database 11*g* Release 2 (11.2), all `VARRAY` data elements are created as tables by default. This provides a significant performance increase at query time. In addition, note the following:

- Tables created prior to 11.2 are not affected by this. The upgrade process retains storage parameters. This only affects tables created in 11.2 or later.

- You can retain the pre-11.2 default of `VARRAY` storage as LOBs if you have small `VARRAY` data elements and you read and or write the full `VARRAY` all at once. You have two options to revert to the pre-11.2 behavior:

- Re-register the schema with `xdb:storeVarrayAsTable=FALSE`. This affects the default and non-default tables.

- Or, when creating the table (for non default tables), you can use the `STORE ALL VARRAYS AS LOBS` clause to override the default for all `VARRAY` data elements in the XMLType. This clause can only be used during table creation. It will return an error if used in the `table_props` at schema registration time.

- For schemas registered prior to 11.2 (when the default storage for `VARRAY` data elements was `LOB`), you can use `STORE ALL VARRAYS AS TABLES` clause to override the default for all `VARRAY` data elements in the XMLType.

2.30 PL/SQL

Note the following when working with PL/SQL.

2.30.1 Secure Database By Default Changes

The `PUBLIC EXECUTE` privilege for the `DBMS_RANDOM` PL/SQL package will be deprecated in future Oracle Database releases (reference Bug 7591837). Users who need to run this PL/SQL package should be given explicit `EXECUTE` privileges.

2.30.2 DBMS_SCHEDULER PL/SQL Package

Consider the following when working with the `DBMS_SCHEDULER` PL/SQL package.

2.30.2.1 Event-Based Jobs The Oracle Scheduler supports event-based jobs. These are jobs that are started when a message is enqueued into a database queue. When the database queue used for this purpose is a secure queue, the `QUEUE_SPEC` attribute of event-based jobs and schedules is a pair of the form *queue_name, agent_name*. The reason an agent has to be specified is to help the scheduler determine which messages the user can see. In the case of nonsecure queues, if the user has dequeue privileges on a queue, then he or she can dequeue any agent that is subscribed to the queue. This means that he or she can see all of the messages in the queue. In the case of secure queues, the user has to be granted privileges to not only the queue but to agents as well.

In addition to granting dequeue privileges, the queue owner has to make a call to `dbms_aqadm.enable_db_access()` to enable the user to dequeue as a specific agent. The agent could be subscribed to the queue using a rule that filters messages that the user should not see.

The scheduler always subscribes to every queue (secure or otherwise) using its own agent `SCHEDULER$_EVENT_AGENT`. Because the scheduler runs as user `SYS`, it can see all the messages in the queue. In the case of secure queues, however, this should not be done because the purpose of secure queues is to enable the queue owner to provide some type of limited access to the queue by users. Therefore, the scheduler requires an agent name in the `QUEUE_SPEC` attribute. This agent name is used solely for the purpose of determining which messages the user can see. The scheduler does not dequeue messages as this agent. When the scheduler receives a notification that a message has arrived in a secure queue, it performs the following checks for each event-based job that depends on the queue.

1. It checks to determine if the job owner has dequeue privileges on the queue.

2. It checks to see if the agent name provided in the queue spec can be used by the job owner.

3. It checks to see if the agent is currently subscribed to the queue.

4. It checks if the incoming message is visible to the agent.

If all of these checks succeed, then the scheduler launches the event-based job. Regardless of whether the checks pass, the message is dequeued using the agent `SCHEDULER$_EVENT_AGENT`.

When providing an agent for the `QUEUE_SPEC` attribute, the user can provide an existing agent name or have a new one created for this purpose. In either case, it is the user's or queue owner's responsibility to ensure that the agent dequeues its messages.

2.30.2.2 Remote Jobs The Oracle Scheduler supports running jobs on remote hosts and transferring files to or from remote hosts. To achieve this goal, it is first necessary to go through the setup described in the "Enabling and Disabling Remote External Jobs" section of Chapter 28 of the *Oracle Database Administrator's Guide*. This feature also requires that Oracle XML DB HTTP Server be enabled on the database.

To check whether Oracle XML DB HTTP Server is enabled, you can use the following command:

```
SQL> SELECT DBMS_XDB.GETHTTPPORT() FROM DUAL;
```

If this statement returns `0`, then you need to enable Oracle XML DB HTTP Server on a non-zero port by logging in as `SYS` and issuing the following commands:

```
SQL> EXEC DBMS_XDB.SETHTTPPORT (port);
SQL> COMMIT;
```

Substitute *port* with Oracle XML DB HTTP Server port you wish to use.

2.30.3 DBMS_SQL Package

In Oracle Database 11*g* Release 1 (11.1), Oracle introduces a number of enhancements to `DBMS_SQL` to improve the security of the package:

- Prevent guessing of open cursor numbers

  A new error, `ORA-29471`, will be raised when any `DBMS_SQL` subprogram is called with a cursor number that does not denote an open cursor. When the error is raised, an alert is issued to the alert log and `DBMS_SQL` becomes inoperable for the life of the session.

  If the actual value for the cursor number in a call to `IS_OPEN` does denote a cursor that is currently open in the session, then the return value is `TRUE`. If the actual is null, then the return value is `FALSE`. Otherwise, you get the `ORA-29471` error.

  Note that the `DBMS_SQL.OPEN_CURSOR` function is the only `DBMS_SQL` subprogram that has no formal parameter for the cursor number. Rather, it returns a cursor number. Therefore, it is not within the scope of the rules.

- Prevent inappropriate use of a cursor

  Cursors are now better protected from security breaches that subvert known, existing cursors.

  Checks are always made when binding and executing. Optionally, checks may be performed for every single `DBMS_SQL` subprogram call. The check is:

  - `current_user` is the same on calling the subprogram in question as it was on calling the most recent parse.

  - The enabled roles on calling the subprogram must be a superset of the enabled roles on calling the most recent parse.

As is always the case, for definer's right subprograms, roles are irrelevant.

If either check fails, then `ORA-29470` is raised.

The mechanism for defining when checks are performed is a new overload for the `OPEN_CURSOR` subprogram which takes a formal parameter, `security_level`, with allowed values `NULL`, `1` and `2`.

When `security_level = 1` (or is `NULL`), the checks are made only when binding and executing. When `security_level = 2`, the checks are always made.

This security regime is stricter than in 10.2 and previous releases. As a consequence, users of `DBMS_SQL` may encounter runtime errors on upgrade. While the regime makes for more secure applications, users may want to relax the security checks temporarily as they migrate to 11.1. If so, consult with Oracle Support Services on steps to relax the security checks.

2.30.4 UTL_FILE Package

The `UTL_FILE` package no longer opens a file if that file is a symbolic link. This new restriction closes a known security hole in the `UTL_FILE` package.

If this security restriction poses an undue burden during migration, then contact Oracle Support Services for assistance.

2.30.5 Interoperability Between Releases

Interoperability between an 11.1 database or Forms client and a 10.1 or 10.2 database requires a minimum patchset level of 10.1.0.5 (for 10.1) or 10.2.0.2 (for 10.2).

An attempt to reference a 10.1 or 10.2 PL/SQL unit or view under the following circumstances fails with a `PLS-801[55916]` error unless the 10.1 or 10.2 environment has been patched to the correct level:

- A PL/SQL unit, anonymous block, trigger, call statement, or SQL statement on an 11.1 database runs a PL/SQL unit on a 10.1 or 10.2 database across a database link.

- A PL/SQL unit, anonymous block, trigger, or call statement on an 11.1 database references a view on a 10.1 or 10.2 database across a database link and the view directly or indirectly references a PL/SQL function or an object type.

- An 11.1 Forms client runs a PL/SQL unit in a 10.1 or 10.2 database using RPC.

To avoid the `PLS-801[55916]` error, a minimum patchset level of 10.1.0.5 (for 10.1) or 10.2.0.2 (for 10.2) is required.

There is no interoperability problem between 10.1 and 10.2 environments.

2.31 Pro*C

The Pro*C readme file is located at:

`ORACLE_HOME/precomp/doc/proc2/readme.doc`

2.32 Pro*COBOL

The Pro*COBOL readme file is located at:

`ORACLE_HOME/precomp/doc/procob2/readme.doc`

2.33 SQLJ

SQLJ in Oracle Database 11*g* Release 2 (11.2) is supported with JDK 5.0 and JDK 6.0.

2.34 SQL*Plus

The SQL*Plus readme file is located at:

ORACLE_HOME/sqlplus/doc/README.htm

2.35 Summary Management

Note the following items when working with Summary Management.

2.35.1 Query Rewrite

The following items apply to Query Rewrite.

2.35.1.1 Fine Grained Auditing (FGA) If Fine Grained Auditing (FGA) is enabled on a table in the query, then Query Rewrite will not occur for this query.

2.35.1.2 Use of PARTITION Clause Query rewrite does not occur for queries that use the `PARTITION` clause in the `FROM` clause to access table partitions. In order for query rewrite to rewrite such queries, the `PARTITION` clauses must first be converted into equivalent selection predicates and added to the `WHERE` clause.

2.35.2 Feature Availability

The following are available in Enterprise Edition:

- Creation and refresh features of materialized views

- Query rewrite and materialized view advice from the SQL Access Advisor

2.35.3 NLS Parameters

When using or refreshing certain materialized views, you must ensure that your NLS parameters are the same as when you created the materialized view. Materialized views that fall under this restriction contain the following constructs:

1. **Expressions that may return different values, depending on NLS parameter settings**

   It is recommended to write such expressions in the NLS-independent way. For example:

   ```
   (date > DATE '2003-01-02')
   ```

   Or:

   ```
   (rate <= 2.150)
   ```

2. **Equijoins where one side of the join is character data**

   The result of this equijoin depends on collation which can change on a session basis, giving an incorrect result in the case of query rewrite or an inconsistent materialized view after a refresh operation.

3. **Expressions that generate internal conversion to character data in the select list of a materialized view, or inside an aggregate of a materialized aggregate view**

This restriction does not apply to expressions that involve only numeric data; for example, a+b where a and b are numeric values.

2.36 Oracle Streams

Note the following when working with Oracle Streams.

2.36.1 Oracle Streams Apply Handler Modifications Necessary for Handling New Error Messages That Replace ORA-1403 Errors

Customized DML and error handlers for Oracle Streams require modification to catch the additional Oracle errors ORA-26786 and ORA-26787 in place of the ORA-01403 No data found message. An ORA-26787 error is raised if the row to be updated or deleted does not exist in the target table. An ORA-26786 error is raised when the row exists in the target table, but the values of some columns do not match those of the LCR.

2.36.2 Oracle Streams Advanced Queuing

Subscription names for Advanced Queuing notification will become canonical. Thus, registrations for scott.queue1:sub and SCOTT.QUEUE1:SUB are for the same entity and are represented in canonical form as "SCOTT"."QUEUE1":"SUB".

2.36.3 UTL_SPADV Load Privileges

An additional privilege is required to load the UTL_SPADV package into the Oracle Streams administrator schema. The Oracle Streams administrator schema must have EXECUTE privilege on DBMS_LOCK in order to load this package.

2.37 Documentation Addendum

This section contains corrections to the following Oracle Documentation for this release:

- Section 2.37.1, "Oracle Database Heterogeneous Connectivity User's Guide"

- Section 2.37.2, "Oracle Database PL/SQL Packages and Types Reference"

- Section 2.37.3, "Oracle Grid Infrastructure Installation Guide"

2.37.1 Oracle Database Heterogeneous Connectivity User's Guide

The following information should be removed from Chapter 4, "Using Heterogeneous Services Agents", section "Determining the Heterogeneous Services Parameters":

The Distributed Access Manager has a refresh capability available through the menu and toolbar that allows users to rerun queries if necessary and update the data. When the data is refreshed, the tool verifies that the set of registered agents remains the same. If it is not, then the global view is updated. See *Oracle Enterprise Manager Administrator's Guide* and online help for more information about the Distributed Access Manager.

2.37.2 Oracle Database PL/SQL Packages and Types Reference

The listno parameter of the DBMS_UTILITY was inadvertently excluded from the *Oracle Database PL/SQL Packages and Types Reference*. For example, the GET_PARAMETER_VALUE function should read as follows:

GET_PARAMETER_VALUE Function

This function gets the value of specified initialization parameter.

Syntax

```
DBMS_UTILITY.GET_PARAMETER_VALUE (
    parnam      IN       VARCHAR2,
    intval      IN OUT   BINARY_INTEGER,
    strval      IN OUT   VARCHAR2,
    listno      IN       BINARY_INTEGER DEFAULT 1)
  RETURN BINARY_INTEGER;
```

Parameters

Parameter	Description
parnam	Parameter name.
intval	Value of an integer parameter or the value length of a string parameter.
strval	Value of a string parameter.
listno	List item number. If retrieving parameter values for a parameter that can be specified multiple times to accumulate values, use this parameter to get each individual parameter.

2.37.3 Oracle Grid Infrastructure Installation Guide

In *Oracle Grid Infrastructure Installation Guide*, in the section "Creating Disk Groups for Oracle Database Data Files," the guide indicates that logical volume managers are not supported with Oracle RAC. This is incorrect. The correct support information is as follows:

> Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. In addition, Oracle RAC requires a cluster logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

2.38 Open Bugs

This section lists known bugs for this release. A supplemental list of bugs may be found as part of the release documentation specific for your platform.

2.38.1 Oracle Automatic Storage Management (Oracle ASM) Known Bugs

Bug 9655918

Increase in Oracle ASM disk group size for Oracle ACFS Database Home may be required.

Workaround: When creating a Database Home on Oracle ACFS, the ASMCA tool will default to a minimum of 6 GB for the Oracle ACFS Database Home. However, depending on the status and configuration of the disk group, additional storage may be required. Therefore, the recommendation is to use a minimum of 10 GB for the disk group to avoid any such issues, as reported by these messages:

```
ORA-15032: not all alterations performed
ORA-15041: diskgroup "dg_name" space exhausted
```

Bug 9413827

An 11.2.0.1 Oracle Clusterware rolling upgrade to 11.2.0.2 fails when Oracle Cluster Registry (OCR) is on Oracle ASM.

Workaround: Apply the patch for bug 9413827 on 11.2.0.1 Oracle Grid Infrastructure for a cluster home before performing the upgrade.

Bug 9276692

Cannot permanently stop the Oracle ASM instance.

Workaround: If the Oracle ASM instance is disabled using SRVCTL, you must unregister Oracle ACFS-related resources to avoid restarting the Oracle ASM instance. Do this by executing the following command as root:

```
acfsroot unregister
```

2.38.2 Deinstallation Tool Known Bugs

Bug 9925724

If *Grid_home* is created directly under a root-owned directory, the deinstallation tool cannot remove the top-level home directory. An empty Oracle home directory remains at the end of the deinstallation.

Workaround: Run rmdir *ORACLE_HOME* using the root user on all nodes.

Bug 9924588

The deinstallation tool is not deleting the old Grid home after upgrading to a new release.

Workaround: After upgrading from previous releases, if you want to deinstall the previous release Oracle Grid Infrastructure Grid home, then you must first change the permission and ownership of the previous release Grid home. Log in as root, and change the permission and ownership of the previous release Grid home using the following command syntax, where oldGH is the previous release Grid home, jsmith is the Oracle Grid Infrastructure installation owner, and oldGHParent is the parent directory of the previous release Grid home:

```
#chmod -R 755 oldGH
#chown -R jsmith oldGH
#chown jsmith oldGHParent
```

As the installation user, use the 11.2.0.2 standalone deinstallation tool to remove oldGH.

Bug 9871334

If the deinstallation tool is launched from Oracle home, it fails to delete the *ORACLE_HOME*/deinstall directory. This issue is not seen with the standalone deinstallation tool. An empty *ORACLE_HOME*/deinstall directory remains at the end of deinstallation.

Workaround: Manually delete empty directory *ORACLE_HOME*/deinstall and *ORACLE_HOME* at the end of the deinstallation.

Bug 9406768

The 11.2 deinstallation utility is removing all the homes under Oracle base if these homes are not using the same central inventory and the deinstallation utility finds this home is the only one registered in inventory.

Workaround: While installing 11.2 products:

- Oracle does not recommend using multiple central inventories. Avoid this if possible.

- If for some reason a different central inventory is required, use a different Oracle base directory for each central inventory.

Bug 8644344

When running the deinstallation tool to deinstall the database, you will be prompted to expand the Oracle home and to select a component. If you select the top level component, `Oracle Database Server`, and do not select the Oracle home, OUI does not show the message to run the deinstall utility and proceeds with the deinstallation of the database.

Workaround: Run the deinstallation tool to deinstall the Oracle home.

Bug 8635356

If you are running the deinstall tool from `ORACLE_HOME` that is installed on shared NFS storage, then you will see errors related to `.nfs` files during `ORACLE_HOME` clean up.

Workaround: To remove the `ORACLE_HOME`, run the `rm -rf ORACLE_HOME` command after the deinstall tool exits. Alternatively, you can use the standalone `deinstall.zip` and specify the location of the `ORACLE_HOME`.

2.38.3 Oracle ACFS Known Bugs

Bug 10019796

When `acfsutil encr info` command is run on a file or a directory on which encryption has never been enabled, the command produces no output. Furthermore, when `acfsutil encr info` is run on a directory with the recursive (`-r`) option and encryption has never been enabled on that directory, the command produces no output for the directory or any files contained in that directory.

Workaround: When `acfsutil encr info` is run on an individual file or directory and the command produces no output, it implies that encryption is not enabled on that file or directory.

When `acfsutil encr info` is run on a directory with the recursive (`-r`) option and the command produces no output for the specified directory or one of its subdirectories, it implies that encryption is not enabled on that directory. To see the encryption status of files in that directory, run `acfsutil encr info` on each file individually.

Bug 10069735

In a cluster with a password-protected key store, when an Oracle ACFS file system using encryption is mounted through the Oracle ACFS mount registry, the administrator is not prompted to enter the key store password. Although the process of mounting the file system succeeds, not all information required for Oracle ACFS encryption to work correctly is made available to the file system. In this case,

encryption is not operational on this file system and any encrypted files in the file system are not available for read or write.

Workaround: In a cluster with a password-protected key store, do not use the Oracle ACFS mount registry for mounting any file systems that are using encryption. If some file systems are already mounted through the Oracle ACFS mount registry, unmount them and remove any such file systems from the mount registry to avoid possible unavailability of encrypted data in the future. Then, remount these file systems without using the Oracle ACFS mount registry, providing the correct password when requested.

Bug 9975343

In some situations when `acfsutil encr set` or `acfsutil sec prepare` commands are run from multiple nodes in succession, the `.Security` directory in the mount point may be created incorrectly.

If this happens, two entries will be shown for `.Security` when listing the contents of the mount point, and the `acfsutil sec prepare` command will fail to add the logging and backup directories to their respective system realms if this command was the last to be run.

Workaround: If `acfsutil encr set` or `acfsutil sec prepare` commands are to be performed in succession, executing the commands from the same node will avoid the problem.

The problem can also be avoided by listing the contents of the `.Security` directory before running the second command.

However, if the problem is encountered, it can be resolved by first removing both security and encryption from the file system using `acfsutil encr set -u` and `acfsutil sec prepare -u` and then removing the `.Security` directory. The removal of the `.Security` directory will have to be performed twice.

Once security and encryption have been removed from the file system and the `.Security` directory has been deleted, the `acfsutil sec prepare` or `acfsutil encr set` commands may be run again on the file system.

Bug 9958114

If the standby file system is initialized for replication, but the primary file system is not, `repl terminate` of the standby will loop forever.

Workaround: Use `acfsutil repl terminate standby immediate` to terminate the standby under these conditions.

Bug 9947559

After a certain number of retries, the `acfsrepl_transport` daemon attempts to failover to another node if it cannot communicate with the remote site. If there is no other node for the daemon to failover to, or after a certain number of failover attempts, Oracle Clusterware stops trying to failover and the transport daemon will not be running. Not running the transport daemon causes replication to stall and eventually the file systems can run out of space.

Workaround: Run `acfsutil repl bg start` *mnpt* after the communication problem is corrected.

Bug 9940090

Unable to communicate with remote site if the disk group is forced dismounted. The root cause of this issue is that the service name associated with the file system is not being unregistered when the disk group is forced dismounted.

Workaround: Unmount the file system whose disk group was forced dismounted. This will unregister the service, thereby allowing the remote site to choose another node with the disk group and file system mounted for communication.

Bug 9896821

The following are alert log messages pointing to trace files that do not indicate an actual problem:

```
ORA-19505: failed to identify file
"/scratch/primary/.ACFS/repl/ready/receipt.transport.cord242.12786023 69"
ORA-27037: unable to obtain file status
Linux-x86_64 Error: 2: No such file or directory
```

Workaround: None.

Bug 8644639

When creating an Oracle ACFS mount point and adding it to the registry, the mount point is not mounted automatically if the following conditions are met:

1. The mount point directory was previously registered with the Oracle ACFS Registry.

2. The mount point directory had been previously mounted.

3. The mount point had then been unmounted and removed from the Oracle ACFS Registry.

4. The `ora.registry.acfs` resource has not been restarted since the mount point was deleted from the registry.

Workaround: Remove the mount point directory from the file `/tmp/.usm_state_file`.

2.38.4 Oracle ASM Dynamic Volume Manager (Oracle ADVM) Known Bugs

Bug 9683229

Oracle ADVM does not support mounting ext3 file systems over Oracle ADVM with the mount barrier option enabled. The mount barrier option is enabled by default on SLES10.

Workaround: Mount ext3 file system with `-o barrier=1`. For example:

```
mount -o barrier=0 /dev/asm/myvol-131 /mnt
```

2.38.5 Oracle Application Express Known Bugs

Refer to "Open Bugs and Known Issues" in *Oracle Application Express Release Notes*.

2.38.6 Oracle Clusterware Known Bugs

Bug 10062301

After a non-rolling clusterware upgrade from 11.2.0.1 to 11.2.0.2, the `CLUSTER_NAME` parameter is not set in the *clusterware_home _path*/crs/install/crsconfig_params. This affects the `addNode` scenario after the upgrade.

Workaround: Execute the following:

1. Execute *clusterware_home_path*/bin/olsnodes -c to get the cluster name.

2. Edit the *clusterware_home_path*/crs/install/crsconfig_params and update the `CLUSTER_NAME` entry as `CLUSTER_NAME=`*clustername_acquired_in_step1*.

Bug 10055232

The `root.sh` scripts fails after downgrading from 11.2.0.2 to 11.2.0.1.

Workaround:

Downgrade using the following steps when Oracle Cluster Registry (OCR) and voting disk *are not* on Oracle ASM:

1. Stop Oracle Clusterware running from the 11.2.0.2 home and verify that it is stopped. For example:

    ```
    11.2.0.2_Home/bin/crsctl stop crs
    11.2.0.2_Home/bin/crsctl check crs
    ```

2. Copy the OLR backup file and verify that the `olrconfig_loc` and the *crs_home* properties in the `olr.loc` file point to the 11.2.0.1.0 home. For example:

    ```
    cd location_of_olr_loc
    cp olr.loc.bkp olr.loc
    ```

    Note that *location_of_olr_loc* is platform-specific. Refer to the platform-specific documentation for the location of `olr.loc` on your system.

3. On the last node to downgrade, format OCR with the 11.2.0.1.0 snapshot. For example:

    ```
    ocrconfig -import 11202_crshome/cdata/ocr11.2.0.1.0
    ```

4. Remove one or more voting files. The voting file is located at the location given by the user while installing 11.2.0.1 prior to the upgrade to 11.2.0.2. Usually, it is located at the same place as OCR. For example:

    ```
    rm vote_file
    ```

5. Start the stack in exclusive mode. For example:

    ```
    11.2.0.1_Home/bin/crsctl start crs -excl
    ```

6. Add the voting file or files. For example:

    ```
    crsctl add css votedisk voting_file
    ```

7. Stop the stack running in exclusive mode and verify it stopped. For example:

    ```
    11.2.0.1_Home/bin/crsctl stop crs -f
    11.2.0.1_Home/bin/crsctl check crs
    ```

8. On all nodes, start the Oracle Clusterware from the 11.2.0.1 home. For example:

```
11.2.0.1_Home/bin/crsctl start crs
```

Downgrade using the following steps when Oracle Cluster Registry (OCR) and voting disk *are* on Oracle ASM:

1. If the `compatible.asm` attribute for the OCR disk group is set to 11.2.0.2, then you must migrate the OCR to a disk group with `compatible.asm` set to 11.2.0.1 using the `ocrconfig -add` and `ocrconfig -delete` commands before starting the downgrade procedure.

2. Stop Oracle Clusterware from running from the 11.2.0.2 home and verify that it is stopped. For example:

```
11.2.0.2_Home/bin/crsctl stop crs
11.2.0.2_Home/bin/crsctl check crs
```

3. Copy the OLR backup file and verify that the olrconfig_loc and the `crs_home` properties in the `olr.loc` file point to the 11.2.0.1.0 home. For example:

```
cd location_of_olr_loc
cp olr.loc.bkp olr.loc
```

Note that *location_of_olr_loc* is platform-specific. Refer to the platform-specific documentation for the location of `olr.loc` on your system.

4. On the last node to downgrade, start the 11.2.0.2.0 Oracle Clusterware stack in exclusive mode. For example:

```
11.2.0.2_Home/bin/crsctl start crs -excl -nocrs
```

5. Delete the voting disk from the Oracle ASM disk group. For example:

```
11.2.0.2_Home/bin/crsctl delete css votedisk <diskgroup>
```

6. Stop Oracle Clusterware running in exclusive mode from the 11.2.0.2 home and verify it is stopped. For example:

```
11.2.0.2_Home/bin/crsctl stop crs -f
11.2.0.2_Home/bin/crsctl check crs
```

7. On the last node to downgrade, start the 11.2.0.1.0 Oracle Clusterware stack in exclusive mode. For example:

```
11.2.0.1_Home/bin/crsctl start crs -excl
```

8. Stop the Oracle Clusterware daemon. For example:

```
11.2.0.1_Home/bin/crsctl stop resource ora.crsd -init -f
```

9. On the last node to downgrade, format OCR with 11.2.0.1.0 snapshot. For example:

```
11.2.0.1_Home/bin/ocrconfig -import 11202_crshome/cdata/ocr11.2.0.1.0
```

10. Verify the integrity of OCR. For example:

```
11.2.0.1_Home/bin/ocrcheck
```

11. Stop the stack running in exclusive mode and verify it is stopped. For example:

```
11.2.0.1_Home/bin/crsctl stop crs -f
11.2.0.1_Home/bin/crsctl check crs
```

12. On all nodes, start Oracle Clusterware from the 11.2.0.1 home. For example:

    11.2.0.1_Home/bin/crsctl start crs

13. Verify the OCR integrity from all nodes. For example:

    11.2.0.1_Home/bin/cluvfy comp ocr -n all -verbose

Bug 10044507

While verifying udev attributes during installation, upgrade or manual execution of cluvfy, an error may be reported even though the expected udev attributes are defined correctly in the appropriate udev rules file. For example:

```
PRVF-5184 : Check of following Udev attributes of "<node>:<path>" failed:
"[Permissions: Found='0600' Expected='0660']"
```

This error is due to the expectation that udev attributes for the device are expected to be in a specific order in the udev rules file.

Workaround: When specifying the device information in the udev rules file, make sure that the OWNER, GROUP and MODE are specified in this order before any other characteristics are specified. For example, if you want to include the ACTION characteristic on the udev line, you should specify the following:

```
KERNEL=="<device>",  OWNER="oracle", GROUP="dba", MODE="660", ACTION=="add|change"
```

And, do not specify the following:

```
KERNEL=="<device>", ACTION=="add|change", OWNER="oracle", GROUP="dba", MODE="660"
```

Bug 10043601

During Oracle Grid Infrastructure for a cluster installation, when a file created by dd on NFS is used as Oracle ASM disks, the prerequisite check for Oracle ASM devices reports an error (with PRVF-5150) that the given disk is not a valid path.

Workaround: This error can be ignored.

Bug 10041397

During Oracle Grid Infrastructure for a cluster installation on Oracle Virtual Machine (Oracle VM), when virtual shared disks are used as Oracle ASM devices, the prerequisite check for Oracle ASM devices gives a warning that these devices are not shared. For example:

```
PRVF-5149 : WARNING: Storage "/dev/xvde1" is not shared on all nodes
```

Workaround: This warning can be ignored.

Bug 10033106

When running addNode.sh with the -silent flag to add a new node to Oracle Clusterware, if there are any failures during cluvfy stage -pre nodeadd, no error messages will be displayed. The addNode operation will be terminated and the new node will not be added.

Workaround: Run cluvfy stage -pre nodeadd command from an existing node of the cluster. Look at the errors reported and fix them. If you want to ignore the errors and

continue with the addNode operation, set the environment variable IGNORE_
PREADDNODE_CHECKS to Y and retry the addNode operation.

Bug 10032601

During Oracle Grid Infrastructure for a cluster installation or during invocation of the
cluvfy stage -pre crsinst command, the prerequisite check for user existence
reports an error if the numerical value of the user ID exceeds 6 digits. This is
applicable to both the local and NIS users.

Workaround: The reported user existence validation error can be ignored.

Bug 10025999

If the Oracle High Availability Services Daemon (OHASD) crashes and restarts while
the Oracle Clusterware stack is already running, subsequent Cluster Verification
Utility (CVU) checks may produce the following informational messages:

```
INFO: PRVF-9656 : The CTSS command to query time offset and reference failed on
node nodename with error message ""
INFO: Query of CTSS for time offset failed
```

Workaround: These informational messages can be ignored.

Bug 9978765

Oracle Grid Infrastructure for a cluster cannot be upgraded to Oracle Grid
Infrastructure 11*g* Release 2 Patchset 2 from Oracle Clusterware releases 10*g*R1, 10*g*R2
and 11*g*R1 when there are multiple private network interfaces configured in the
previous release. This can be verified by running the OIFCFG getif command prior to
upgrading, and inspecting the number of interfaces marked cluster_interconnect.

Workaround: Execute the steps in the following procedure to upgrade Oracle Grid
Infrastructure for a cluster. This example assumes a 2-node (A and B) cluster:

1. Shut down all database instances and Oracle ASM running on nodes A and B in the
    10*g*R1, 10*g*R2 or 11*g*R1 cluster.

2. On node A, identify the private network interface used by Oracle Clusterware
    using the OLSNODES -p command. Delete the rest of the private network
    interfaces marked cluster_interconnect using OIFCFG delif -global *if_name*
    command.

3. Start up all database instances and Oracle ASM running on nodes A and B in the
    10*g*R1, 10*g*R2 or 11*g*R1 cluster.

4. Proceed upgrading the Oracle Grid Infrastructure for the cluster 11*g* Release 2.

5. Verify that the Oracle Grid Infrastructure for a cluster upgrade is complete using
    crsctl query crs activeversion command. Add the private network interfaces
    which were removed in Step 2 using oifcfg setif -global *if_name*:cluster_
    interconnect command.

Bug 9974223

Oracle Grid Infrastructure for a cluster needs multicast communication on 230.0.1.0
addresses working properly between the cluster nodes.

When first attempting to start an 11.2.0.2 cluster node, either during a fresh install or
after an upgrade attempt, it is possible that the node will be unable to join the cluster.

The Oracle Cluster Synchronization Services (OCSSD) will report a failure to join the existing cluster members.

This will be noted as an inability for the CSS to connect to the other active members of the cluster.

The problem occurs because the clusterware in 11.2.0.2 uses multicast messages to initiate communications on the 230.0.1.0 address. If multicast on these addresses is not enabled for the private interfaces, then the clusterware will be unable to connect to it's peer nodes.

Workaround: Contact your network IT department and have them validate that multicast communications are working across all nodes of the cluster on the private interconnect. In particular, ensure that multicast communications on the 230.0.1.0 addresses are working properly between the cluster nodes.

Bug 9951257

If Oracle Cluster Registry (OCR) is on a disk group and that disk group cannot be automatically mounted, it prevents, Cluster Ready Services Daemon (CRSD) from starting (also reference Bug 9799693).

Workaround: You must recover the OCR disk group manually. Determine the reason the disk group has failed to mount, correct the problem and mount the disk group. See also message `CRS-5019` for recovery options.

Bug 9948549

If the agent terminates while starting a database or Oracle ASM instance, it is possible that the instance startup will not complete.

Workaround: Stop and restart the instance using `srvctl` or `sqlplus`.

Bug 9877357

The Oracle Enterprise Manager configuration fails during DBCA if a policy-managed database's name contains '_' or more than 8 characters.

Workaround: Choose a database name accordingly.

Bug 8974576

`REMOTE_LISTENER` parameter was not changed after a `SRVCTL ADD SCAN_LISTENER` or `SRVCTL REMOVE SCAN_LISTENER` command.

In a future release, Oracle Clusterware may start managing `REMOTE_LISTENER` as part of a project to support multiple public networks each with their own `SCAN`. But in release 11.2, because of the need to include node VIPs in the `REMOTE_LISTENER` parameter of pre-11.2 databases that are upgraded, Oracle Clusterware does not have all the data needed to keep `REMOTE_LISTENER` properly updated. This is, instead, implemented by DBCA and DBUA, or by manual means.

Workaround: To manually change the REMOTE_LISTENER parameter, execute the following:

```
SQL> ALTER SYSTEM SET REMOTE_
LISTENER='t2000-cluster0-scan.t2000-18.oraclecorp.com' sid='*' scope=memory;

System altered.
```

```
SQL> ALTER SYSTEM REGISTER;

SQL> SHOW PARAMETER LISTENER;

NAME                                 TYPE        VALUE
------------------------------------ ----------- ------------------------------
remote_listener                      string      t2000-cluster0-scan.t2000-18.o
```

Bug 8737732

If you upgrade a previous Oracle Clusterware release to Oracle Grid Infrastructure for a cluster 11*g* Release 2 (11.2.0.2), Cluster Verification Utility (CVU) may fail if the Oracle Grid Infrastructure for a cluster installation owner does not have correct permissions or ownership of the raw devices on which the Oracle Cluster Registry and voting disk files are placed.

Workaround: Check CVU output in the installation log file, and change permissions or ownership of raw devices to the settings specified in the log file. Click Retry to run the configuration tool again.

Bug 8733944

Due to a problem in Oracle Clusterware starting with release 11.1.0.7, with the patches required for Exadata support or 11.1.0.7 CRS bundle Patch 1, in some cases the CSS daemon may fail when the clusterware is brought down on another node, either due to a shutdown command or a failure.

The symptom is an ASSERT in the CSSD log indicating that a maximum value has been exceeded. For example:

```
Group ID of xxxx exceeds max value for global groups
```

Workaround: Oracle recommends that customers running with the Exadata support patches or 11.1.0.7 CRS Bundle Patch 1 apply the patch for this bug to avoid this problem.

This problem may also be seen during an upgrade from 11.1.0.7 with patches as indicated above. To eliminate the potential of an 11.1.0.7 node failing during upgrade, the patch for this bug may be applied to the 11.1.0.7 nodes prior to upgrade.

When upgrading, it is recommended that the upgrade be completed on all nodes without restarting any non-upgraded nodes during the upgrade. If an 11.1.0.7 node does fail while performing the upgrade, it should be upgraded as opposed to restarted.

Bug 3841387, 8373758

Oracle resources for 10.1 Oracle RAC Databases may not operate properly after upgrading Oracle Clusterware to 11.2.

Workaround: Apply the patches for Bug 3841387 and Bug 8373758 to the Oracle Database home.

Bug 8668083

In Oracle Restart environments, the following commands cannot perform a verification and return an error message if the user equivalence does not exist for the local node:

```
cluvfy stage -pre hacfg
cluvfy stage [-pre | -post] cfs
cluvfy stage [-pre | -post] acfscfg
cluvfy comp admprv
cluvfy stage -pre dbcfg
cluvfy stage -pre dbinst
```

Workaround: Configure SSH in Oracle Restart.

Contact Oracle Support Services to obtain the patches for these bugs and apply the patches to the Oracle Database home.

Bug 8666509

A deinstallation of Oracle Clusterware should ask you to detach any pre-11.2 Oracle RAC homes from the Oracle inventory.

Workaround: After you deconfigure and deinstall an upgraded 11.2 Oracle RAC home and want to continue with deconfiguration and deinstallation of the Oracle Grid Infrastructure for a cluster home, first detach any pre-11.2 Oracle RAC software homes from the central Inventory.

Bug 8657184

If two network interfaces are configured as public network interfaces in the cluster, the failure of one public interface on a node does not result in automatic VIP failover to the other public interface.

Workaround: If multiple public networks interfaces are present, then use interface bonding for high availability. At the Oracle Clusterware installer "Specify Network Interface Usage" screen, choose only one (bonded) interface as public. When configuring public networks with srvctl add nodeapps or srvctl add vip, specify only a single network interface name in the -A or -S argument.

Bug 8262786, 8373758, 8441769, 8406545

Oracle resources for 10.2 and 11.1 Oracle RAC databases may not operate properly after upgrading Oracle Clusterware to 11.2.

Workaround: Apply the patches for Bugs 8262786, 8373758, 8441769, and 8406545 to the Oracle Database home.

Bug 8288940

Creating pre-11.2 Oracle RAC database in 11.2 Oracle Grid Infrastructure for a cluster environment using DBCA may fail with following messages. When using a cluster file system as storage, you see the following message:

```
ORA-00119: invalid specification for system parameter REMOTE_LISTENER
```

When using Oracle ASM as storage, you see the following message:

```
DBCA could not startup the ASM instance configured on this node
```

Workaround: Apply the patch for this bug in pre-11.2 database home. This patch is needed for 10.2.0.4, 11.1.0.6 and 11.1.0.7 database releases. No patch is needed for release 10.2.0.5.

2.38.7 Oracle Data Pump Known Bugs

Bug 10013976

When a database level, schema level, or queue table database export log generated from an 11.2 database is imported into an 11.2 database, `ORA-1403` errors may be raised during an Oracle Streams Advanced Queuing (AQ) procedural import action. Import completes successfully with errors.

Workaround: None. The `ORA-1403` errors in the import log are not fatal. Import of database objects including AQ queues completes successfully.

2.38.8 Oracle Database Vault Known Bugs

Bug 10225918

The stated issue can be reproduced when Database Configuration Assistant (DBCA) attempts to configure Database Vault while the local listener is not properly configured. Since this failure leads to partially populated metadata, some Database Vault features may not function as expected.

Workaround: The following steps explain a workaround to this issue:

1.   Remove Database Vault completely using `rdbms/admin/dvremove.sql`.

2.   Make sure that local listener (`LISTENER`) is configured properly.

3.   Retry Database Vault installation and configuration using DBCA.

Optionally, you can manually execute the failed DVCA operations (for example, `dvca -action addlanguages` and `catmacpost.sql`) to recover from the failure by following the steps previously mentioned. Make sure that local listener (`LISTENER`) is configured properly before running the steps.

Bug 8686258

Database Vault policy cannot be managed in Oracle Enterprise Manager Database Control because the following message is displayed in Database Vault Administration page:

```
"OPERATOR TARGET" privilege does not exist.
"You must have OPERATOR TARGET privilege to perform this operation."
```

Workaround: To manage Database Vault policy using Oracle Enterprise Manager, the Database Vault administrator must have the EM Administrator privilege. If you do not want to grant the EM Administrator privilege to the Database Vault administrator, then use the Database Vault Administrator page directly. For additional information, see *Oracle Database Vault Administrator's Guide*.

Bug 8474817

If a user is granted Database Vault scheduler job authorization by `DBMS_MACADM.AUTHORIZE_SCHEDULER_USER` API, dropping this user does not remove the user from the list of authorized users.

Workaround: Invoke the `DBMS_MACADM.UNAUTHORIZE_SCHEDULER_USER` API call to remove the user's authorization manually.

Bug 8341283

The `ACTION_NAME` entry in the `DVSYS.AUDIT_TRAIL$` table displays `Realm Authorization Audit` for a failed realm enforcement if the audit option is set to audit on success and failure. The `RETURNCODE` will show the correct error code that was triggered.

Workaround: Use the `RETURNCODE` value to determine whether a violation has occurred and the `ACTION_NAME` column to identify whether the audit was generated by a realm enforcement or command rule enforcement.

Bug 8247647

Dropping a user does not automatically remove the user from the list of authorized users for Oracle Data Pump.

Workaround: Invoke the `DBMS_MACADM.UNAUTHORIZE_DATAPUMP_USER` API call to manually remove the user's authorization for Oracle Data Pump.

Bug 7033772

Database Vault Administrator (DVA) does not work after an Enterprise Manager DBControl upgrade.

Workaround: Manually redeploy DVA after DBControl has been upgraded. You can follow the steps described in Appendix C, Section "Deploying Database Vault Administrator to the Database Console OC4J Container" of the *Oracle Database Vault Administrator's Guide*.

2.38.9 Oracle Database Enterprise Edition Known Bugs

Bug 10107649

Creating a draft service request (SR) might fail when uploading a correlated package.

Workaround: Upload the correlated package to an existing SR (normally the same SR as the main package).

Bug 9951827

Transportable tablespace import does not handle timestamp with timezone version change.

If a transportable dumpfile produced in release 11.2.0.2 contains tables with timestamp with timezone columns and the version of the timezone table on the target database is different than that from the source database, the import is be prevented from running.

If a dumpfile produced prior to release 11.2.0.2 had a different timezone table version than that of the target, then the import is be prevented from running.

Workaround: Make sure the timezone tables for the import and export databases are the same.

Bug 9859532

Current implementation of node-specific network interfaces requires complete definition of all networks used by Oracle RAC for that node (that is, either node abides by global network configuration or it defines its own node-specific network configuration).

As a corollary, once the first node-specific network interface is defined for a given node, Oracle RAC will not consider any configured global network interfaces that are already configured and may have applied for the same node.

While this is correct, it presents a problem. If the cluster had a working global network configuration, the moment a user updates it (using `oifcfg`) to define a node-specific `public` interface, a global configuration will not be considered for this node, and it will have only one newly-defined `public` interface. Any cluster interconnects that existed in the global network configuration, and may still resolve fine for this node, will not be considered valid. Thus, the node loses cluster interconnects and PCW stack goes down on that node.

Workaround: If the node belongs to a global cluster network configuration, then if there is an intent to make network configuration node-specific, the first defined node-specific interface must be cluster interconnect, so that node never loses interconnect with other cluster nodes. Then, other node-specific interfaces can be defined as necessary.

Bug 9181675

The `asmgidwrap` script needs to be called if you are creating a database manually on Oracle ASM to avoid a permission error.

Workaround: For a role-separated installation (that is, there is a different user and group for grid and RDBMS), use DBCA to create the database that automatically calls `asmgidwrap` script while creating a database on Oracle ASM. If you choose to create a database manually, the script needs to be called explicitly so the proper group can be set to avoid a permission error.

Bug 8729627

When using 11.1 DBCA to remove a database on a cluster running 11.2 Oracle Clusterware, a `PRKP-1061/CRS-2524` error may be displayed because the database resource is locked.

Workaround: You can ignore the message. Click **OK** to continue.

Bug 8679182

When configuring a database on a cluster that has multiple public subnets defined for its VIPs (for example, using a command similar to `srvctl add vip -k 2 -A ...`), the database agent automatically sets `LOCAL_LISTENER` to the listener on the default network. This may duplicate a listener set in `LISTENER_NETWORKS`.

Workaround: Do not specify listeners in `LISTENER_NETWORKS` that are on the default public subnet.

Bug 8671032

In Oracle Enterprise Manager, the "TIP" obtained from the "Run Analysis" on the Automatic Undo Management page is not correct. For example, the Recommended Undo Tablespace Size should be three times the Minimum Size, to allow for workload fluctuations, but Oracle Enterprise Manager shows different values.

Workaround: The "TIP" should read as follows:

```
Oracle advises that you configure the undo tablespace to be three times the
Recommended Undo Tablespace Size to allow for workload fluctuations
```

Bug 8304261

If compression attributes are specified at the table level and in the sub-partition template, newly created interval partitions do not use the compression attributes defined in the sub-partition template, but use the table level compression instead.

Workaround: Redefine compression attributes in sub-partitions.

2.38.10 Oracle Enterprise Manager Database Control Known Bugs

Bug 10162200

Clone Database offers an option to configure Database Control during the clone process. If that option is specified when creating an Oracle ASM-based clone database, Database Control is not configured.

Workaround: Manually run Oracle Enterprise Manager Configuration Assistant (EMCA) to setup Database Control.

Bug 9766628

`emctl` commands did not return valid results as expected.

Workaround: The `emctl` command needs to be run from an Oracle Database home. Do not invoke this command from the Oracle Grid Infrastructure for a cluster home.

Bug 8674920

If the installation owners for the Oracle Grid Infrastructure for a cluster and Oracle Database are different, then the owners of Oracle ASM binaries and Oracle Enterprise Manager Agent binaries are also different. When you start Support Workbench, the error message `Error Operation failed - Operation failed` might appear, because the Oracle Enterprise Manager Agent is running as a different user, and Support Workbench does not have permissions for the Oracle ASM target.

Workaround: None.

Bug 8350597

If Database Control is running in an IPv6 environment, then you cannot use it to monitor Exadata cells and you should not add Exadata cells as targets.

Workaround: None.

2.38.11 Oracle Database QoS Management Known Bugs

Bug 10019209

This bug applies to platforms that support the Cluster Health Monitor (CHM). If an Oracle Clusterware-managed database service is in a stopped but not disabled state, it will be started by Oracle Database QoS Management if the server hosting that service is not detected to be in a memory overcommitted state. If memory is overcommitted, then all enabled services will be stopped even if they were manually started. The desired behavior is to only start services on the transition from a memory overcommitted state (red) to a normal state (green). If a service is manually started when the server is in the red state, that service should not be shut down.

Workaround: Stop and disable services that you want to remain in the stopped state or disable QoS Management from the Oracle Enterprise Manager Console.

Bug 10019186

The Oracle Database QoS Management feature that protects servers from overcommitting memory is only available on platforms that also support Cluster Health Monitor (CHM).

Workaround: None.

2.38.12 Oracle OLAP Known Bugs

Bug 9917299

If the database is installed using the seed provided in the installation kit, and the OLAP option is not selected, then either at the end of the installation or some time later, the OLAP Analytic Workspace and OLAP API components will be reported as invalid.

This will not affect the running of the instance in any way, other than the error messages.

Workaround: Do one of the following as a workaround:

- Ignore the error.

- Enable OLAP (or the offending option).

- Create and use your own seed database that does not include OLAP.

Bug 9545221

Importing an materialized view-enabled cube or cube dimension whose source table is not part of the target schema fails with an `Object not found` error.

Workaround: Disable materialized views for the failing object prior to the import, then reenable them when the source tables are present.

2.38.13 Oracle SQL*Loader Known Bugs

Bug 9301862

When external table code reads very large files on disks served by NFS, the I/O performance of the read can slow down over time. This is caused by NFS caching blocks from the file in memory as it is read. Since these blocks are not re-read, the time spent maintaining the cache slows down the I/O operations.

Workaround: The current behavior (of not using the `O_DIRECT` flags) remains the default. You can enable the use of the `O_DIRECT` flag in the following ways:

- Enable fix control for this bug and set it to `ON` with the following command:

  ```
  ALTER SESSION SET "_fix_control"='9301862:ON';
  ```

  When fix control is enabled, the external table code looks at the `FILESYSTEMIO_OPTIONS` configuration parameter and if it is set to either `DIRECTIO` or `SETALL`, then the `ORACLE_LOADER` access driver will specify the `O_DIRECT` flag when opening data files for reading. If the `FILESYSTEMIO_OPTIONS` parameter is not set or if it is set to

other values, then the access driver will not attempt to use `O_DIRECT` unless you choose the following option.

- Use the new `IO_OPTIONS` clause in the access driver to specify direct I/O. The clause is part of the larger `RECORDS` clause. The syntax is:

```
IO_OPTIONS (DIRECTIO | NODIRECTIO)
```

If `DIRECTIO` is specified, then the access driver uses `O_DIRECT` flag when opening the file. If `NODIRECTIO` is specified, then the access driver does not use the `O_DIRECT` flag. Note that the action specified by `IO_OPTIONS` is performed regardless of the setting of `_fix_control` for this bug.

Note that the first option is a way to enable the use of `O_DIRECT` for all external tables while the second option allows `DIRECTIO` to be used or not used for specific external tables.

2.38.14 Oracle Streams Known Bugs

Bug 10115681

When the client application detaches from an inbound server while transmitting a transaction with more than 1000 logical change records (LCRs), the coordinator process aborts because it receives an invalid dequeue low-watermark from the reader server.

Workaround: Avoid detaching from the inbound server in the middle of a large transaction. Wait until the `COMMIT` LCR for the large transaction is sent and then detach.

2.38.15 Oracle Universal Installer Known Bugs

Bug 10253337

`addnode.sh` fails if `DISPLAY` is not set.

Workaround: Run the `addnode.sh` command with `-silent` flag.

Bug 10096522

Running the `configToolAllCommands` configuration step after the silent installation of Oracle Grid Infrastructure software or database software may fail if the installer session was restarted while choosing to apply software updates.

Workaround: Change the value for `JRE_LOCATION` property in *ORACLE_HOME*/oui/bin/runconfig.sh to *ORACLE_HOME*/jdk/jre, and then run `configToolAllCommands` again.

Bug 10053840

`clone` or `addNode` operations may fail if the installer session was restarted while choosing to apply software updates in source Oracle home.

Workaround: Change the value for `JRE_LOCATION` property in *ORACLE_HOME*/oui/oraparam.ini to ../../jdk/jre.

Bug 10053090

During the prerequisite checks step of installation, the OUI may display an error with text `-1`, and the prerequisite page result shows `No results`.

Workaround: Reexecute the prerequisites by clicking on the `Check Again` button.

Bug 8729326

When upgrading to 11.2 Clusterware, the Installer invokes ASMCA in silent mode to upgrade Oracle ASM into Oracle Grid Infrastructure for a cluster home. Oracle ASM upgrade is handled in rolling fashion when upgrading from 11.1.0.7. Prior versions of Oracle ASM instances are upgraded in non-rolling fashion and Oracle ASM-based databases are bounced without any prior warning.

Workaround: You can plan your database outage to be the point where you acknowledge the Installer prompt after executing `root.sh` on all nodes. At this point, CRS is upgraded in rolling fashion and the Installer will be calling ASMCA to upgrade Oracle ASM, which will bounce databases as part of Oracle ASM upgrade.

Bug 8666656

The Oracle Universal Installer (OUI) `runInstaller` script that resides in the Oracle home (`ORACLE_HOME`/oui/bin/runInstaller) cannot be used to install the 11.2.0.1 releases of Oracle Database, Oracle Grid Infrastructure for a cluster, and Oracle Database Client.

Workaround: Use Oracle Universal Installer on the respective 11.2.0.1.0 product media to install each product.

Bug 8638708

If you select the database configuration Desktop Class in Oracle Universal Installer (OUI), listener and database control are configured with `'localhost'` as the host name. The Oracle Enterprise Manager Database Control `start` and `stop` operations using `emctl` may fail.

Workaround: For Database Control start and stop operations that use `emctl` in that home, set the `ORACLE_HOSTNAME` environment variable to `'localhost'`.

Bug 8407818

After adding a new node to a shared Oracle database home using `addNode.sh`, the `/etc/oratab` on the newly added node gets an entry of the source database name that exists on the source node from where `addNode.sh` was run. The `/etc/oratab` file on the new node is supposed to get the database entry after the database instance is added for the new node using DBCA.

Workaround: Before invoking DBCA from the source node to add a new database instance for the new node, open the `/etc/oratab` file on the new node using an editor and remove the entry made for the source database name.

2.38.16 Vendor and Operating System Known Bugs

Bug 8256753

A connect using `SCAN` and `EZCONNECT` on one client machine can be requested to use a specific `SCAN` listener. Therefore, load balancing by round-robin DNS is not possible.

Workaround: Connect to a database using the following configuration specifying `LOAD_BALANCE=on` in `tnsnames.ora`:

```
ORCL =
```

```
(DESCRIPTION =
  (LOAD_BALANCE=on)
  (ADDRESS = (PROTOCOL = TCP)(HOST = stscan1)(PORT = 1521))
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = srv.world)
  )
)
```

2.38.17 Oracle Wallet Manager Known Bugs

Bug 9880435

Oracle Wallet Manager fails to upload wallet to Directory service when the wallet password and the directory user password are different.

Workaround: Use the same password for the wallet and the directory user.

2.38.18 Oracle Warehouse Builder Known Bugs

Refer to "Open Bugs and Known Issues" in *Oracle Warehouse Builder Release Notes*.

Bug 10010873

Oracle Warehouse Builder (OWB) 11.2.0.1 Design Center client is not compatible with OWB 11.2.0.2 server-side components. Using OWB 11.2.0.1 Design Center to connect to an 11.2.0.2 OWB repository can cause unpredictable behavior, possibly including repository corruption.

Workaround: To avoid these conflicts, configure OWB 11.2.0.1 to check the version of the target repository at connection time.

Edit `OWB_CLIENT_HOME/owb/bin/admin/Preference.properties` on each client installation, and update the `OverrideRuntimeVersionCheck` and `OverrideRepositoryVersionCheck` preferences as follows:

```
OverrideRuntimeVersionCheck=false
OverrideRepositoryVersionCheck=false
```

This change to `Preference.properties` should be applied to the 11.2.0.1 Design Center installation on every client host.

To use OWB Design Center with 11.2.0.2 databases, use the OWB 11.2.0.2 Linux 32-bit Design Center until the OWB 11.2.0.2 Windows 32-bit Design Center is available.

If you cannot use OWB 11.2.0.2 Linux 32-bit Design Center client in your environment, contact Oracle Support Services for further assistance.

2.38.19 Oracle XML Database Known Bugs

Bug 10018899

Because operator arguments are not checked for tracking changes to database objects, an existing result cache cannot be invalidated when one of those objects is updated.

The `xmlexists` operator is affected by this problem as well as any operator that accepts a database object as an argument.

Workaround: Disable or flush the result cache for queries that involve such operators.

Bug 9586264

In order to fully optimize some `XMLQUERY` or `XMLTABLE` queries, `OPTIMIZER_FEATURE_ ENABLE` should be set to 11.1.0.6 or above.

Workaround: None.

3 Readme Information for Oracle Database 11*g* Release 2 (11.2.0.3)

> **Note:**   If you are on Oracle Database 11*g* Release 2 (11.2.0.3), this is the Readme section that you need to read.

This section of the Readme contains the following sub-sections:

3.1 Compatibility, Upgrading, Downgrading, and Installation

For late-breaking updates and best practices about preupgrade, post-upgrade, compatibility, and interoperability discussions, see Note 785351.1 on My Oracle Support (at https://support.oracle.com) that links to the "Upgrade Companion" web site for Oracle Database 11*g* Release 2.

> **Caution:** After installation is complete, do not manually remove or run `cron` jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running. If you remove these files, then Oracle software can encounter intermittent hangs. Oracle Grid Infrastructure for a cluster and Oracle Restart installations fail with the following error:
>
> ```
> CRS-0184: Cannot communicate with the CRS daemon.
> ```

3.1.1 Downgrading Release 11.2.0.3 to 11.2.0.2 Results in an Error When You Run catdwgrd.sql

When downgrading from release 11.2.0.3 to 11.2.0.2, the following error is raised when you run `@catdwgrd.sql` (reference Bug 11811073):

```
ORA-20000: Upgrade from version 11.2.0.2.0 cannot be downgraded to version
```

Apply patch 11811073 for release 11.2.0.2 which provides an updated version of `catrelod.sql`. Applying this patch must be done prior to executing `@catdwgrd.sql` in the 11.2.0.2 environment.

3.1.2 Downgrade to Release 11.1.0.6

If you anticipate downgrading back to release 11.1.0.6, then apply the patch for Bug 7634119. This action avoids the following `DBMS_XS_DATA_SECURITY_EVENTS` error:

```
PLS-00306: wrong number or types of arguments in call
to 'INVALIDATE_DSD_CACHE' DBMS_XS_DATA_SECURITY_EVENTS
PL/SQL: Statement ignored
```

Apply this patch prior to running `catrelod.sql`.

3.1.3 Invalid Objects Seen After Downgrading from 11.2.0.3 to 11.2 or 10.2

After downgrading from 11.2.0.3 to 11.1 or 10.2, the following invalid object may be seen:

```
CTX_FILTER_CACHE_STATISTICS (synonym)
CTX_FILTER_CACHE_STATISTICS (view)
```

In the higher release of Oracle, after running `@catdwgrd.sql` and before running `@catrelod.sql`, issue the following two commands:

```
SQL> drop public synonym ctx_filter_cache_statistics;
SQL> drop view ctx_filter_cache_statistics;
```

3.1.4 Performing -force Upgrade Results in an Incorrect Grid Home Node List in Inventory

When a node crash occurs during an upgrade, a `-force` upgrade can be performed to upgrade a partial cluster minus the unavailable node (reference Bug 12933798).

After performing a -force upgrade, the node list of the Grid home in inventory is not in sync with the actual Oracle Grid Infrastructure deployment. The node list still contains the unavailable node. Because the node list in inventory is incorrect, the next upgrade or node addition, and any other Oracle Grid Infrastructure deployment, fails.

After performing a -force upgrade, manually invoke the following command as a CRS user:

```
$GRID_HOME/oui/bin/runInstaller -updateNodeList "CLUSTER_NODES={comma_separated_
alive_node_list}" ORACLE_HOME=$GRID_HOME CRS=true
```

3.1.5 Downgrading a Database Having Database Control Configured

Consider the following when downgrading the database while having Database Control configured (reference Bug 9922349):

1. If you are upgrading from 11.2.0.1 to 11.2.0.3 and then plan to downgrade to 11.2.0.1, you need to apply the following patches in order to downgrade Database Control as part of the database downgrade:

   - 11.2.0.1 PSU2 bundle

   - One-off patch for Bug 8795792

   Without these patches, the emdwgrd utility would fail with IMPORT (impdp) errors when restoring Database Control data.

2. When running emdwgrd on 11.2.0.1 Oracle RAC databases, you may need to pass an additional parameter, -serviceAlias, if you do not have system identifier (SID) aliases defined in tnsnames.ora. This is also needed for single instance if SID and database names are different. For example:

   ```
   emdwgrd -save [-cluster] -sid SID [-serviceAlias tns_alias] -path save_
   directory
   emdwgrd -restore -tempTablespace TEMP [-cluster] -sid SID [-serviceAlias tns_
   alias] -path save_directory
   ```

3. In the case of in-place downgrade from 11.2.0.3 to 11.2.0.1 using the same Oracle home, you do not need to run emca -restore before running emdwngrd -restore.

3.1.6 Deinstallation Restrictions

The following sections describe deinstallation and deconfiguration restrictions. See for additional information.

3.1.6.1 Deinstall Upgraded 11.2 Oracle RAC and Oracle Grid Infrastructure for a Cluster Homes

After you deconfigure and deinstall an upgraded Oracle Database 11*g* Release 2 (11.2) Oracle RAC home and to deconfigure and deinstall an 11.2 Oracle Grid Infrastructure for a cluster home, you must detach any pre-11.2 Oracle RAC software homes from the central Inventory (reference Bug 8666509).

Detach the pre-11.2 Oracle RAC homes from the central inventory with the following command:

```
ORACLE_HOME/oui/bin/runInstaller -detachHome ORACLE_HOME_NAME=pre-11.2_ORACLE_
HOME_NAME ORACLE_HOME=pre-11.2_ORACLE_HOME
```

3.1.7 Multiple Interconnects and Oracle ACFS

If you have Oracle ACFS file systems on Oracle Grid Infrastructure for a cluster 11g release 2 (11.2.0.1), you upgrade Oracle Grid Infrastructure to 11*g* release 2 (11.2.0.2) or 11*g* release 2 (11.2.0.3), and you take advantage of Redundant Interconnect Usage and

add one or more additional private interfaces to the private network, then you must restart the Oracle ASM instance on each upgraded cluster member node (reference Bug 9969133).

3.1.8 Oracle ASM Rolling Upgrade

Oracle Automatic Storage Management (Oracle ASM) rolling upgrade check does not allow rolling upgrade to be done from 11.1.0.6 to any later release (reference Bug 6872001). The following message is reported in the alert log:

```
Rolling upgrade from 11.1.0.6 (instance instance-number) to 11.x.x.x is not
supported
```

ORA-15156 is signalled by LMON which will then terminate the instance.

When trying to upgrade Oracle ASM from 11.1.0.6 to a later release of Oracle ASM, apply the patch for this bug to 11.1.0.6 instances before rolling upgrade starts. This patch can be applied to 11.1.0.6 instances in a rolling fashion.

3.1.9 The ora.acfs.registry Resource May Be in an Inconsistent State After an Upgrade to 11.2.0.3.0 from a Previous Release

TARGET and STATE for ora.registry.acfs will be set to either ONLINE if the Oracle Automatic Storage Management Cluster File System (Oracle ACFS) registry resource existed in the previous release, or set to OFFLINE if the Oracle ACFS registry resource did not exist in the previous release (reference Bug 12812838 and Bug 9878976).

To disable Oracle ACFS, enter the command acfsroot disable which will set ora.registry to STATE OFFLINE, TARGET OFFLINE after a CRS stack restart.

3.1.10 Oracle ACFS Registry May Be in an Inconsistent State After Installing or Upgrading to 11.2.0.3.0 or After An Oracle Clusterware Restart

If Oracle ASM is not used as the voting disk and quorum disk, the Oracle Automatic Storage Management Cluster File System (Oracle ACFS) registry resource will report OFFLINE after an install (reference Bug 9876173 and Bug 9864447). This occurs because the Oracle ACFS registry requires that Oracle ASM be used in order to provide Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes.

3.1.11 Upgrading a Database With Oracle Data Mining (ODM)

If you upgrade a database with the Data Mining option from 11.2.0.1 to 11.2.0.3, make sure that the DMSYS schema does not exist in your 11.2.0.1 database. If it does, you should drop the DMSYS schema and its associated objects from the database as follows:

```
SQL> CONNECT / AS SYSDBA;
SQL> DROP USER DMSYS CASCADE;
SQL> DELETE FROM SYS.EXPPKGACT$ WHERE SCHEMA = 'DMSYS';
SQL> SELECT COUNT(*) FROM DBA_SYNONYMS WHERE TABLE_OWNER = 'DMSYS';
```

If the above SQL returns non-zero rows, create and run a SQL script as shown in the following example:

```
SQL> SET HEAD OFF
SQL> SPOOL dir_path/DROP_DMSYS_SYNONYMS.SQL
SQL> SELECT 'Drop public synonym ' ||'"'||SYNONYM_NAME||'";' FROM DBA_SYNONYMS
WHERE TABLE_OWNER = 'DMSYS';
SQL> SPOOL OFF
SQL> @dir_path/DROP_DMSYS_SYNONYMS.SQL
SQL> EXIT;
```

If you upgrade a database from 10*g* to 11.2, all Data Mining metadata objects are migrated from `DMSYS` to `SYS`. After the upgrade, when you determine that there is no need to perform a downgrade, set the initialization parameter `COMPATIBLE` to 11.2 and drop the `DMSYS` schema and its associated objects as described above.

3.1.12 Oracle ACFS File Systems Must be Manually Dismounted Prior to Upgrade, Deinstallation, or Direct Shutdown of Oracle Clusterware or Oracle ASM

Oracle ACFS file systems must be unmounted, on any given node, prior to upgrade, deinstallation, or direct shutdown of Oracle Clusterware or Oracle ASM on that node (reference Bug 8594128 and Bug 12726434). Use `srvctl stop filesystem` and `umount` on UNIX or `srvctl stop filesystem` and `acfsdismount` on Windows.

When upgrading, all Oracle ACFS file systems must be stopped before beginning the upgrade. For UNIX, this can happen on a node-by-node basis before beginning the upgrade on each node. For Windows, Oracle ACFS file systems must be unmounted across the cluster due to the unattended nature of the rolling upgrade on Windows.

Use the `lsof` and `fuser` commands (Linux and UNIX) or the `handle` and `wmic` commands (Windows) to identify processes which are active on the Oracle ACFS file systems. To ensure that these processes are no longer active, dismount all Oracle ACFS file systems and issue Oracle Clusterware shutdown. Otherwise errors may be issued during Oracle Clusterware shutdown relating to activity on Oracle ACFS file systems which will stop the successful shutdown of Oracle Clusterware.

3.1.13 catrelod.sql Fails if the Time Zone File Version Used by the Database Does Not Exist in Oracle Home

The following error is returned when `catrelod.sql` is run as part of the downgrade process if you previously installed a recent version of the time zone file and used the `DBMS_DST` PL/SQL package to upgrade `TIMESTAMP WITH TIME ZONE` data to that version (reference Bug 9803834):

```
ORA-00600: internal error code, arguments: [qcisSetPlsqlCtx:tzi init],
 [], [], [], [], [], [], [], [], [], [], []
```

See Step 2 of 'Downgrade the Database' in Chapter 6 of the *Oracle Database Upgrade Guide* for more details.

If you previously installed a recent version of the time zone file and used the `DBMS_DST` PL/SQL package to upgrade `TIMESTAMP WITH TIME ZONE` data to that version, then you must install the same version of the time zone file in the release to which you are downgrading. For example, the latest time zone files that are supplied with Oracle Database 11*g* Release 2 (11.2) are version 14. If, after the database upgrade, you had used `DBMS_DST` to upgrade the `TIMESTAMP WITH TIME ZONE` data to version 14, then install the version 14 time zone file in the release to which you are downgrading. This ensures that your `TIMESTAMP WITH TIME ZONE` data is not logically corrupted during retrieval. To find which version your database is using, query `V$TIMEZONE_FILE`.

Also see the *Oracle Database Globalization Support Guide* for more information on installing time zone files.

3.1.14 Data Pump Export Operations Do Not Work if the DMSYS Schema is Not Removed as Part of the Upgrade to 11.2.0.3

Data Pump Export operations do not work if the `DMSYS` schema is not removed as part of the upgrade to release 11.2.0.3 (reference Bug 10007411). The reported error is similar to the following:

```
Processing object type SCHEMA_EXPORT/TABLE/TABLE_DATA
ORA-39126: Worker unexpected fatal error in
KUPW$WORKER.GET_TABLE_DATA_OBJECTS []
ORA-31642: the following SQL statement fails:
BEGIN "DMSYS"."DBMS_DM_MODEL_EXP".SCHEMA_CALLOUT(:1,0,1,'10.01.00.05.00');
END;
ORA-06512: at "SYS.DBMS_SYS_ERROR", line 86
ORA-06512: at "SYS.DBMS_METADATA", line 1245
ORA-04063: package body "DMSYS.DBMS_DM_MODEL_EXP" has errors
ORA-06508: PL/SQL: could not find program unit being called:
"DMSYS.DBMS_DM_MODEL_EXP"
```

The pre-upgrade checks for 11.2.0.3 report the action that should be taken before the upgrade:

```
The DMSYS schema exists in the database. Prior to performing an upgrade Oracle
recommends that the DMSYS schema, and its associated objects be removed from the
database. Refer to the Oracle Data Mining Administration Guide for the
instructions on how to perform this task.
```

Until this step is taken, Data Pump Export will not work.

3.1.15 Upgrading With Objects in the Recycle Bin

The recycle bin must be empty during an upgrade to avoid possible `ORA-00600` deadlock errors, as well as to minimize the time required to perform the upgrade (reference Bug 8632581).

To avoid this deadlock, use the `PURGE DBA_RECYCLEBIN` statement to remove items and their associated objects from the recycle bin and release their storage space prior to upgrading your database.

3.1.16 INVALID Materialized View

A materialized view has a status of `INVALID` after both `@catupgrd.sql` and `@utlrp.sql` have been run (reference Bug 12530178). You can see this using the following command:

```
SELECT object_name, object_id, owner FROM all_objects WHERE
object_type='MATERIALIZED VIEW' and status='INVALID';

OBJECT_NAME                     OBJECT_ID   OWNER
----------------------------- ----------- -------

FWEEK_PSCAT_SALES_MV                51062   SH
```

If, after running both `@catupgrd.sql` to upgrade the database and `@utlrp.sql` to recompile invalid objects, there still exists an invalid materialized view, then issue the following SQL statement:

```
ALTER MATERIALIZED VIEW sh.FWEEK_PSCAT_SALES_MV COMPILE;
```

3.1.17 rootupgrade.sh Script Might Fail

During an upgrade from release 11.2.0.2 to 11.2.0.3, if the `rootupgrade.sh` script exits while the CRS stack is shutting down, re-running the `rootupgrade.sh` script might fail (reference Bug 12721330).

Manually start the old CRS stack and then re-run the `rootupgrade.sh` script.

3.1.18 ora.ons Status May Show UNKNOWN

After upgrading from release 11.2.0.2 to 11.2.0.3, you may see `ora.ons` status showing `UNKNOWN` with the explanation being `CHECK TIMED OUT` (reference Bug 12861771).

The workaround is to kill the Oracle Notification Services (ONS) process and execute `srvctl start nodeapps`.

3.1.19 PRVF-10037 Error During Installation of Oracle Grid Infrastructure 11.2.0.3 software

During storage verification when performing an 11.2.0.3 install or upgrade of Grid Infrastructure, the following error message may be displayed:

```
PRVF-10037 : Failed to retrieve storage type for "<devicepath>" on node "<node>"
Could not get the type of storage
```

This issue could be the result of the CVUQDISK package not being installed or an incorrect version of the CVUQDISK package being installed. Please ensure that the correct package (typically `cvuqdisk-1.0.9-1.rpm`) is installed and if it is not, install this version of the package and run the install or upgrade again (reference bug 12881575).

Perform the following procedure to install the CVUQDISK package:

1. Login as the `root` user.

2. Copy the package, `cvuqdisk-1.0.9-1.rpm`, to a local directory. You can find this package in the rpm subdirectory of the top-most directory in the Oracle Grid Infrastructure installation media. For example, you can find `cvuqdisk-1.0.9-1.rpm` in the directory /*mountpoint*/clusterware/rpm/ where *mountpoint* is the mounting point for the disk on which the Oracle Grid Infrastructure installation media is located.

3. Set the environment variable `CVUQDISK_GRP` to the operating system group that should own the CVUQDISK package binaries. It is recommended that you set this group to the installation group. If `CVUQDISK_GRP` is not set, then by default the `oinstall` group is used for the group that owns the CVUQDISK package binaries.

4. Determine whether previous versions of the CVUQDISK package are installed by running the command `rpm -q cvuqdisk`. If you find previous versions of the CVUQDISK package, then remove them by running the following command, where *previous_version* is the identifier of the previous CVUQDISK version:

   ```
   rpm -e cvuqdisk previous_version
   ```

5. Install the latest CVUQDISK package by running the following command:

   ```
   rpm -iv cvuqdisk-1.0.9-1.rpm
   ```

3.2 Features Not Available or Restricted in 11.2.0.3

The following is a list of components that are not available or are restricted in this release of Oracle Database 11*g* Release 2 (11.2.0.3):

- Starting with Oracle Database 11*g* Release 2 (11.2.0.3), the Data Mining Java API is deprecated. For information about the Data Mining Java API, see Chapter 7 of *Oracle Data Mining Application Developer's Guide*. See Section 3.8.5, "Data Mining Features Not Available or Deprecated with Oracle Database 11g" for information about other deprecated features of Oracle Data Mining.

- Security-Enhanced Linux (SELinux) is not supported on Oracle Automatic Storage Management Cluster File System (Oracle ACFS) file systems (reference Bug 12754448). (SELinux was also not supported on Oracle ACFS in release 11.2.0.1 and release 11.2.0.2.)

- Certain Oracle Text functionality based on third-party technologies, including `AUTO_LEXER` and `CTX_ENTITY`, have been disabled in this release (reference Bug 12618046). For `BASIC_LEXER`, the usage of the `INDEX_STEMS` attribute values that depend on third-party technologies, is also affected. If this impacts an existing application, contact Oracle Support Services for guidance.

- Oracle Data Mining now supports a new release of Oracle Data Miner. The earlier release, Oracle Data Miner Classic, is still available for download on OTN, but it is no longer under active development. For information about the new release, Oracle Data Miner 11*g* Release 2, go to
  http://www.oracle.com/technetwork/database/options/odm/index.html.

- All Oracle Grid Infrastructure patch set upgrades must be out-of-place upgrades, in which case you install the patch set into a new Oracle Grid home (reference Bug 10210246). In-place patch set upgrades are not supported.

- Oracle Database release 11.2.0.1 or 11.2.0.2 upgrade to Oracle Clusterware release 11.2.0.3 is not supported if the 11.2.0.1 or 11.2.0.2 release of Oracle Grid Infrastructure for a cluster is installed in a non-shared Oracle home and the 11.2.0.3 release of Oracle Grid Infrastructure for a cluster is installed in a shared Oracle home (reference Bug 10074804). The original and upgraded releases of Oracle Clusterware should both be installed in either a shared or non-shared Oracle home.

3.3 Default Behavior Changes

This section describes some of the differences in behavior between Oracle Database 11*g* Release 2 (11.2) and previous releases. The majority of the information about upgrading and downgrading is already included in the *Oracle Database Upgrade Guide*.

3.3.1 Non-Uniform Memory Access Optimizations and Support Disabled in 11.2

With Oracle Database 11*g* Release 2 (11.2), non-uniform memory access support is disabled by default. This restriction applies to all platforms and operating systems (reference Bug 8450932).

Non-uniform memory access optimizations and support in the Oracle Database are only available for specific combinations of Oracle version, operating systems, and platforms. Work with Oracle Support Services and your hardware vendor to enable non-uniform memory access support.

3.3.2 FILE_ACCESS_ROLE Default Behavior Change

The default behavior of the CTX system parameter `FILE_ACCESS_ROLE` has changed (reference Bug 8360111). Customers with existing Oracle Text indexes that use the file or URL datastore must take action to continue to use the indexes without error. The changes are as follows:

- If `FILE_ACCESS_ROLE` is null (the default), then access is not allowed. By default, users who were previously able to create indexes of this type will not be able to create these indexes after the change.

- `FILE_ACCESS_ROLE` is now checked for index synchronization and document service operations. By default, users will not be able to synchronize indexes of this

type or use document service calls such as `ctx_doc.highlight` who were allowed to prior to this change.

- Only SYS will be allowed to modify `FILE_ACCESS_ROLE`. Calling `ctx_adm.set_parameter (FILE_ACESS_ROLE, role_name)` as a user other than SYS will now raise the new error:

  `DRG-10764: only SYS can modify FILE_ACCESS_ROLE`

- Users can set `FILE_ACCESS_ROLE` to `PUBLIC` to explicitly disable this check (which was the previous default behavior).

3.3.3 Use of the Append Hint Runs Out of Memory When Loading Many Partitions

Use of direct-path `INSERT` to load a large number of partitions can exceed memory limits, especially when data compression is specified (reference Bug 6749894). Starting in 11.2, the number of partitions loaded at the same time will be limited, based on the `PGA_AGGREGATE_TARGET` initialization parameter, to preserve memory. Rows that are not stored in the partitions that are currently being loaded are saved in the temporary tablespace. After all rows are loaded for the current set of partitions, other partitions are loaded from rows that are saved in the temporary tablespace.

This behavior helps prevent the direct-path `INSERT` from terminating because of insufficient memory.

3.4 Database Security

Note the following changes in Database Security.

3.4.1 Protection Against Password-Guessing Attacks

Starting in 11.2.0.3, the database authentication protocol has been strenghtened against certain types of password-guessing attacks. In order to force the use of this more secure behavior, both the database client and the database server must be upgraded to release 11.2.0.3 or later and the server's `SQLNET.ORA` configuration file `SQLNET.ALLOWED_LOGON_VERSION` parameter should be set to a value of `12` to force the new protocol behavior.

If the SQLNET.ALLOWED_LOGON_VERSION initialization parameter is set to 12 on the server without upgrading all clients to release 11.2.0.3 or later, password authentication will fail with an `ORA-28040: No matching authentication protocol` error because older clients do not support this new protocol behavior.

If the SQLNET.ALLOWED_LOGON_VERSION initialization parameter is set to 12 on the server without upgrading the server to release 11.2.0.3, password authentication will fail with an `ORA-28040: No matching authentication protocol` error because older server software does not support this new protocol behavior.

3.5 Oracle Automatic Storage Management (Oracle ASM)

The following sections describe information pertinent to Oracle Automatic Storage Management (Oracle ASM) in Oracle Database 11*g* Release 2 (11.2).

3.5.1 Oracle Database Data Files and Oracle ACFS

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is the preferred file manager for non-database files. It is optimized for general purpose files. Oracle ACFS does not support any file type that can be directly stored in Oracle ASM.

However, starting with Oracle Automatic Storage Management 11*g* Release 2 (11.2.0.3), Oracle ACFS supports, without snapshots, RMAN backups, archive logs, and Data Pump dumpsets.

3.5.2 Oracle Homes on Oracle ACFS Supported Starting With Release 11.2

Placing Oracle homes on Oracle ACFS is supported starting with Oracle Database release 11.2 (reference Bug 10144982). Oracle ACFS can result in unexpected and inconsistent behavior if you attempt to place Oracle homes on Oracle ACFS on database versions prior to 11.2.

3.6 Java and Web Services

Note the following items when working with Java.

3.6.1 Restricted Access to com.sun.imageio.* Packages

Access is restricted to `com.sun.imageio.*` packages for applications that need to use these routines (reference Bug 12583785). If you need to use this package and you encounter a permission error, the SQL statement needed to grant the permission is displayed. Execute the statement.

3.7 Oracle Application Express

To learn more about Oracle Application Express, refer to the *Oracle Application Express Release Notes* and the *Oracle Application Express Installation Guide*.

3.8 Oracle Data Mining

Note the following items when working with Oracle Data Mining.

3.8.1 Scoring Capabilities in Oracle Exadata Storage Server Layer

Oracle Data Mining scoring functions in Oracle Database 11*g* Release 2 are also available in Oracle Exadata Storage Server Software. Scoring capabilities in the storage layer permit very large data sets to be mined quickly, thus further increasing the competitive advantage already gained from Oracle in-database analytics. For information about Oracle Exadata Storage Server Software, see http://www.oracle.com/technology/products/bi/db/exadata/index.html.

3.8.2 Installation of Oracle Data Mining

The Data Mining Option, as an embedded feature of the database, is automatically installed with the Oracle Enterprise Edition Database. When installing the database with the Data Mining Option, choose the Data Warehouse configuration type for the most appropriate default initialization parameters.

3.8.3 Oracle Data Mining Models

In Oracle Database 11*g*, Data Mining models are implemented as data dictionary objects in the `SYS` schema. The `DMSYS` schema no longer exists.

Data Mining users must have the `CREATE MINING MODEL` privilege to create mining models in their own schema. Additional privileges are required for other data mining activities, as described in the *Oracle Data Mining Administrator's Guide*.

New data dictionary views for Oracle Data Mining were introduced in Oracle Database 11*g* Release 1 (11.1):

- USER/ALL/DBA_MINING_MODELS

- USER/ALL/DBA_MINING_MODEL_ATTRIBUTES

- USER/ALL/DBA_MINING_MODEL_SETTINGS

3.8.4 Demo Programs for the Data Mining APIs

Demo programs that illustrate the Data Mining APIs (PL/SQL and Java) are installed with Oracle Database Examples. Instructions are in the *Oracle Data Mining Administrator's Guide*.

3.8.5 Data Mining Features Not Available or Deprecated with Oracle Database 11g

- The Oracle Data Mining Scoring Engine Option, a separately installed database option in Oracle Database 10*g*, is not available in Oracle Database 11*g*. However, all functionality of the Data Mining Scoring Engine Option is offered in the Data Mining Option of Oracle Database 11*g*.

- The Basic Local Alignment Search Tool (BLAST), previously supported by Oracle Data Mining, is not available in Oracle 11*g*.

- Starting with release 11.2.0.3 of Oracle Database, the Data Mining Java API is deprecated. For information about the Data Mining Java API, see Chapter 7 of *Oracle Data Mining Application Developer's Guide*.

- Oracle Data Mining now supports a new release of Oracle Data Miner. The earlier release, Oracle Data Miner Classic, is still available for download on Oracle Technology Network (OTN), but it is no longer under active development. For information about the new release, Oracle Data Miner 11*g* Release 2, go to http://www.oracle.com/technetwork/database/options/odm/index.html.

3.9 Oracle Database Vault

Note the following items when working with Oracle Database Vault.

3.9.1 New Language for Oracle Database Vault

To add a new language for Oracle Database Vault, connect as a user and specify the `DV_ADMIN` or `DV_OWNER` role. Run the following command:

```
DVSYS.DBMS_MACADM.ADD_NLS_DATA("<language>");
```

Where `<language>` is one of the following:

```
ENGLISH
GERMAN
SPANISH
FRENCH
ITALIAN
JAPANESE
KOREAN
BRAZILIAN PORTUGUESE
SIMPLIFIED CHINESE
TRADITIONAL CHINESE
```

3.10 Oracle Grid Infrastructure for a Cluster

Note the following items when working with Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM), which are installed with an Oracle Grid Infrastructure for a cluster installation.

3.10.1 Oracle ACFS and Oracle Clusterware Stack Shut Down

When attempting to shut down Oracle Clusterware, the Oracle Clusterware stack may report that it did not successfully stop on selected nodes (reference Bug 8651848). If the database home is on Oracle ACFS, then you may receive the following error:

```
CRS-5014: Agent orarootagent.bin timed out starting process acfsmount for action
```

This error can be ignored.

Alternatively, the Oracle Clusterware stack may report that it did not successfully stop on selected nodes due to the inability to shut down the Oracle ACFS resources. If this occurs, take the following steps:

- Ensure that all file system activity to Oracle ACFS mount points is quiesced by shutting down programs or processes and retry the shutdown.

- If the `ora.registry.acfs` resource check function times out, or the resource exhibits a state of `UNKNOWN` or `INTERMEDIATE`, then this may indicate an inability to access the Oracle Cluster Registry (OCR). The most common cause of this is a network failure. The commands `acfsutil registry` and `ocrcheck` may give you a better indicator of the specific error. Clear this error and attempt to stop Oracle Clusterware again.

3.10.2 In-Place Patch Set Upgrades Are Not Supported

All Oracle Grid Infrastructure patch set upgrades must be out-of-place upgrades, in which case you install the patch set into a new Oracle Grid home (reference Bug 10210246). In-place patch set upgrades are not supported.

3.11 Oracle Multimedia

The name Oracle *inter*Media was changed to Oracle Multimedia in Oracle Database 11*g* Release 1 (11.1). The feature remains the same, only the name has changed. References to Oracle *inter*Media were replaced with Oracle Multimedia, however some references to Oracle *inter*Media or *inter*Media may still appear in graphical user interfaces, code examples, and related documents in the Oracle Database documentation library for 11*g* Release 2 (11.2).

For additional information, refer to the Oracle Multimedia Readme file located at:

```
ORACLE_HOME/ord/im/admin/README.txt
```

3.12 Oracle Net Listener

Note the following items when working with Oracle Net Listener.

3.12.1 Deprecation of SNMP Support in Oracle Net Listener

Oracle is deprecating SNMP support in Oracle Net Listener in Oracle Database 11*g* Release 2 (11.2). Oracle recommends not using SNMP in new implementations.

3.13 Oracle ODBC Driver

The Oracle ODBC Driver Readme file is located at:

`ORACLE_HOME/odbc/html/ODBCRelnotesUS.htm`

3.14 Oracle Real Application Clusters

Note the following items when working with Oracle RAC.

3.14.1 Using the Same Cluster for Administrator-Managed and Policy-Managed Databases

If you are creating an administrator-managed database on a cluster that already hosts policy-managed databases, then you must carefully select the nodes for the administrator-managed database (reference Bug 10027250). This is because the nodes that you select for an administrator-managed database that are in policy-managed server pools will be moved into the Generic server pool as part of this process. If you select nodes that already run other policy-managed database instances, then DBCA prompts you with a message that lists the instances and services that will be shut down when DBCA creates the administrator-managed database. If you select the `Yes` button on the dialog box when DBCA asks "`Do you want to continue?`," then your policy-managed database instances and services will be shut down as a result of the administrator-managed database creation process.

Note: This is also true if you use the `srvctl add instance` command, which gives a similar error message indicating that the databases would be shut down. If you also use the force option (`-f`) with the `srvctl add instance` command, then this is the same as choosing `Yes` on the DBCA dialog. Doing this shuts down any policy-managed databases that are running on the node before moving the node into the Generic server pool.

3.14.2 Moving root Owned Binaries that need setuid to Local Nodes from NFS

If you install an Oracle RAC database into a shared Oracle home on an NFS device, then you must copy the ORADISM binary (`oradism`) into a local directory on each node (reference Bug 7210614).

To move `oradism`, take the following steps:

1. Copy the `ORACLE_HOME/bin/oradism` binary to an identical directory path on all cluster nodes. The path (for example, `/u01/local/bin` in the example in Step 2) must be local and not NFS. For example:

    ```
    cp -a ORACLE_HOME/bin/oradism /u01/local/bin
    ```

2. Run the following commands, as the root user, to set ownership and permissions of the `oradism` executable:

    ```
    $ chown root /u01/local/bin/oradism
    $ chmod 4750 /u01/local/bin/oradism
    ```

3. Create a symbolic link from the NFS shared home to the local `oradism` directory path. This needs to be done from one node only. Each node can then reference its own `oradism` using the `symlink` from the shared Oracle home. For example:

```
$ cd /nfs/app/oracle/product/11.2.0/db_1/bin
$ rm -f oradism
$ ln -s /u01/local/bin/oradism oradism
```

4. If the Oracle home is an Oracle Database home directory, then repeat steps 1-3 for other binaries such as `extjob`, `jssu`, `nmb`, `nmhs` and `nmo`. You do not need to perform this step if the Oracle home is an Oracle Grid Infrastructure home directory.

3.15 Oracle Spatial

The Oracle Spatial readme file supplements the information in the following manuals: *Oracle Spatial Developer's Guide*, *Oracle Spatial Topology and Network Data Models Developer's Guide*, and *Oracle Spatial GeoRaster Developer's Guide*. The Oracle Spatial readme file is located at:

`ORACLE_HOME/md/doc/README.txt`

3.16 Oracle SQL Developer

The Oracle SQL Developer readme file is located at:

`ORACLE_HOME/sqldeveloper/readme.html`

3.17 Oracle Text

Note the following items when working with Oracle Text. You should also check entries for the *Oracle Text Application Developer's Guide* in the Documentation Addendum.

3.17.1 Change to Supported Features

Certain Oracle Text functionality based on third-party technologies, including `AUTO_LEXER` and `CTX_ENTITY`, have been disabled in this release (reference Bug 12618046). For `BASIC_LEXER`, the usage of the `INDEX_STEMS` attribute values that depend on third-party technologies, is also affected. If this impacts an existing application, contact Oracle Support Services for guidance.

3.17.2 Oracle Text Supplied Knowledge Bases

An Oracle Text knowledge base is a hierarchical tree of concepts used for theme indexing, `ABOUT` queries, and deriving themes for document services. The following Oracle Text services require that a knowledge base be installed:

- Index creation using a `BASIC_LEXER` preference where `INDEX_THEMES=YES`

- `SYNC`ing of an index where `INDEX_THEMES=YES`

- `CTX_DOC.THEMES`

- `CTX_DOC.POLICY_THEMES`

- `CTX_DOC.GIST`

- `CTX_DOC.POLICY_GIST`

- `CTX_QUERY.HFEEDBACK`

- `CTX_QUERY.EXPLAIN`, if using `ABOUT` or `THEMES` with `TRANSFORM`

- `CTX_DOC.SNIPPET` (if using the `ABOUT` operator)

- `CTX_DOC.POLICY_SNIPPET` (if using the `ABOUT` operator)

- `CONTAINS` queries that use `ABOUT` or `THEMES` with `TRANSFORM`

- The Knowledge Base Extension Compiler, `ctxkbtc`

- Clustering and classification services, if themes are specified

If you plan to use any of these Oracle Text features, then you should install the supplied knowledge bases, English and French, from the Oracle Database Examples media, available for download on OTN.

Note that you can extend the supplied knowledge bases, or create your own knowledge bases, possibly in languages other than English and French. For more information about creating and extending knowledge bases, refer to the *Oracle Text Reference*.

For information about how to install products from the Oracle Database Examples media, refer to the *Oracle Database Examples Installation Guide* that is specific to your platform.

3.18 Oracle Warehouse Builder

For additional information about Oracle Warehouse Builder (OWB) in Oracle Database 11*g* Release 2 (11.2), refer to the *Oracle Warehouse Builder Release Notes.*

3.19 Oracle XML DB

The following features are not supported with Oracle XML DB:

- Flashback Archive

- Editioning Views

- SecureFiles LOB Encryption

- Oracle Label Security (OLS) with a hybrid structured and unstructured XMLIndex on the same XML document.

3.19.1 Change in Semantics of xdb:defaultTable Annotation

There is a change in behavior in the semantics of `xdb:defaultTable` annotation while registering Oracle XML DB schemas in 11.2 as compared to 11.1 (reference Bug 7646934). If you specify `xdb:defaultTable="MY_TAB"` without specifying `xdb:sqlInline="false"`, Oracle XML DB creates the table as requested and implicitly marks it as an out-of-line table. This behavior is different from 11.1 where the `defaultTable` annotation was ignored when the `sqlInline` setting was missing.

3.19.2 Additional SSL_CIPHER_SUITE Settings for HTTPS Configuration

Starting in release 11.2.0.3, configuring HTTPS with Oracle XML DB requires that you first set up `SSL_CIPHER_SUITES` to include `SSL_DH_anon` (reference Bug 8403366). The parameter can be set to any one of the following values:

```
SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
SSL_DH_anon_WITH_RC4_128_MD5
SSL_DH_anon_WITH_DES_CBC_SHA
```

See the section titled "Configuring Secure Sockets Layer Authentication" in *Oracle Database Advanced Security Administrator's Guide* 11*g* Release 2 for more details.

This has to be followed by the normal HTTPS configuration steps for Oracle XML DB documented in the section titled "Accessing the Repository using Protocols" in the *Oracle XML DB Developer's Guide* 11*g* Release 2 .

3.19.3 VARRAY Storage Default Change

In Oracle Database 11*g* Release 1 (11.1), the default value for `xdb:storeVarrayAsTable` changed from `FALSE` to `TRUE` for XMLType object-relational storage. This default applied to the default table, but not when creating XMLType object-relational tables and columns after the schema registration (reference Bug 6858659). In Oracle Database 11*g* Release 2 (11.2), all `VARRAY` data elements are created as tables by default. This provides a significant performance increase at query time. In addition, note the following:

- Tables created prior to 11.2 are not affected by this. The upgrade process retains storage parameters. This only affects tables created in 11.2 or later.

- You can retain the pre-11.2 default of `VARRAY` storage as LOBs if you have small `VARRAY` data elements and you read and or write the full `VARRAY` all at once. You have two options to revert to the pre-11.2 behavior:

  - Re-register the schema with `xdb:storeVarrayAsTable=FALSE`. This affects the default and non-default tables.

  - Or, when creating the table (for non default tables), you can use the `STORE ALL VARRAYS AS LOBS` clause to override the default for all `VARRAY` data elements in the XMLType. This clause can only be used during table creation. It will return an error if used in the `table_props` at schema registration time.

- For schemas registered prior to 11.2 (when the default storage for `VARRAY` data elements was `LOB`), you can use `STORE ALL VARRAYS AS TABLES` clause to override the default for all `VARRAY` data elements in the XMLType.

3.20 Pro*C

The Pro*C readme file is located at:

`ORACLE_HOME/precomp/doc/proc2/readme.doc`

3.21 Pro*COBOL

The Pro*COBOL readme file is located at:

`ORACLE_HOME/precomp/doc/procob2/readme.doc`

3.22 SQL*Plus

The SQL*Plus readme file is located at:

`ORACLE_HOME/sqlplus/doc/README.htm`

3.23 Documentation Addendum

This section contains corrections to the following Oracle Documentation for this release:

- Section 3.23.1, "Oracle Database Heterogeneous Connectivity User's Guide"

3.23.1 Oracle Database Heterogeneous Connectivity User's Guide

The following information should be removed from Chapter 4, "Using Heterogeneous Services Agents", section "Determining the Heterogeneous Services Parameters":

> The Distributed Access Manager has a refresh capability available through the menu and toolbar that allows users to rerun queries if necessary and update the data. When the data is refreshed, the tool verifies that the set of registered agents remains the same. If it is not, then the global view is updated. See *Oracle Enterprise Manager Administrator's Guide* and online help for more information about the Distributed Access Manager.

3.23.2 Oracle Database PL/SQL Packages and Types Reference

The `listno` parameter of the `DBMS_UTILITY` was inadvertently excluded from the *Oracle Database PL/SQL Packages and Types Reference*. For example, the `GET_PARAMETER_VALUE` function should read as follows:

GET_PARAMETER_VALUE Function

This function gets the value of specified initialization parameter.

Syntax

```
DBMS_UTILITY.GET_PARAMETER_VALUE (
   parnam      IN       VARCHAR2,
   intval      IN OUT   BINARY_INTEGER,
   strval      IN OUT   VARCHAR2,
   listno      IN       BINARY_INTEGER DEFAULT 1)
 RETURN BINARY_INTEGER;
```

Parameters

Parameter	Description
parnam	Parameter name.
intval	Value of an integer parameter or the value length of a string parameter.
strval	Value of a string parameter.
listno	List item number. If retrieving parameter values for a parameter that can be specified multiple times to accumulate values, use this parameter to get each individual parameter.

3.23.3 Oracle Grid Infrastructure Installation Guide

In *Oracle Grid Infrastructure Installation Guide*, in the section "Creating Disk Groups for Oracle Database Data Files," the guide indicates that logical volume managers are not supported with Oracle RAC. This is incorrect. The correct support information is as follows:

> Although you can specify a logical volume as a device in an Oracle ASM disk group, Oracle does not recommend their use because it adds a layer of complexity that is unnecessary with Oracle ASM. In addition, Oracle RAC requires a cluster

logical volume manager in case you decide to use a logical volume with Oracle ASM and Oracle RAC.

3.24 Open Bugs

This section lists known bugs for this release. A supplemental list of bugs may be found as part of the release documentation specific for your platform.

3.24.1 Oracle Automatic Storage Management (Oracle ASM) Known Bugs

Bug 12837037

If you try to put the DIAGNOSTIC_DEST initialization parameter on Oracle ACFS by modifying the DIAGNOSTIC_DEST parameter in DBCA's *"All initialization Parameters"* page, DBCA creates the database successfully, but does not add the Oracle ACFS resource dependency for the database.

Workaround: After creating the database successfully, manually add the Oracle ACFS dependency using the following command:

```
srvctl modify database -d db_unique_name -j acfs_path_list
```

Bug 12539000

A synchronization problem in the Interprocess Communication (IPC) state of some Oracle processes causes a fatal error during rolling migration. The following error is seen in the alert log:

```
processes are not on active shared page
```

Workaround: None.

Bug 12332603

Oracle Automatic Storage Management (Oracle ASM) loses the rolling migration state if Cluster Ready Services (CRS) shuts down on all nodes. If this occurs, one of the Oracle ASM versions will fail with either the ORA-15153 or ORA-15163 error message.

Workaround: Consider the following scenario of 4 nodes (node1, node2, node3, and node4) that are at release 11.2.0.2 and being upgraded to release 11.2.0.3:

- node1 and node2 are upgraded to 11.2.0.3 and running.

- node3 and node 4 are still at 11.2.0.2 and running.

- Now consider that there is an outage where all CRS stacks are down which leaves the cluster in a heterogeneous state (that is, two nodes at 11.2.0.2 and two nodes at 11.2.0.3).

To proceed with the upgrade, execute one of the following steps (depending on the node that was started as the first node):

- If node3 or node4 was started as the first node (for example, as an 11.2.0.2 node), you need to execute the ALTER SYSTEM START ROLLING MIGRATION TO '11.2.0.3' command on the Oracle ASM instance on node3 or node4 before you can bring up an 11.2.0.3 node.

- If node1 or node2 was started as the first node, you need to execute the ALTER SYSTEM START ROLLING MIGRATION TO '11.2.0.2' command on the Oracle ASM instance on node1 or node2 before you can bring up any 11.2.0.2 node.

Continue the upgrade procedure as already documented from this point forward. Note that before executing one of the above steps to bring the Oracle ASM cluster back into rolling migration, you cannot start two nodes of different versions in the cluster. If you do so, one of the Oracle ASM versions will fail with either the `ORA-15153` or `ORA-15163` error message.

Bug 9413827

An 11.2.0.1 Oracle Clusterware rolling upgrade to 11.2.0.3 fails when Oracle Cluster Registry (OCR) is on Oracle ASM.

Workaround: Apply the patch for bug 9413827 on 11.2.0.1 Oracle Grid Infrastructure for a cluster home before performing the upgrade.

Bug 9276692

Cannot permanently stop the Oracle ASM instance.

Workaround: If the Oracle ASM instance is disabled using SRVCTL, you must unregister Oracle ACFS-related resources to avoid restarting the Oracle ASM instance. Do this by executing the following command as root:

```
acfsroot disable
```

3.24.2 Oracle ASM Dynamic Volume Manager (Oracle ADVM) Known Bugs

Bug 9683229

Oracle ADVM does not support mounting ext3 file systems over Oracle ADVM with the mount barrier option enabled. The mount barrier option is enabled by default on SLES11 SP1.

Workaround: Mount ext3 file system with `-o barrier=1`. For example:

```
mount -o barrier=0 /dev/asm/myvol-131 /mnt
```

3.24.3 Deinstallation Tool Known Bugs

Bug 12769015

After upgrading from 11.2.0.1 or 11.2.0.2 to 11.2.0.3, deinstallation of the Oracle home in the previous version may result in the deletion of the old Oracle base that was associated with it. This may also result in the deletion of data files, audit files, etc., that are stored under the old Oracle base.

Workaround: Before deinstalling the Oracle home in the previous version, edit the `orabase_cleanup.lst` file found in the `<Oracle Home>/utl` directory and remove the `oradata` and `admin` entries. Then, deinstall the Oracle home using the 11.2.0.3 deinstallation tool.

Bug 12762927

When using the deinstallation tool to deinstall a shared Oracle RAC home, some of the files or directories may not get deleted.

Workaround: To remove the `ORACLE_HOME`, run the `rm -rf $ORACLE_HOME` command after the deinstallation tool exits.

Bug 9925724

If `Grid_home` is created directly under a root-owned directory, the deinstallation tool cannot remove the top-level home directory. An empty Oracle home directory remains at the end of the deinstallation.

Workaround: Run `rmdir` `ORACLE_HOME` using the root user on all nodes.

Bug 9406768

The 11.2 deinstallation utility is removing all the homes under Oracle base if these homes are not using the same central inventory and the deinstallation utility finds this home is the only one registered in inventory.

Workaround: While installing 11.2 products:

- Oracle does not recommend using multiple central inventories. Avoid this if possible.

- If for some reason a different central inventory is required, use a different Oracle base directory for each central inventory.

Bug 8666509

A deinstallation of Oracle Clusterware should ask you to detach any pre-11.2 Oracle RAC homes from the Oracle inventory.

Workaround: After you deconfigure and deinstall an upgraded 11.2 Oracle RAC home and want to continue with deconfiguration and deinstallation of the Oracle Grid Infrastructure for a cluster home, first detach any pre-11.2 Oracle RAC software homes from the central Inventory.

Bug 8644344

When running the deinstallation tool to deinstall the database, you will be prompted to expand the Oracle home and to select a component. If you select the top level component, `Oracle Database Server`, and do not select the Oracle home, OUI does not show the message to run the deinstall utility and proceeds with the deinstallation of the database.

Workaround: Run the deinstallation tool to deinstall the Oracle home.

Bug 8635356

If you are running the deinstall tool from `ORACLE_HOME` that is installed on shared NFS storage, then you will see errors related to `.nfs` files during `ORACLE_HOME` clean up.

Workaround: To remove the `ORACLE_HOME`, run the `rm -rf` `ORACLE_HOME` command after the deinstall tool exits. Alternatively, you can use the standalone `deinstall.zip` and specify the location of the `ORACLE_HOME`.

3.24.4 Oracle ACFS Known Bugs

Bug 12842804

After executing the following sequence of commands, file system security will be left in an incomplete state:

```
acfsutil sec prepare -m <mount-path> -u
acfsutil snap create -w <snap-name> <mount-path>
```

```
acfsutil sec prepare -m <mount-path>
```

Specifically, the problem is executing the `sec prepare` command after `sec prepare -u` and `snap create` commands.

Workaround: Contact Oracle Support Services to help restore the security status of the file system.

Bug 12827493

During execution of the `acfsutil sec save` command, the XML file that is generated has a timestamp formatted in accordance to NLS language settings. For instance, a Japanese environment will have a timestamp in Japanese characters in the format `Dy DD-MON-YYYY HH24:MI:SS`, where `Dy` is the day of the week (for example, `Mon`, `Tue`, `Wed`, and so on) and `MON` is the month of the year (for example, `JAN`, `FEB`, and so on).

The XSD validation fails because unknown characters (apart from English characters) are observed in the XML file. Consequently, the `acfsutil sec load` command fails. The problem is also seen with the automatically generated XML file `secbackup.xml` during `acfsutil sec` commands.

Workaround: Change the date to English in the XML file. Contact Oracle Support Services to fix the `secbackup.xml` file before performing a `acfsutil sec load` command.

Bug 12704678

When upgrading from release 11.2.0.2 to 11.2.0.3, entries in the Oracle ACFS registry do not automatically have the proper dependencies set for the Oracle ACFS registry resource.

Workaround: To ensure that proper dependencies are set, delete each Oracle ACFS registry entry and reenter it using `acfsutil registry`. This can be done while the file systems are mounted.

Bug 12690672

In releases prior to 11.2.0.3.0, it is possible to put the database home on an Oracle Automatic Storage Management Cluster File System (Oracle ACFS). If the database home is on an Oracle ACFS file system, the database will have a hard start and stop dependency on the corresponding Oracle ACFS file system.

After upgrading Oracle Grid Infrastructure or Oracle RAC to release 11.2.0.3.0, the dependency between the database and the Oracle ACFS file system, which stored the previous version's database home, is not deleted.

If using a different Oracle ACFS file system than was used to store the previous version's database home, the database fails to start.

Workaround: After a database upgrade, if using a different Oracle ACFS file system for the database home, Oracle recommends that you review the list of Oracle ACFS file systems that you are using for the database, and update the database dependencies on the file systems using the `srvctl modify database -d` *db_unique_name* `-j` *acfs_path_list* command (instead of the `srvctl modify` *filesystem* `-j` *filesystem-list* command).

Bug 10069735

In a cluster with a password-protected key store, when an Oracle ACFS file system using encryption is mounted through the Oracle ACFS mount registry, the administrator is not prompted to enter the key store password. Although the process of mounting the file system succeeds, not all information required for Oracle ACFS encryption to work correctly is made available to the file system. In this case, encryption is not operational on this file system and any encrypted files in the file system are not available for read or write.

Workaround: In a cluster with a password-protected key store, do not use the Oracle ACFS mount registry for mounting any file systems that are using encryption. If some file systems are already mounted through the Oracle ACFS mount registry, unmount them and remove any such file systems from the mount registry to avoid possible unavailability of encrypted data in the future. Then, remount these file systems without using the Oracle ACFS mount registry, providing the correct password when requested.

Bug 8644639

When creating an Oracle ACFS mount point and adding it to the registry, the mount point is not mounted automatically if the following conditions are met:

1. The mount point directory was previously registered with the Oracle ACFS Registry.

2. The mount point directory had been previously mounted.

3. The mount point had then been unmounted and removed from the Oracle ACFS Registry.

4. The `ora.registry.acfs` resource has not been restarted since the mount point was deleted from the registry.

Workaround: Remove the mount point directory from the file `/tmp/.usm_state_file`.

3.24.5 Oracle Clusterware Known Bugs

Bug 13110641

While installing Oracle RAC software on a cluster configured with Grid Naming Service (GNS), the Prerequisite's page might show a warning status for the GNS Integrity check even when GNS is working fine.

The message appears to be of the following type:

```
PRVF-5217 : An error occurred while trying to look up IP address for
"<gns-subdomain-extended-name>"
```

Workaround: Execute `nslookup` on the fully qualified names that are listed in the error message. If `nslookup` returns an IP address for the name with a non-authoritative answer, then this warning can be ignored. If the name does not resolve to an IP address, then follow the steps mentioned in Action part of the error message.

Bug 12900070

If you are preparing to upgrade Oracle Clusterware, and you use the Cluster Verification Utility (CVU) command `runcluvfy.sh stage -pre crsinst -upgrade`, then you may encounter the following error:

```
Unable to retrieve nodelist from Oracle Clusterware
```

The cause of this error is that `olsnodes` cannot return a list of nodes when Oracle Clusterware is down.

Workaround: Run the `cluvfy.sh stage crsinst -upgrade` command using the `-n` flag, and provide a comma-delimited list of cluster member nodes. For example:

```
runcluvfy.sh stage -pre crsinst -upgrade -n node1, node2, node3
```

Bug 12769576

In release 11.2.0.3, the default `RETENTION_TIME` size of a Cluster Health Monitor (CHM) repository in number of seconds is 30823 for a 4-node cluster or is (30823*4) divided by the number of nodes for other clusters. When upgrading from 11.2.0.2 to 11.2.0.3, the `RETENTION_TIME` is 6311 for a 4-node cluster.

Workaround: Oracle recommends changing the `RETENTION_TIME` size from 6311 to 30823 for a 4-node cluster after upgrading from 11.2.0.2 to 11.2.0.3 by using the following `oclumon` command:

```
oclumon manage -repos resize 30823
```

Bug 12710059

When upgrading Oracle Grid Infrastructure to 11.2.0.3 with Oracle Cluster Registry (OCR) and voting disk files on Network File Storage (NFS), if the network card (for example, eth0) that the cluster nodes use to connect to the NFS server goes down, then the `rootupgrade.sh` scriptwill fail to upgrade the Oracle stack.

Workaround: Restore the network interface and make sure that the old Oracle clusterware stack is up and actively running prior to running the `rootupgrade.sh` script.

Bug 12380213

An error exception occurs when installing an 11.2.0.2 database with data files on 11.2.0.3 Oracle Restart.

Workaround: To install release 11.2.0.2 database against 11.2.0.3 Oracle Restart, you need to invoke 11.2.0.2 `runInstaller` with `-ignorePrereq` and then complete the 11.2.0.2 database installation.

Bug 9948549

If the agent terminates while starting a database or Oracle ASM instance, it is possible that the instance startup will not complete.

Workaround: Stop and restart the instance using `srvctl` or `sqlplus`.

Bug 8733944

Due to a problem in Oracle Clusterware starting with release 11.1.0.7, with the patches required for Exadata support or 11.1.0.7 CRS bundle Patch 1, in some cases the CSS daemon may fail when the clusterware is brought down on another node, either due to a shutdown command or a failure.

The symptom is an `ASSERT` in the CSSD log indicating that a maximum value has been exceeded. For example:

```
Group ID of xxxx exceeds max value for global groups
```

Workaround: Oracle recommends that customers running with the Exadata support patches or 11.1.0.7 CRS Bundle Patch 1 apply the patch for this bug to avoid this problem.

This problem may also be seen during an upgrade from 11.1.0.7 with patches as indicated above. To eliminate the potential of an 11.1.0.7 node failing during upgrade, the patch for this bug may be applied to the 11.1.0.7 nodes prior to upgrade.

When upgrading, it is recommended that the upgrade be completed on all nodes without restarting any non-upgraded nodes during the upgrade. If an 11.1.0.7 node does fail while performing the upgrade, it should be upgraded as opposed to restarted.

Bug 8657184

If two network interfaces are configured as public network interfaces in the cluster, the failure of one public interface on a node does not result in automatic VIP failover to the other public interface.

Workaround: If multiple public networks interfaces are present, then use interface bonding for high availability. At the Oracle Clusterware installer "Specify Network Interface Usage" screen, choose only one (bonded) interface as public. When configuring public networks with `srvctl add nodeapps` or `srvctl add vip`, specify only a single network interface name in the `-A` or `-S` argument.

Bug 8288940

Creating pre-11.2 Oracle RAC database in 11.2 Oracle Grid Infrastructure for a cluster environment using DBCA may fail with following messages. When using a cluster file system as storage, you see the following message:

```
ORA-00119: invalid specification for system parameter REMOTE_LISTENER
```

When using Oracle ASM as storage, you see the following message:

```
DBCA could not startup the ASM instance configured on this node
```

Workaround: Apply the patch for this bug in pre-11.2 database home. This patch is needed for 10.2.0.4, 11.1.0.6 and 11.1.0.7 database releases. No patch is needed for release 10.2.0.5.

Bug 8262786, 8373758, 8441769, 8406545

Oracle resources for 10.2 and 11.1 Oracle RAC databases may not operate properly after upgrading Oracle Clusterware to 11.2.

Workaround: Apply the patches for Bugs 8262786, 8373758, 8441769, and 8406545 to the Oracle Database home.

Bug 3841387, 8373758

Oracle resources for 10.1 Oracle RAC Databases may not operate properly after upgrading Oracle Clusterware to 11.2.

Workaround: Apply the patches for Bug 3841387 and Bug 8373758 to the Oracle Database home.

3.24.6 Oracle Data Pump Known Bugs

Bug 12866600

Database links imported from an 11.2.0.3 database into a version prior to 11.2.0.3 (including 11.2.0.2) will not be usable in the import database. Any attempt to use a database link will cause the following ORA-600 error:

```
ORA-00600 [kzdlk_zt2 err], [18446744073709551601]
```

Workaround: The database links that are imported need to be re-created before they can be used.

3.24.7 Oracle Database Enterprise Edition Known Bugs

Bug 9951827

Transportable tablespace import does not handle timestamp with timezone version change.

If a transportable dumpfile produced in release 11.2.0.3 contains tables with timestamp with timezone columns and the version of the timezone table on the target database is different than that from the source database, the import is be prevented from running.

If a dumpfile produced prior to release 11.2.0.3 had a different timezone table version than that of the target, then the import is be prevented from running.

Workaround: Make sure the timezone tables for the import and export databases are the same.

Bug 9859532

Current implementation of node-specific network interfaces requires complete definition of all networks used by Oracle RAC for that node (that is, either node abides by global network configuration or it defines its own node-specific network configuration).

As a corollary, once the first node-specific network interface is defined for a given node, Oracle RAC will not consider any configured global network interfaces that are already configured and may have applied for the same node.

While this is correct, it presents a problem. If the cluster had a working global network configuration, the moment a user updates it (using `oifcfg`) to define a node-specific `public` interface, a global configuration will not be considered for this node, and it will have only one newly-defined `public` interface. Any cluster interconnects that existed in the global network configuration, and may still resolve fine for this node, will not be considered valid. Thus, the node loses cluster interconnects and PCW stack goes down on that node.

Workaround: If the node belongs to a global cluster network configuration, then if there is an intent to make network configuration node-specific, the first defined node-specific interface must be cluster interconnect, so that node never loses interconnect with other cluster nodes. Then, other node-specific interfaces can be defined as necessary.

Bug 9181675

The `asmgidwrap` script needs to be called if you are creating a database manually on Oracle ASM to avoid a permission error.

Workaround: For a role-separated installation (that is, there is a different user and group for grid and RDBMS), use DBCA to create the database that automatically calls `asmgidwrap` script while creating a database on Oracle ASM. If you choose to create a database manually, the script needs to be called explicitly so the proper group can be set to avoid a permission error.

Bug 8729627

When using 11.1 DBCA to remove a database on a cluster running 11.2 Oracle Clusterware, a `PRKP-1061/CRS-2524` error may be displayed because the database resource is locked.

Workaround: You can ignore the message. Click **OK** to continue.

Bug 8679182

When configuring a database on a cluster that has multiple public subnets defined for its VIPs (for example, using a command similar to `srvctl add vip -k 2 -A ...`), the database agent automatically sets `LOCAL_LISTENER` to the listener on the default network. This may duplicate a listener set in `LISTENER_NETWORKS`.

Workaround: Do not specify listeners in `LISTENER_NETWORKS` that are on the default public subnet.

3.24.8 Oracle Enterprise Manager Grid Control Known Bugs

Bug 12793336

When attempting to upgrade Cluster Ready Services (CRS) or Oracle ASM to release 11.2 using the Oracle ASM Configuration Assistant (ASMCA), the upgrade succeeds, but it may fail to update the new clusterware home for cluster targets in existing agent homes due to permission issues. As a result, Oracle Enterprise Manager Grid Control and Database Control cannot monitor the Oracle ASM and CRS targets.

Workaround: Modify `ORACLE_HOME` property of Oracle ASM and cluster targets and reload the agents.

Bug 12725253

Database upgrade from release 10.2.0.1 to 11.2.0.3 could run into error `ORA-01722` `(Invalid number)` when running in locales that do not use the period (.) as the decimal separator.

Workaround: Run the upgrade in a locale that uses the period (.) as the decimal separator. After the upgrade, switch to your preferred locale.

3.24.9 Oracle Database QoS Management Known Bugs

Bug 12767103

If a user creates a performance class with two or more services in its classifier(s) and these services are not all specified to run in the same server pool, the metrics graphs for that performance class on the Enterprise Manager Quality of Service (QoS) Management Performance Class details page are incorrect. The Resource Use Time and Resource Wait Time graphs will only display metrics from one server pool. The other graphs will correctly display metrics for all server pools.

This bug will not affect the correct management or recommended actions associated with this type of performance class.

Bug 12792222

This bug applies to recommendations for CPU resources managed by Oracle Database QoS Management. If the number of configured CPUs for all instances on a server is less than the number of physical CPUs for that server, then the nonallocated, or "free", CPUs are not detected by Oracle Database QoS Management and no recommendation is made to increase the number of configured CPUs. Only those "slices" that host databases are considered as donors for the target slice. Adding one of the non-allocated CPUs should be the first-ranked Move CPU action.

Workaround: Make sure the sum of CPU counts configured for each database instance on each server is the same as the number of physical CPUs.

Bug 10019209

This bug applies to platforms that support the Cluster Health Monitor (CHM). If an Oracle Clusterware-managed database service is in a stopped but not disabled state, it will be started by Oracle Database QoS Management if the server hosting that service is not detected to be in a memory overcommitted state. If memory is overcommitted, then all enabled services will be stopped even if they were manually started. The desired behavior is to only start services on the transition from a memory overcommitted state (red) to a normal state (green). If a service is manually started when the server is in the red state, that service should not be shut down.

Workaround: Stop and disable services that you want to remain in the stopped state or disable QoS Management from the Oracle Enterprise Manager Console.

3.24.10 Oracle Database Vault Known Bugs

Bug 7033772

Database Vault Administrator (DVA) does not work after an Enterprise Manager DBControl upgrade.

Workaround: Manually redeploy DVA after DBControl has been upgraded. You can follow the steps described in Appendix C, Section "Deploying Database Vault Administrator to the Database Console OC4J Container" of the *Oracle Database Vault Administrator's Guide*.

Bug 8686258

Database Vault policy cannot be managed in Oracle Enterprise Manager Database Control because the following message is displayed in Database Vault Administration page:

```
"OPERATOR TARGET" privilege does not exist.
"You must have OPERATOR TARGET privilege to perform this operation."
```

Workaround: To manage Database Vault policy using Oracle Enterprise Manager, the Database Vault administrator must have the EM Administrator privilege. If you do not want to grant the EM Administrator privilege to the Database Vault administrator, then use the Database Vault Administrator page directly. For additional information, see *Oracle Database Vault Administrator's Guide*.

Bug 8341283

The `ACTION_NAME` entry in the `DVSYS.AUDIT_TRAIL$` table displays `Realm Authorization Audit` for a failed realm enforcement if the audit option is set to audit on success and failure. The `RETURNCODE` will show the correct error code that was triggered.

Workaround: Use the `RETURNCODE` value to determine whether a violation has occurred and the `ACTION_NAME` column to identify whether the audit was generated by a realm enforcement or command rule enforcement.

Bug 9766628

`emctl` commands did not return valid results as expected.

Workaround: The `emctl` command needs to be run from an Oracle Database home. Do not invoke this command from the Oracle Grid Infrastructure for a cluster home.

3.24.11 Oracle Enterprise Manager Database Control Known Bugs

Bug 8674920

If the installation owners for the Oracle Grid Infrastructure for a cluster and Oracle Database are different, then the owners of Oracle ASM binaries and Oracle Enterprise Manager Agent binaries are also different. When you start Support Workbench, the error message `Error Operation failed - Operation failed` might appear, because the Oracle Enterprise Manager Agent is running as a different user, and Support Workbench does not have permissions for the Oracle ASM target.

Workaround: None.

Bug 8350597

If Database Control is running in an IPv6 environment, then you cannot use it to monitor Exadata cells and you should not add Exadata cells as targets.

Workaround: None.

3.24.12 Oracle OLAP Known Bugs

Bug 9917299

If the database is installed using the seed provided in the installation kit, and the OLAP option is not selected, then either at the end of the installation or some time later, the OLAP Analytic Workspace and OLAP API components will be reported as invalid.

This will not affect the running of the instance in any way, other than the error messages.

Workaround: Do one of the following as a workaround:

- Ignore the error.
- Enable OLAP (or the offending option).
- Create and use your own seed database that does not include OLAP.

Bug 9545221

Importing an materialized view-enabled cube or cube dimension whose source table is not part of the target schema fails with an `Object not found` error.

Workaround: Disable materialized views for the failing object prior to the import, then reenable them when the source tables are present.

3.24.13 Oracle Real Application Clusters Known Bugs

Bug 10369108

After converting from an administrator-managed database to a policy-managed database, you may need to update the database password file.

Workaround: To update the database password file, execute the following steps:

1. Copy the existing password file `orapw$ORACLE_SID` to `orapw<db_unique_name>` on the node where the administrator-managed database was running.

2. Copy this file, `orapw<db_unique_name>`, to the same location on every cluster node.

3.24.14 Oracle SQL*Loader Known Bugs

Bug 9301862

When external table code reads very large files on disks served by NFS, the I/O performance of the read can slow down over time. This is caused by NFS caching blocks from the file in memory as it is read. Since these blocks are not re-read, the time spent maintaining the cache slows down the I/O operations.

Workaround: The current behavior (of not using the `O_DIRECT` flags) remains the default. You can enable the use of the `O_DIRECT` flag in the following ways:

- Enable fix control for this bug and set it to `ON` with the following command:

  ```
  ALTER SESSION SET "_fix_control"='9301862:ON';
  ```

  When fix control is enabled, the external table code looks at the `FILESYSTEMIO_OPTIONS` configuration parameter and if it is set to either `DIRECTIO` or `SETALL`, then the `ORACLE_LOADER` access driver will specify the `O_DIRECT` flag when opening data files for reading. If the `FILESYSTEMIO_OPTIONS` parameter is not set or if it is set to other values, then the access driver will not attempt to use `O_DIRECT` unless you choose the following option.

- Use the new `IO_OPTIONS` clause in the access driver to specify direct I/O. The clause is part of the larger `RECORDS` clause. The syntax is:

  ```
  IO_OPTIONS (DIRECTIO | NODIRECTIO)
  ```

  If `DIRECTIO` is specified, then the access driver uses `O_DIRECT` flag when opening the file. If `NODIRECTIO` is specified, then the access driver does not use the `O_DIRECT` flag. Note that the action specified by `IO_OPTIONS` is performed regardless of the setting of `_fix_control` for this bug.

  Note that the first option is a way to enable the use of `O_DIRECT` for all external tables while the second option allows `DIRECTIO` to be used or not used for specific external tables.

3.24.15 Oracle Universal Installer Known Bugs

You should also review Section 3.1, "Compatibility, Upgrading, Downgrading, and Installation" for other issues related to installation and upgrades.

Bug 12930328

If the central inventory location is different on different nodes of a cluster, `addnode.sh` does not update the inventory correctly on remote nodes of the cluster.

Workaround: Adding nodes to a cluster requires the central inventory location to be the same on all the nodes of the cluster. Please ensure that this is the case prior to running `addnode.sh`.

Bug 12887720

A file system supporting Access Control Lists (ACL) should not be used for staging Oracle software as this may cause an error when copying the files from the staging area to a temporary directory which does not support Access Control Lists. ACL permissions cannot be preserved while copying the files and this causes the copy to fail.

Workaround: None.

Bug 12719356

After installing Oracle RAC, you might see the following error message in the installation log files:

```
OiiolLogger.addFileHandler:Error while adding file handler -
Logs/remoteInterfaces<time>.log
```

Workaround: None. This error message can be ignored.

Bug 12711224

If Oracle Universal Installer (OUI) crashes during a node reboot or crashes while you are executing the rootupgrade script, OUI cannot resume post-upgrade tasks.

Workaround: You have to manually take care of the following tasks and complete the upgrade:

- If you are upgrading from a pre-11.2 to 11.2.0.3 release:

  1. Update inventory

  2. Orace Net Configuration Assistant

  3. Automatic Storage Management Configuration Assitant

  4. Enterprise Manager Configuration Upgrade Utility

  5. Oracle Cluster Verification Utility

- If you are upgrading from a post-11.2 to 11.2.0.3 release:

  1. Update inventory

  2. Enterprise Manager Configuration Upgrade Utility

  3. Oracle Cluster Verification Utility

Bug 8729326

When upgrading to 11.2 Clusterware, the Installer invokes ASMCA in silent mode to upgrade Oracle ASM into Oracle Grid Infrastructure for a cluster home. Oracle ASM upgrade is handled in rolling fashion when upgrading from 11.1.0.7. Prior versions of Oracle ASM instances are upgraded in non-rolling fashion and Oracle ASM-based databases are bounced without any prior warning.

Workaround: You can plan your database outage to be the point where you acknowledge the Installer prompt after executing `root.sh` on all nodes. At this point, CRS is upgraded in rolling fashion and the Installer will be calling ASMCA to upgrade Oracle ASM, which will bounce databases as part of Oracle ASM upgrade.

Bug 8666656

The Oracle Universal Installer (OUI) `runInstaller` script that resides in the Oracle home (*ORACLE_HOME*/oui/bin/runInstaller) cannot be used to install the 11.2.0.1 releases of Oracle Database, Oracle Grid Infrastructure for a cluster, and Oracle Database Client.

Workaround: Use Oracle Universal Installer on the respective 11.2.0.1.0 product media to install each product.

Bug 8638708

If you select the database configuration Desktop Class in Oracle Universal Installer (OUI), listener and database control are configured with 'localhost' as the host name. The Oracle Enterprise Manager Database Control `start` and `stop` operations using `emctl` may fail.

Workaround: For Database Control start and stop operations that use `emctl` in that home, set the `ORACLE_HOSTNAME` environment variable to 'localhost'.

Bug 8407818

After adding a new node to a shared Oracle database home using `addNode.sh`, the `/etc/oratab` on the newly added node gets an entry of the source database name that exists on the source node from where `addNode.sh` was run. The `/etc/oratab` file on the new node is supposed to get the database entry after the database instance is added for the new node using DBCA.

Workaround: Before invoking DBCA from the source node to add a new database instance for the new node, open the `/etc/oratab` file on the new node using an editor and remove the entry made for the source database name.

3.24.16 Oracle Wallet Manager Known Bugs

Bug 9880435

Oracle Wallet Manager fails to upload wallet to Directory service when the wallet password and the directory user password are different.

Workaround: Use the same password for the wallet and the directory user.

3.24.17 Oracle XML Database Known Bugs

Bug 12836913

When using `oracle.xdb.XMLType` proprietary constructors in a Java stored procedure with JDK6, the error `Invalid version of the XMLType` could be returned.

Workaround: Either do not use Java stored procedures. Or, use JDK5 instead of JDK6.

Bug 12834970

Starting with release 11.2.0.3, the `MOVEXDB_TABLESPACE` and `REBUILDHIERARCHICALINDEX` procedures were moved from the `DBMS_XDB` package to the `DBMS_XDB_ADMIN` package. These procedures are no longer available in the `DBMS_XDB` package anymore.

Workaround: None.

Bug 9586264

In order to fully optimize some `XMLQUERY` or `XMLTABLE` queries, `OPTIMIZER_FEATURE_ENABLE` should be set to 11.1.0.6 or above.

Workaround: None.

3.24.18 Vendor and Operating System Known Bugs

Bug 8256753

A connect using `SCAN` and `EZCONNECT` on one client machine can be requested to use a specific `SCAN` listener. Therefore, load balancing by round-robin DNS is not possible.

Workaround: Connect to a database using the following configuration specifying `LOAD_BALANCE=on` in `tnsnames.ora`:

```
ORCL =
   (DESCRIPTION =
     (LOAD_BALANCE=on)
     (ADDRESS = (PROTOCOL = TCP)(HOST = stscan1)(PORT = 1521))
     (CONNECT_DATA =
       (SERVER = DEDICATED)
       (SERVICE_NAME = srv.world)
     )
   )
```

4 Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

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