

A Oracle Setup for the Model Repository

This appendix explains how to install, configure, and maintain an Oracle database to support the SA Model Repository.

Supported Oracle Versions and Operating Systems

Support for the Model Repository is limited to certain versions of Oracle running on certain versions of operating systems. HP strongly recommends that you also apply the latest Oracle CPU or PSU patches.

See the *SA 10.0 Support and Compatibility Matrix* for a list of supported Oracle versions and operating systems.

System Requirements

The following sections list the system requirements for Oracle 11g. The SA Installer performs an automated check to ensure that these requirements are met on the Oracle host.

The system requirements and configurations listed in this section apply both to the SA-supplied Oracle RDBMS software as well as to non-SA-supplied Oracle RDBMS and software installations.

-
- ▶ If you create the database using the Oracle Universal Installer rather than the SA Installer, you must check for these packages and patches manually.
-
- ▶ The Oracle database must be installed either on its own host or on a server that has the SA Infrastructure Component bundle installed.
-

Database Server Time Requirements

Core Servers (either Single Core or Multimaster) and Satellite Core Servers must meet the following requirements. These time requirements do not apply to Managed Servers.

- All SA Core Servers must have their time zone set to Coordinated Universal Time (UTC).
- All SA Core Servers must maintain synchronized system clocks. Typically, you will synchronize the system clocks through an external server that uses NTP (Network Time Protocol) services.

Linux Time Configuration

To configure the time zone on a Linux server, perform the following tasks:

- 1 Copy or link
`/usr/share/zoneinfo/UTC` to `/etc/localtime`.
- 2 Ensure that the `/etc/sysconfig/clock` file contains the following lines:
`ZONE="UTC"`
`UTC=true`

Hostname Setup

- 1 You must be able to ping the database server hostname. To verify this, enter the following command:
`# ping <hostname>`
- 2 Check that the database server name is FQDN by using the following command:
Red Hat Linux
`# hostname`
SUSE Linux
`# hostname -f`

If the hostname is not configured correctly, Oracle will not start and you will encounter the following error:

```
ORA-00600: internal error code, arguments: [keltnfy-ldmInit], [46], [1], [],  
[], [], [], []
```

Hardware Requirements

The server that will host the Oracle database for the Model Repository must meet the hardware requirements listed in this section.

Linux Requirements

The following are hardware requirements for running Oracle 11g under Linux.

- For detailed Linux requirements, see the *Oracle® Database Quick Installation Guide 11g Release 2 (11.2) for Linux x86-64* (Part Number E24326-02) available at:
<http://docs.oracle.com>
- Determine the processor type to verify that the processor's architecture matches the Oracle software release you will install. Use the following command to check system architecture:
`# uname -m`
- The recommended physical memory is 16 GB or more of RAM. An HP-supplied Oracle installation will use a minimum of 2 GB memory. The Oracle SGA memory can be increased after database installation. You can use the following command to check memory status:
`# grep MemTotal /proc/meminfo`

- Required available swap space is shown in [Table 46](#):

Table 46 Required Available RAM Swap Space

RAM Between:	Available Swap Space
4 GB and 16 GB	Equal to the size of RAM
More than 16 GB	16 GB

You can use the following command to check swap space:

```
# grep SwapTotal /proc/meminfo
```

- As of Oracle 11g, Automatic Memory Management (ASMM) requires more shared memory (/dev/shm) and file descriptors. Shared memory should be sized to be at least the greater of MEMORY_MAX_TARGET and MEMORY_TARGET for each Oracle instance on a database server.

You can use the following command to check available shared memory:

```
# df -h /dev/shm/
```

- Free tmp space should be 1GB or more of /tmp directory space

You can use the following command to check tmp space:

```
# df -h /tmp
```

Oracle Solaris, HP-UX and IBM Requirements

Refer to the *Checking the Hardware Requirements* section in the following Oracle documents:

- *Database Quick Installation Guide for Oracle Solaris on SPARC (64 Bit)*
- *Database Quick Installation Guide for HP-UX Itanium*
- *Database Quick Installation Guide for IBM AIX on POWER Systems (64-Bit)*

Model Repository (Database) Disk Space Requirements

Additional disk space is required for the Oracle software and the Model Repository data files. Keep in mind that storage requirements for the database grow as the number of managed servers and database activity grows.

As a benchmark figure, you should allow an additional 3.5 GB of database storage for every 1,000 servers in the facility that SA manages. When sizing the tablespaces, follow the general guidelines described in [Table 47](#). If you need to determine a more precise tablespace sizing, contact your technical support representative.

Table 47 Tablespace Sizes

Tablespace	MB/1000	Recommended Minimum Tablespace Size
AAA_DATA	256 MB	2000 MB
AAA_INDX	256 MB	2000 MB
AUDIT_DATA	256 MB	2000 MB
AUDIT_INDX	256 MB	2000 MB
LCREP_DATA	3000 MB	8000 MB

Table 47 Tablespace Sizes (cont'd)

Tablespace	MB/1000	Recommended Minimum Tablespace Size
LCREP_INDX	2000 MB	8000 MB
TRUTH_DATA	1500 MB	4000 MB
TRUTH_INDX	500 MB	4000 MB
STRG_DATA	1300 MB	2000 MB
STRG_INDX	400 MB	2000 MB

Software Requirements

This section lists the requirements for running Oracle 11g under Red Hat Enterprise Linux, Oracle Enterprise Linux and SUSE Linux Enterprise Server.

Linux Requirements

The following are software requirements for running Oracle 11g under Red Hat Enterprise Linux, Oracle Enterprise Linux and SUSE Linux Enterprise Server:

- Required operating system version:
 - Oracle Linux 5 Update 2 (with Red Hat Compatible Kernel)
 - Oracle Linux 6 (with Red Hat Compatible Kernel)
 - Red Hat Enterprise Linux 5 Update 2
 - Red Hat Enterprise Linux 6
 - SUSE Linux Enterprise Server 10 SP2
 - SUSE Linux Enterprise Server 11

You can use the following command to determine the distribution and version of Linux installed:

```
# cat /proc/version
```

- Required Kernel version for Oracle Database 11g Release 2 (11.2):
 - Oracle Linux 5 Update 2
2.6.18 or later (with Red Hat Compatible Kernel)
 - Oracle Linux 6
2.6.32-71.el6.x86_64 or later (with Red Hat Compatible Kernel)
 - Red Hat Enterprise Linux 6
2.6.32-71.el6.x86_64 or later
 - Red Hat Enterprise Linux 5 Update 2
2.6.18 or later
 - SUSE Linux Enterprise Server 10
2.6.16.21 or later

— SUSE Linux Enterprise Server 11

2.6.27.19 or later

You can use the following command to check the kernel versions:

```
# uname -r
```

You can use the following command to check the platform:

```
# uname -mi
```

You can use the following command to check the processor type:

```
# grep "model name" /proc/cpuinfo
```

Linux Package Requirements

Red Hat Enterprise Linux 5 and Oracle Enterprise Linux 5



Starting with Oracle Database 11g Release 2 (11.2.0.2), all the *32-bit packages*, excepting `gcc-32bit-4.3`, listed in the following table are no longer required for installing a database on Linux x86-64. Only the 64-bit packages are required. However, for any Oracle Database 11g release before 11.2.0.2, both the 32-bit *and* 64-bit packages listed in the following table are required.

The following or later package versions for Red Hat Enterprise Linux 5 and Oracle Enterprise Linux 5 (with Red Hat compatible kernel) must be installed (shaded rows indicate 32-bit packages):

Table 48 Required Packages for Red Hat Enterprise Linux 5 and Oracle Enterprise Linux 5

Required Packages	Version
binutils	2.17.50.0.6
compat-libstdc++	33-3.2.3
compat-libstdc++	33-3.2.3 (32-bit)
elfutils-libelf	0.125
elfutils-libelf-devel	0.125
gcc	4.1.2
gcc-c++	4.1.2
glibc	2.5-24
glibc	2.5-24 (32-bit)
glibc-common	2.5
glibc-devel	2.5
glibc-devel	2.5 (32-bit)
glibc-headers	2.5
ksh	20060214

Table 48 Required Packages for Red Hat Enterprise Linux 5 and Oracle Enterprise Linux 5 (cont'd)

Required Packages	Version
libaio	0.3.106
libaio	0.3.106 (32-bit)
libaio-devel	0.3.106
libaio-devel	0.3.106 (32-bit)
libgcc	4.1.2
libgcc	4.1.2 (32-bit)
libstdc++	4.1.2
libstdc++	4.1.2 (32-bit)
libstdc++-devel	4.1.2
make	3.81
sysstat	7.0.2

Red Hat Enterprise Linux 6 and Oracle Enterprise Linux 6

The following or later package versions for Red Hat Enterprise Linux 6 and Oracle Enterprise Linux 6 (with Red Hat compatible kernel) must be installed:

Table 49 Required Packages for Red Hat Enterprise Linux 6 and Oracle Enterprise Linux 6

Required Packages	Version
binutils	2.20.51.0.2-5.11.el6 (x86_64)
compat-libcap1	1.10-1 (x86_64)
compat-libstdc++	33-3.2.3-69.el6 (x86_64)
compat-libstdc++	33-3.2.3-69.el6.i686
gcc	4.4.4-13.el6 (x86_64)
gcc-c++	4.4.4-13.el6 (x86_64)
glibc	2.12-1.7.el6 (x86_64)
glibc	2.12-1.7.el6 (i686)
glibc-devel	2.12-1.7.el6 (x86_64)
glibc-devel	2.12-1.7.el6.i686
ksh	NA
libaio	0.3.107-10.el6 (x86_64)
libaio	0.3.107-10.el6.i686
libaio-devel	0.3.107-10.el6 (x86_64)

Table 49 Required Packages for Red Hat Enterprise Linux 6 and Oracle Enterprise Linux 6 (cont'd)

Required Packages	Version
libaio-devel	0.3.107-10.el6.i686
libgcc	4.4.4-13.el6 (x86_64)
libgcc	4.4.4-13.el6 (i686)
libstdc++	4.4.4-13.el6 (x86_64)
libstdc++	4.4.4-13.el6.i686
libstdc++-devel	4.4.4-13.el6 (x86_64)
libstdc++-devel	4.4.4-13.el6.i686
make	3.81-19.el6
sysstat	9.0.4-11.el6 (x86_64)

SUSE Linux Enterprise Server 10

The following or later package versions for SUSE Linux Enterprise Server 10 must be installed:

Table 50 Required Packages for SUSE Linux Enterprise Server 10

Required Packages	Version
binutils	2.16.91.0.5
compat-libstdc++	5.0.7
gcc	4.1.0
gcc-c++	4.1.2
glibc	2.4-31.63
glibc-devel	2.4-31.63
glibc-devel	32bit-2.4-31.63
ksh	93r-12.9
libaio	0.3.104
libaio	32bit-0.3.104
libaio-devel	0.3.104
libaio-devel	32bit-0.3.104
libelf	0.8.5
libgcc	4.1.2
libstdc++	4.1.2
libstdc++-devel	4.1.2

Table 50 Required Packages for SUSE Linux Enterprise Server 10 (cont'd)

Required Packages	Version
make	3.80
numactl	0.9.6.x86_64
sysstat	8.0.4

SUSE Linux Enterprise Server 11

The following or later package versions for SUSE Linux Enterprise Server 11 must be installed:

Table 51 Required Packages for SUSE Linux Enterprise Server 11

Required Packages	Version
binutils	2.19
gcc	4.3
gcc	32bit-4.3
gcc-c++	4.3
glibc	2.9
glibc	32bit-2.9
glibc-devel	2.9
glibc-devel	32bit-2.9
ksh	93t
libaio	0.3.104
libaio	32bit-0.3.104
libaio-devel	0.3.104
libaio-devel	32bit-0.3.104
libgcc43	4.3.3_20081022
libstdc++-devel	4.3
libstdc++33	3.3.3
libstdc++33	32bit-3.3.3
libstdc++43	4.3.3_20081022
libstdc++43	32bit-4.3.3_20081022
libstdc++43-devel	4.3.3_20081022
libstdc++43-devel	32bit-4.3.3_20081022
make	3.81
sysstat	8.1.5

Verifying that Packages are Installed

To verify that RPMs are installed under Linux, enter the following command:

```
# rpm -q --qf '%{NAME}-%{VERSION}-%{RELEASE} (%{ARCH})\n' <rpm_name>
```



See your Oracle documentation to determine requirements for supported databases earlier than Oracle 11.2

Oracle Solaris SPARC, HP-UX and IBM AIX Requirements

For HP-UX and IBM AIX operating system, compiler, patch and any additional software requirements, see the *Checking the Software Requirements* section in the *Oracle® Database Quick Installation Guide* for your operating system.

Non-SA-Supplied Oracle Software and Database Setup



If you plan to install the SA-supplied Oracle RDBMS software and database, you do not need to perform the tasks in this section. The SA Installer performs all the tasks discussed below. For information about installing the SA-supplied Oracle software and database, see [SA-Supplied Oracle RDBMS Software and Database Setup](#) on page 235.

If you plan to use a non-SA-supplied Oracle database with the SA Model Repository, the following steps are required for compatibility with SA. You should also review [System Requirements](#) on page 219 before proceeding with this section.

Modifiable Kernel Parameters

If you manually install the Oracle database, or use an existing database, you must insure that all kernel parameter values are specified correctly for your environment but also within the limitations required by SA.

You can find additional information about kernel parameter configuration in the *Configuring Kernel Parameters* section of the *Oracle® Database Quick Installation Guide*.

Modifiable Kernel Parameter Values for Linux

This section identifies the kernel parameters you can change for supported Linux operating systems.

You can change values for the following parameters in `/etc/sysctl.conf`. If the current value of any parameter is higher than the value listed in this table, then do not change the value of that parameter:

```
# SA Oracle parameters begin
fs.aio-max-nr=1048576
fs.file-max=6815744
kernel.shmmax=2147483648
kernel.shmall=2097152
kernel.shmni=4096
kernel.sem=250 32000 100 128
```

```

net.core.rmem_default=262144
net.core.rmem_max=4194304
net.core.wmem_default=262144
net.core.wmem_max=1048586
net.ipv4.ip_local_port_range=9000 65500
net.ipv4.tcp_wmem=262144 262144 262144
net.ipv4.tcp_rmem=4194304 4194304 4194304
# SA Oracle parameters end

```

You can change values for the following parameters in `/etc/security/limits.conf`:

```

# SA Oracle parameters begin
oracle soft nofile 1024
oracle hard nofile 65536
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft stack 10240
oracle hard stack 32768
# SA Oracle parameters end

```

You can change values for the following parameters in `/etc/pam.d/login`:

```

session required /lib/security/pam_limits.so

```

You can change values for the following parameters in `/etc/fstab`:

```

shmfs /dev/shm tmpfs size=4g 0

```

You can change values for the following parameters in `/etc/selinux/config`:

```

# SA Oracle parameters begin
SELINUX=disabled
# SA Oracle parameters end

```

Modifiable Kernel Parameter Values for SUSE Linux x86_64

This section identifies additional required settings for SUSE Linux x86_64 when running Oracle 11g:

- Enter the following command to cause the system to read the `/etc/sysctl.conf` file when it restarts:

```
# /sbin/chkconfig boot.sysctl on
```
- You must enter the GID of the `oinstall` group as the value for the parameter `/proc/sys/vm/hugetlb_shm_group`. Doing this grants members of `oinstall` a group permission to create shared memory segments.

For example, where the `oinstall` group GID is 501:

```
# echo 501 > /proc/sys/vm/hugetlb_shm_group
```

After running this command, use `vi` to add the following text to `/etc/sysctl.conf`, and enable the `boot.sysctl` script to run on system restart:

```
vm.hugetlb_shm_group=501
```



Only one group can be defined as the `vm.hugetlb_shm_group`.

Modifiable Kernel Parameter Values for Oracle SPARC Solaris (64 bit) HP-UX and IBM AIX

Refer to the *Configuring Kernel Parameters* section in the following Oracle documents:

- *Database Quick Installation Guide for Oracle Solaris on SPARC (64 Bit)*
- *Database Quick Installation Guide for HP-UX Itanium*
- *Database Quick Installation Guide for IBM AIX on POWER Systems (64-Bit)*

Oracle Database Installation Steps

In order to install an Oracle database for use with the SA Model Repository, you must perform the following tasks which are explained in more detail in the following sections:

- 1 Create the database with the UTF8 database character set
- 2 Set the database with `TIME_ZONE` to '+00:00'
- 3 Create the database with the required initialization (`init.ora`) parameters
- 4 Create the database with required tablespaces
- 5 Create the database user `opsware_admin`
- 6 `tnsnames.ora` file requirements
- 7 File linking requirements
- 8 Enable Oracle Daylight Savings Time (DST)

1. UTF8 Database Character Set

Create the database with the UTF8 database character set:

```
CHARACTER SET UTF8
```

2. Set the Database `TIME_ZONE`

Create the database with `TIME_ZONE` set to '+00:00':

```
SET TIME_ZONE = '+00:00'
```

3. Specify the Required Initialization (`init.ora`) Parameters

Create the database instance with the following initialization (`init.ora`) parameters. For parameters not listed, SA assumes that the default Oracle parameters are used.

All Oracle 11g

```
db_file_multiblock_read_count := suggested to be >= 16
job_queue_processes := required to be >= 1000
remote_login_passwordfile := required to be = EXCLUSIVE
optimizer_mode := 'required to be = ALL_ROWS
recyclebin := required to be = OFF
event := required to be = 12099 trace name context forever, level 1
_complex_view_merging := required to be = FALSE
db_block_size := required to be >= 8192
processes := required to be >= 1024
undo_management := should be = AUTO
session_cached_cursors := required to be >= 50
```

undo_tablespace := should be = UNDO or other UNDO tablespace
optimizer_index_cost_adj := required to be = 20
optimizer_index_caching := required to be = 80
nls_length_semantics := required to be = CHAR
nls_sort := required to be = GENERIC_M

Oracle 11.2.0.x

deferred_segment_creation := required to be = FALSE
compatible := required to be >= 11.2.0
open_cursors := required to be >= 1500
cursor_sharing := required to be = FORCE
memory_target := required to be >= 1879048192 (1.75GB)
log_buffer := required to be >= 5242880
db_file_multiblock_read_count := required to be >= 16

Oracle 11.1.0.x

compatible := >= 11.1.0
open_cursors := required to be >= 1000
cursor_sharing := should be = SIMILAR (preferred) or
EXACT (recommended only if you encounter Oracle Bug No. 3102053)
memory_target := required to be >= 1694498816 (1616M)
log_buffer := required to be >= 1048576

4. Create the Required Tablespaces

The following tablespaces must be created to support SA. For tablespace disk space requirements, see [Model Repository \(Database\) Disk Space Requirements](#) on page 221.

- LCREP_DATA
- LCREP_INDX
- TRUTH_DATA
- TRUTH_INDX
- AAA_DATA
- AAA_INDX
- AUDIT_DATA
- AUDIT_INDX
- STRG_DATA
- STRG_INDX

5. Create the Database User `opsware_admin`

Create the database user 'opsware_admin' with the following privileges.

```
SQL create user opsware_admin identified by opsware_admin
      default tablespace truth_data
      temporary tablespace temp
      quota unlimited on truth_data;
```

```
SQL grant alter session to opsware_admin with admin option;
SQL grant create procedure to opsware_admin with admin option;
SQL grant create public synonym to opsware_admin with admin option;
SQL grant create sequence to opsware_admin with admin option;
SQL grant create session to opsware_admin with admin option;
SQL grant create table to opsware_admin with admin option;
SQL grant create trigger to opsware_admin with admin option;
SQL grant create type to opsware_admin with admin option;
SQL grant create view to opsware_admin with admin option;
SQL grant delete any table to opsware_admin with admin option;
SQL grant drop public synonym to opsware_admin with admin option;
SQL grant select any table to opsware_admin with admin option;
SQL grant select_catalog_role to opsware_admin with admin option;
SQL grant query rewrite to opsware_admin with admin option;
SQL grant restricted session to opsware_admin with admin option;
SQL grant create job to opsware_admin with admin option;
SQL grant manage scheduler to opsware_admin with admin option;
```

```
SQL grant execute on dbms_utility to opsware_admin with grant option;
SQL grant analyze any to opsware_admin;
SQL grant insert, update, delete, select on sys.aux_stats$ to opsware_admin;
SQL grant gather_system_statistics to opsware_admin;
SQL grant create any directory to opsware_admin;
SQL grant drop any directory to opsware_admin;
```

```
SQL grant alter system to opsware_admin;
SQL grant create role to opsware_admin;
SQL grant create user to opsware_admin;
SQL grant alter user to opsware_admin;
SQL grant drop user to opsware_admin;
SQL grant create profile to opsware_admin;
SQL grant alter profile to opsware_admin;
SQL grant drop profile to opsware_admin;
```

6. `tnsnames.ora` File Requirements

The `tnsnames.ora` file enables resolution of database names used internally by the core components. SA has the following requirements for the `tnsnames.ora` file:

- The file must reside in the following locations:
 - `/var/opt/oracle/tnsnames.ora`
 - `ORACLE_HOME/network/admin`
- If the core is installed across multiple servers, a copy of the file must reside on the servers hosting the following components:
 - Model Repository

- Infrastructure Component bundle (required by the Data Access Engine, Model Repository Multimaster Component, Software Repository Store)
- Slice Component bundle (required by the Command Center, Web Services Data Access Engine, Global File System)
- For a core installed on multiple servers, the directory path of the `tnsnames.ora` file must be the same on each server.
- In a Single Core installation, the `tnsnames.ora` file must contain an entry for the Model Repository, as in the following example:

```
truth = DESCRIPTION= (ADDRESS= (HOST=magenta.example.com) (PORT=1521)
(PROTOCOL=tcp)) (CONNECT_DATA= (SERVICE_NAME=truth)) )
```

tnsnames.ora: Multimaster Mesh Requirements

In a Multimaster Mesh, the `tnsnames.ora` file must be set up for a Source Core and a Destination Core using the following guidelines.

Source Core

The `tnsnames.ora` file must contain an entry for its own Model Repository. The port number must be set to the port that you have designated that the Oracle listener process use, such as 1521 (default), 1526, and so on.

The `tnsnames.ora` file must also contain an entry that specifies the Source Core Management Gateway. This port is used by the Data Access Engine for Multimaster traffic. The port number is derived from the following formula: (20000) + (facility ID of the Destination Core).

Example: In the following example, the TNS service name of the Source Core is `orange_truth`, which runs on the host `orange.example.com`. The TNS name of the Destination Core is `cyan_truth`, which has a facility ID of 556. Note that the entry for `cyan_truth` specifies `orange.example.com`, which is the host running the Source Core's Management Gateway.

```
orange_truth= (DESCRIPTION= (ADDRESS= (HOST=orange.example.com) (PORT=1521)
(PROTOCOL=tcp)) (CONNECT_DATA= (SERVICE_NAME=truth)) )
cyan_truth= (DESCRIPTION= (ADDRESS= (HOST=orange.example.com) (PORT=20556)
(PROTOCOL=tcp)) (CONNECT_DATA= (SERVICE_NAME=truth)) )
```

Destination Core

The `tnsnames.ora` file must contain an entry for its own Model Repository. The port number must be set to the port that you have designated that the Oracle listener process use, such as 1521 (default), 1526, and so on. The `tnsnames.ora` file does not require any entries for other cores in the mesh.

Example: In the following example, the TNS service name of the Destination Core is `cyan_truth`, and the core runs on the host, `cyan.example.com`.

```
cyan_truth= (DESCRIPTION= (ADDRESS= (HOST=cyan.example.com) (PORT=1521)
(PROTOCOL=tcp)) (CONNECT_DATA= (SERVICE_NAME=truth)) )
```

7. File Linking Requirements:

After creating the database, but before installing the Model Repository with the SA Installer, perform the following tasks:

- 1 Create the `tnsnames.ora` file in the following directory:

```
# ORACLE_HOME/network/admin
```

Verify that the file conforms to the rules listed in [6. tnsnames.ora File Requirements](#) on page 231.

- 2 If it does not exist, create the following directory:

```
# mkdir -p /var/opt/oracle
```

- 3 Create the following symbolic link:

```
# ln -s $ORACLE_HOME/network/admin/tnsnames.ora \  
/var/opt/oracle/tnsnames.ora
```

Ensure that the Oracle Unix user has read-write permission on the `tnsnames.ora` file.

For Linux:

Create a symbolic link for the `oratab` file:

```
# ln -s /etc/oratab /var/opt/oracle/oratab
```

Finally:

- 1 Copy the sample `opsware-oracle` script to `/etc/init.d/`.
- 2 Link `/etc/init.d/opsware-oracle` to corresponding scripts in the `/etc/rc*` directories. For example:

```
# ln -s /etc/init.d/opsware-oracle \  
/etc/rc0.d/K02opsware-oracle  
# ln -s /etc/init.d/opsware-oracle \  
/etc/rc1.d/K02opsware-oracle  
# ln -s /etc/init.d/opsware-oracle \  
/etc/rc2.d/S60opsware-oracle  
# ln -s /etc/init.d/opsware-oracle \  
/etc/rcS.d/K02opsware-oracle
```

8. Enable Oracle Daylight Savings Time (DST)

To enable Daylight Saving Time for the Oracle database, you must apply database tier patches. To apply these patches, perform the following steps:

- 1 Use MetaLink Note 359145.1 to apply Oracle Database time zone fixes specific to your database version.
- 2 Use MetaLink Note 359145.1 to apply time zone fixes to the Oracle Java Virtual Machine (JVM) in the Oracle Database specific to your E-Business Suite database version.

SA Database Installation Sample Scripts

HP Support can provide sample scripts for steps 1 through 5 of the Oracle Database Installation Steps.

Oracle/SA Installation Scripts, SQL Scripts, and Configuration Files

- **create_os_user_oracle:** A script that lists the commands for creating the OS user `oracle`.
- **truth.sh:** A shell script that creates directories and then launches the `truth.sql` script. Running this script causes all the scripts to be run automatically, in the correct order.

- **truth.sql**: Prompts for passwords of the SYS and SYSTEM users and launches the remainder of the SQL scripts in this list.
- **CreateDB.sql**: Creates a database with the UTF8 character set and TIME_ZONE set to '+00:00'
- **CreateDBFiles.sql**: Creates the following tablespaces that are required by SA:

```
LCREP_DATA
LCREP_INDX
TRUTH_DATA
TRUTH_INDX
AAA_DATA
AAA_INDX
AUDIT_DATA
AUDIT_INDX
STRG_DATA
STRG_INDX
```

See [Model Repository \(Database\) Disk Space Requirements](#) on page 51 for additional tablespace sizing information.

- **CreateDBCatalog.sql**: Runs Oracle scripts to create data system catalog objects.
- **JServer.sql**: Sets up the Oracle Java environment.
- **CreateAdditionalDBFiles.sql**: Adds data and index files to certain tablespaces and allocates additional disk space. This script is optional, but recommended.
- **CreateUserOpware_Admin.sql**: Creates the opware_admin database user and grants permissions (privileges) to this user (required by SA).
- **postDBCcreation.sql**: Creates the spfile from the pfile (parameter file).
- **init.ora**: Contains initialization parameters for the database. See [3. Specify the Required Initialization \(init.ora\) Parameters](#) on page 229.
- **tnsnames.ora**: Enables resolution of database names used internally by SA.
- **listener.ora**: Contains configuration parameters for the listener. SA by default listens on port 1521. You can change the default port during installation or by editing the tsnames.ora file.



The SA-supplied Oracle 11.2.02 database has a new listener.ora parameter:

```
SUBSCRIBE_FOR_NODE_DOWN_EVENT_LISTENER=
```

Default is OFF. This parameter must be set to OFF for non-RAC installations. For more information about this parameter, see the Oracle documents IDs 372959.1 and 437598.1.

- **bash_profile** or **profile**: Sets environment variables and sets shell limits for the oracle Unix user.
- **opware-oracle**: A script residing in /etc/init.d that starts up and shuts down the database and listener.



The /etc/init.d/opware-sas start script, which starts and stops the SA components, does not start and stop the database and listener. For more information on the opware-sas start script, see “Start Script for SA” in the *SA Administration Guide*.

Creating the Database using the SA-Supplied Scripts

To create the Oracle database using the SA-supplied scripts, perform the following steps:

- 1 Obtain the database creation scripts from your HP support representative.
- 2 Make any required changes to the scripts.
- 3 As root, create the Unix user `oracle` and log in to the server as the user `oracle`.
- 4 Copy the SA-supplied files to the following directory:
`$ORACLE_BASE/admin/truth/create`
- 5 Change the mode of the SA-supplied `truth.sh` script:
`# chmod 755 truth.sh`
- 6 Launch the SQL scripts that create the database by running the `truth.sh` script:
`# ./truth.sh`
- 7 After the scripts launched by `truth.sh` complete, check the log files in the following directory for errors:
`/u01/app/oracle/admin/truth/scripts/*.log`

SA-Supplied Oracle RDBMS Software and Database Setup



If you plan to install the Oracle RDBMS software and database yourself, you do not need to perform the tasks in this section. See [Non-SA-Supplied Oracle Software and Database Setup](#) on page 227.

If you plan to use a SA-supplied Oracle database with the SA Model Repository, you should read the following sections for information about what the SA Installer does when installing the Oracle software and database during SA installation. The SA Installer performs all the tasks discussed below. You should also review [System Requirements](#) on page 219 before proceeding with this section.

SA-Supplied RDBMS Configuration Details

When you install the SA-supplied Oracle RDBMS using the SA Installer Oracle installation option, the installer:

- Checks that all requirements are met on the host server (see [System Requirements](#) on page 219).
- Sets certain kernel parameters to required values (see [Modifiable Kernel Parameters](#) on page 227).
- Creates the Unix user `oracle` locally in `/etc/passwd`.
- Creates the Unix groups `dba` and `oinstall` locally in `/etc/group`.
- Sets the `$ORACLE_HOME` environment variable to the following directory:
`/u01/app/oracle/product/11.2.0/db_2`
- Sets the `$ORACLE_SID` environment variable to `truth`.

- Creates a database with the UTF8 character set, `TIME_ZONE` set to '+00:00' and with required `init.ora` parameters.
- Creates the tablespaces and data and index files under the following directories:

```
/u01/oradata/truth
/u02/oradata/truth
/u03/oradata/truth
/u04/oradata/truth
```

The system administrator can configure the `/u01`, `/u02`, `/u03`, `/u04` directories before installing the Oracle RDBMS software.

- Gets the service name (TNS name) from the SA Installer interview (`truth.servicename` prompt) and inserts it into the `tnsnames.ora` file in `$ORACLE_HOME/network/admin` and `/var/opt/oracle`. The SA Installer changes the value of the `host` parameter in `tnsnames.ora` to the value returned by the Unix `hostname` command.
- In the `/$ORACLE_HOME/network/admin/listener.ora` file, changes the value of the `host` parameter to the value returned by the Unix `hostname` command.

The listener is password protected and OS authenticated. (The default password is `opsware`.) By default, it listens on port 1521.

- Creates the `/etc/init.d/opsware-oracle` script, which you can use to start up and shut down the database and listener.

This script is linked to corresponding scripts in the `/etc/rc*.d` directories.

- Creates the user `opsware_admin` with the required privileges.
- After installation is complete, you can examine the logs that are created here:

```
/var/log/opsware/install_opsware
```

Security

SA recommends that you change the default passwords for the following:

- the Unix user `oracle`
- the Oracle database users `SYS` and `SYSTEM`
- the Oracle listener

SA does not use the `SYS` and `SYSTEM` users.

In the `/$ORACLE_HOME/network/admin/listener.ora` file, SA sets the value of the `host` parameter to the value returned by the Unix `hostname` command. The listener is password protected and OS authenticated. The default password is `opsware`. By default, the Oracle listener uses port 1521.

SA-Supplied Oracle Installation Procedure

SA supports the following SA/Oracle database configurations:

- SA Core and Oracle database on a single host
- SA Core with the Oracle database on a remote database server

See [Chapter 2, “SA Core Configurations Supported For Customer Installation”](#) for a description of supported SA Core/Oracle database configurations and [Chapter 5, “SA Core Installation”](#) for installation instructions.

Installing the Model Repository Database on a Remote Server

To install or upgrade the Model Repository Oracle database on a remote server, perform the following tasks:

- 1 Perform the following tasks on the server on which you will run the SA Installer:
 - a Install the Oracle Full Client software.

The steps below use `/u01/app/oracle/product/11.2.0/client_1` as the Oracle Full Client home.



The Oracle Full Client must be the same version as the Oracle database.

- b Ensure that the Oracle Full Client software is owned by the OS user `oracle`.
 - c Copy the database server's `/var/opt/oracle/tnsnames.ora` file to the client machine's `$ORACLE_HOME/network/admin/tnsnames.ora`. Ensure that the hostname in the file resolves properly.
 - d If it does not exist, create the following directory:

```
# mkdir -p /var/opt/oracle
```
 - e Create the following symbolic link:

```
# ln -s $ORACLE_HOME/network/admin/tnsnames.ora /var/opt/oracle/tnsnames.ora
```
 - f Ensure that the Unix user `oracle` has read-write permission on the `tnsnames.ora` file.
 - g Ensure that the SA Installer Core Definition File (CDF) has the correct path to the client `tnsnames.ora` file (`%truth.tnsdir`), oracle client home (`%db.orahome`), database server name/IP (`%db.host`), listener port (`%db.port`), SA Installer machines subdomain (`%truth.dcSubDom`), and so on. Based on the above steps your parameter values will be:
 - `%truth.tnsdir=/var/opt/oracle`
 - `%db.orahome=/u01/app/oracle/product/11.2.0/client_1`
 - `%db.port=1521`
 - `%truth.dcSubDom=prod.example.com`
 - `db.host=192.168.9.99` (server on which the Oracle database is installed)
 - h Ensure that the `COMPATIBLE` parameter is set correctly and that it matches the database version. For example, for database software that is version 11.2.0.2 ensure that `COMPATIBLE=11.2.0.2`. SA uses Oracle's Export Data Pump and Import Data Pump utilities during secondary core creation. These utilities require the `COMPATIBLE` parameter be specified correctly.
- 2 Perform the following tasks on the Model Repository host:
 - a Log in as the user `oracle`.
 - b Change to the `$ORACLE_HOME/network/admin` directory.

- c Ensure that the listener is started with the command:

```
# lsnrctl start <your_listener_name>
```

Oracle RAC Support

SA supports Oracle Real Application Clusters (RAC).



Oracle RAC support requires a new installation of both Oracle and SA. Therefore, in order to enable Oracle RAC support in SA, you must first install Oracle RAC 11g, configured as described in the following sections.

Supported Oracle Versions and Operating Systems

Support for the Model Repository is limited to certain versions of Oracle running on certain versions of operating systems. HP strongly recommends that you also apply the latest Oracle CPU or PSU patches.

- See the *SA 10.0 Support and Compatibility Matrix* for a list of supported Oracle versions and operating systems.

System Requirements

See [System Requirements](#) on page 219.

Set up the Oracle RAC Database/Instances

SA supports any valid Oracle RAC configuration, such as any number of nodes, ASM or regular disks, and so on.

However, the Oracle database must be configured for use with SA. You may require your Oracle DBA's help to configure the Oracle RAC/instances, the required initialization parameters, the required tablespaces, the `opsware_admin` database user, and the `listener.ora` and `tnsnames.ora` files.

Create the Database with the Required Initialization Parameters

Follow the procedure described in [Non-SA-Supplied Oracle Software and Database Setup](#) on page 227.

You should perform the following tasks listed in the [Non-SA-Supplied Oracle Software and Database Setup](#) section:

- Modifiable Kernel Parameters
- Oracle Database Installation Steps
 - a Create the database with the UTF8 database character set
 - b Set the database with `TIME_ZONE` to `' +00:00 '`
 - c Create the database with the required initialization (`init.ora`) parameters

- d Create the database with required tablespaces
- e Create the database user `opsware_admin`



Use the `tnsnames.ora` file and file linking requirements listed in the following sections since they differ for the Oracle RAC environment from those listed in [Non-SA-Supplied Oracle Software and Database Setup](#).

[Non-SA-Supplied Oracle Software and Database Setup](#) on page 227 describes the required database setup, Oracle initialization parameters, required tablespaces, database user `opsware_admin`, etc.

Installing the Model Repository

In most production environments with Oracle RAC, you can perform the Model Repository installation from any SA server. The database server or RAC nodes in this case are considered to be remote.

The examples used in the following sections assume an SA server (`rac1sa.dev.opsware.com`) on which SA will be installed and a 2 node RAC configuration shown in [Table 52](#)

Table 52 Example RAC Configurations

Identity	Host Note	Name	Type	Address	Address Static or Dynamic	Resolved By
Node 1 Public	rac1pub	rac1pub	Public	192.168.173.210	Static	DNS
Node 1 Virtual	Selected by Oracle Clusterware	rac1-vip	Virtual	192.168.173.212	Static	DNS and/or host file
Node 1 Private	rac1pub	rac1prv	Private	172.16.1.100	Static	DNS, host file or none
Node 2 Public	rac2pub	rac2pub	Public	192-168-173-211	Static	DNS
Node 2 Virtual	Selected by Oracle Clusterware	rac2-vip	Virtual	192.168.173.213	Static	DNS and/or host file
Node 2 Private	rac2pub	rac2prv	Private	172.16.1.101	Static	DNS, host file or none
SCAN vip 1	Selected by Oracle Clusterware	sa_cluster1-scan	Virtual	192.168.173.216	Static	DNS
SCAN vip 2	Selected by Oracle Clusterware	sa_cluster1-scan	Virtual	192.168.173.217	Static	DNS
SCAN vip 3	Selected by Oracle Clusterware	sa_cluster1-scan	Virtual	192.168.173.218	Static	DNS

Installing the Model Repository in a RACed Environment

In an Oracle RAC environment, only one of the RAC nodes is used during the SA installation/upgrade process. The SA Installer connects to only one Oracle RAC instance to install/modify the Model Repository. During the regular SA operations, all RAC nodes are used.

Perform the following tasks on the SA server on which you will run the SA Installer, for example `rac1sa.dev.opsware.com`.

1 Model Repository Hostname Resolution

Ensure that all RAC database related host names, scan addresses and RAC instance hosts are resolvable by:

- adding the domain and name servers to `/etc/resolv.conf` or by
- adding the RAC hosts entries to `/etc/hosts`

On the server where you will run the SA Installer, ensure that the Model Repository hostname `truth` resolves to the remote database server, not to the server on which you will be running the SA Installer:

In `/etc/hosts`, enter the public IP address of one of the RAC nodes/instances. For example the `/etc/hosts` file on `rac1sa.dev.opsware.com` would have the following entry:

```
192.168.173.210 truth rac1pub rac1pub.dev.opsware.com
```



If you have set up Oracle Clusterware, you should use the Clusterware IP address rather than a single database node IP address. For example:

```
192.168.173.216 truth sacluster1-scan sacluster1-scan.dev.opsware.com
```

2 Install the Oracle 11g Full Client on the SA Server



For Oracle 11.2.0.3, use the Oracle Full Client version 11.2.0.3.

- a The SA Installer uses the Oracle Full Client to connect to the SA server and install the Model Repository. Below are sample commands for installing the Oracle full client.

Create the database user `oracle` for the Oracle Full Client installation:

```
root@rac1sa ~]# mkdir -p /u01/app/oracle
root@rac1sa ~]# mkdir -p /u01/app/orainventory
root@rac1sa ~]# groupadd oinstall
root@rac1sa ~]# groupadd dba
root@rac1sa ~]# useradd -c "Oracle Client software owner" -g oinstall -G
dba -d /u01/app/oracle -s /bin/bash oracle
root@rac1sa ~]# chown -R oracle:oinstall /u01/app
root@rac1sa ~]# chmod -R 775 /u01/app
root@rac1sa ~]#passwd oracle (change oracle user password)
```

- b Create the `.bash_profile` file

In `/u01/app/oracle` create the `.bash_profile` file.

Temporarily comment out `ORACLE_HOME` and `ORACLE_PATH`. You will uncomment these entries after the Oracle client installation is complete.

Sample .bash_profile File

```
# .bash_profile
# Get the aliases and functions
if [ -f ~/.bashrc ]; then
. ~/.bashrc
fi

# User specific environment and startup programs
PATH=$PATH:$HOME/bin
export PATH

# SA-OracleRAC parameters begin
# unset USERNAME
export ORACLE_BASE=/u01/app/oracle
# export ORACLE_HOME=$ORACLE_BASE/product/11.2.0/client_1
# PATH=$ORACLE_HOME/bin:$ORACLE_HOME/OPatch:$PATH
export PATH

if [ -t ]; then
stty intr ^C
fi

umask 022
# SA-OracleRAC parameters end
```

c Install the Oracle Full Client.

Install the Oracle Full Client as described in your Oracle documentation. You can create a share to access the Oracle Full Client binaries.

d Set Up Terminals.

You will need two X window terminals to install the Oracle Full Client:

Terminal 1: log in as root and enter the commands:

Terminal 1> xhost +

Terminal 2: ssh -X oracle@<new_oracle_full_client_host>

e Start Oracle Full Client installation

From Terminal 2, run the Oracle Universal Installer (OUI). The Oracle Full Client is installed in:

/u01/app/oracle/product/11.2.0/client_1

f Run the Oracle Universal Installer to install Oracle Full Client. The directories in this example assume an Oracle 11g Full Client on Linux.

1. cd /location_of_oracle_full_client.

2. ./runInstaller.

3. At the Welcome Screen, click Next.

4. Specify the Inventory Directory and Credentials (/u01/app/oraInventory and /u01/app/oinstall).

5. For Select Installation Type, choose Administrator, click Next.

6. For ORACLE_BASE select: /u01/app/oracle, click Next.

7. The Oracle Universal Installer performs some checks. If the checks are not successful, fix the issue and re-run this step. If the checks are successful click Next.
8. The Oracle OUI will list the products that are to be installed. Click Install.
9. The OUI shows the progress bar while installing.
10. On the 'Welcome to Oracle Net Configuration Assistant' window click on Next.
11. Click Finish when the installation completes.
12. You must run the following two configuration scripts as root after installation completes:

```

-u01/app/oraInventory/orainstRoot.sh
-u01/app/oracle/product/11.2.0/client_1/root.sh

```

- g Verify that the `.bash_profile` file for the user `oracle` is correct.
- h Uncomment `$ORACLE_HOME` and `$ORACLE_PATH`.

3 Making Changes to `tnsnames.ora` on an SA Server (Use `tnsnames.ora-install_upgrade` File)

By default SA expects the `tnsnames.ora` file to be located in `/var/opt/oracle`.

- a Login as root on the SA server from which the installer will be run.
- b Enter the command:

```
# mkdir -p /var/opt/oracle
```

- c Copy `tnsnames.ora` from the remote database server to the directory you created above.

For the RAC environment, copy `tnsnames.ora` from RAC Node 1 (for example, `rac1pub.dev.opsware.com`).

To accommodate the remote Model Repository installation process, two sets of `tnsnames.ora` files are required on the SA server due to an SA Installer limitation in recognizing RAC environments.

- **`tnsnames.ora-install_upgrade`** – this copy of `tnsnames.ora` is used during SA installation/upgrade. The file can be renamed.
- **`tnsnames.ora-operational`** – this copy of `tnsnames.ora` is used during normal SA operation. The file can be renamed.

You can use softlinks to point `tnsnames.ora` to either `tnsnames.ora-install_upgrade` or `tnsnames.ora-operational`. For example:

```
# ln -s tnsnames.ora-install_upgrade tnsnames.ora
```

`tnsnames.ora-install_upgrade` Sample File

```

# tnsnames.ora Network Configuration File: /u01/app/oracle/product/11.2.0/db_2/
network/admin/tnsnames.ora
# Generated by Oracle configuration tools.
RAC1SA_TRUTH =
(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP) (HOST = rac1pub.dev.opsware.com)
(PORT = 1521)) (CONNECT_DATA =) (SERVER = DEDICATED) (SID = truth1))

```

Testing the Connection from the SA host to the Database

Before starting the Model Repository installation/upgrade, you can perform the following tests to verify that your `tnsnames.ora` file is configured correctly and if the SA Installer can connect to the database.

- 1 Verify that the SA server's `/var/opt/oracle/tnsnames.ora` file is configured correctly as described in [Making Changes to tnsnames.ora on an SA Server \(Use tnsnames.ora-install_upgrade File\)](#) on page 242.
- 2 On the SA server:
 - a Login as `oracle` or `root` or `su - twist/spin` – if these users exist.
 - b `# export ORACLE_HOME=/u01/app/oracle/product/11.2.0/client_1` (or where you installed the Oracle Full Client)
 - c `# export LD_LIBRARY_PATH=$ORACLE_HOME/lib`
 - d `# export TNS_ADMIN=/var/opt/oracle`
 - e `# set path=$PATH:$ORACLE_HOME/bin`
 - f `# sqlplus sys/password@RAC1SA_TRUTH as sysdba;`
where `rac1sa_truth` is the `service_name` or entry from the `tnsnames.ora` file
 - g `SQL> connect opsware_admin/<password>@RAC1SA_truth`
If you are able to logon to the database then all files are configured correctly.

SA Installer Core Definition File (CDF)

You can now start the installation of the SA Model Repository. Use the SA Installer's Advanced Interview mode to set the various parameter values. Ensure that you have the correct parameter values for the installation interview or that you have a previous Core Definition File (CDF).

- `%db.sid: truth1` (Oracle SID of the instance where SA installer is going to connect to.)
- `%db.oraHome: /u01/app/oracle/product/11.2.0/client_1` (Oracle client home)
- `%db.port: 1521` (Oracle listener port)
- `%db.host: 192.168.173.210` (server where Oracle RDBMS is installed)
- `%truth.servicename: rac1sa_truth` (value of service name from `tnsnames.ora` file)

You can now install the SA Core as described in the [Chapter 5, "SA Core Installation"](#).

During the Secondary Core installation you will run the `hpsa_add_dc_to_mesh.sh` script. Ensure that you have the correct parameter values for the Secondary Core.

- `newCore.servicename: value of the service name for the Secondary Core from tnsnames.ora file)`
- `%db.sid: truth1` (In the Secondary Core environment, the Oracle SID of the instance where the SA Installer connects.)

Post SA Installation Process

After you install the SA Core, perform the following tasks in order to use all the nodes in the Oracle RAC environment.

Making Changes to `tnsnames.ora` on the SA Server (Use `tnsnames.ora-operational` File)

After SA Core installation is complete, the `tnsnames.ora` file should point/link to the `tnsnames.ora-operational` file.

In an Oracle RAC environment, only one of the RAC nodes or instances is used during the installation/upgrade process. The SA Installer connects to only one Oracle instance to modify the Model Repository. During normal SA operations, all the RAC nodes are used.

To accommodate the remote database installation process, two sets of `tnsnames.ora` files are required on the SA server.

- **`tnsnames.ora-install_upgrade`** – this copy of `tnsnames.ora` is used during SA installation/upgrade. You can rename the file.
- **`tnsnames.ora-operational`** – this copy of `tnsnames.ora` is used during normal SA operation. You can rename the file.

You can use softlinks to point `tnsnames.ora` to either `tnsnames.ora-install_upgrade` or `tnsnames.ora-operational`:

```
# ln -s tnsnames.ora-operational tnsnames.ora
```

`tnsnames.ora-operational` Sample File



If you have set up Oracle Clusterware, you should use the Clusterware IP address rather than a single database node IP address.

Make a note of the text that is in **bold** letters. This `tnsnames.ora` file is used during normal SA operation and contains the RAC parameters.

`tnsnames.ora-operational` sample File - with a SCAN IP Address

If you have set up an Oracle Single Client Access Name (SCAN) IP address, use the following:

#This entry is for connecting to RAC virtual machines. This entry is used by SA during operation of SA.

```
RAC1SA_TRUTH=  
(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=sacluster1-scan.dev.opsware.com)  
(PORT=1521))(CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=TRUTH)))
```

#This entry is for connecting to node2 via `service_name`. This is for DBA convenience. This is not used by SA.

```
TRUTH2 =  
(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=rac2pub.dev.opsware.com)  
(PORT=1521))(CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=truth)))
```

#This entry is for connecting to node1 via `service_name`. This is for DBA convenience. This is not used by SA.

```
TRUTH1 =  
(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=rac1pub.dev.opsware.com)  
(PORT=1521))(CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=truth)))
```

During installation, the SA Installer adds an SA Gateway entry into `tnsnames.ora` file (linked to `tnsnames.ora.install_upgrade`) on the Primary SA Core.

After installation completes, copy that entry into the `tnsname.ora.operational` file. If this entry is not present in the `tnsname.ora.operational` file, Multimaster Mesh transactions will not flow. The following is a sample gateway entry from `tnsnames.ora`:

```
RAC2SA_TRUTH=(DESCRIPTION=(ADDRESS=(HOST=192.168.173.214)
(PORT=20002)(PROTOCOL=tcp))(CONNECT_DATA=(SERVICE_NAME=truth)))
```

tnsnames.ora-operational Sample File - Without a SCAN IP Address

If you have not set up an Oracle Single Client Access Name (SCAN) IP address, use the following:

#This entry is for connecting to RAC virtual machines.

```
RAC1SA_TRUTH =
(DESCRIPTION =
(ADDRESS = (PROTOCOL = TCP)(HOST = rac1pub.dev.opsware.com) (PORT =
1521))(ADDRESS = (PROTOCOL = TCP)(HOST = rac2pub.dev.opsware.com) (PORT =
1521))(LOAD_BALANCE = yes)(CONNECT_DATA =(SERVER = DEDICATED) (SERVICE_NAME =
truth)(FAILOVER_MODE =)(TYPE = SELECT)(METHOD = Preconnect) (RETRIES =
180)(DELAY = 5)))
```

#This entry is for connecting to node2 via `service_name`. This entry is optional. This is for DBA convenience. This is not used by SA.

```
TRUTH2 =
(DESCRIPTION =(ADDRESS = (PROTOCOL = TCP)(HOST = rac2pub.dev.opsware.com)
(PORT = 1521))(CONNECT_DATA =(SERVER = DEDICATED)(SERVICE_NAME = truth)
(INSTANCE_NAME = truth2)))
```

#This entry is for connecting to node1 using `service_name`. This entry is optional. This is for DBA convenience. This is not used by SA.

```
TRUTH1 =
(DESCRIPTION =(ADDRESS = (PROTOCOL = TCP)
(HOST = rac1pub.dev.opsware.com)(PORT = 1521))
(CONNECT_DATA =(SERVER = DEDICATED)(SERVICE_NAME = truth)
(INSTANCE_NAME = truth1)))
```

Use softlinks to link the file to `tnsnames.ora` file after SA installation is complete and you are ready to start SA in operational mode.

```
# ln -s tnsnames.ora-operational tnsnames.ora
```

vault.conf File Changes



If you have set up Oracle Clusterware, you should use the Clusterware IP address rather than a single database node IP address.

In an Oracle RAC environment, the `vault.conf` file must be modified after SA installation is complete. Modify `/etc/opt/opsware/vault/vault.conf` to specify the complete `tnsname.ora` definition instead of the SID. For example:

- *If you have set up an Oracle Single Client Access Name (SCAN) IP address, use the following:*
 - Before:

```
db.sid: truth
```

— After:

```
#truth.sid: truth1
truth.sid: (DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)
(HOST = sacluster1-scan.dev.opsware.com) (PORT = 1521))
```

- *If you have not set up an Oracle Single Client Access Name (SCAN) IP address, use the following:*

```
#truth.sid: truth1
DESCRIPTION = (ADDRESS = (PROTOCOL = TCP) (HOST = rac1pub.dev.opsware.com)
(PORT = 1521)) (ADDRESS = (PROTOCOL = TCP) (HOST = rac2pub.dev.opsware.com)
(PORT = 1521)) (LOAD_BALANCE = yes) (CONNECT_DATA = (SERVER = DEDICATED)
(SERVICE_NAME = truth) (FAILOVER_MODE =) (TYPE = SELECT)
(METHOD = Preconnect) (RETRIES = 180) (DELAY = 5)) ) )
```

- Also, ensure that these values are correct:

```
truth.port: 1521
truth.host: 192.168.173.210 (database server IP)
truth.servicename: rac1sa_truth (tnsnames.ora enter)
```

- Restart the vaultdaemon:

```
# /etc/init.d/opsware-sas restart vaultdaemon
```

Upgrading the Model Repository in a RACed Environment

To upgrade the Model Repository in an Oracle RAC environment, follow the same procedure as [Installing the Model Repository](#) on page 239. If you are doing a remote database installation, ensure that you modify the `tnsnames.ora` file on the server on which the SA Installer is run. HP recommends that you test the connection as described in [Testing the Connection from the SA host to the Database](#) on page 243.

Garbage Collection

The Garbage Collector (GC) is a stored procedure written in PL/SQL that runs in the database on a schedule. The GC procedures look at the `AUDIT_PARAMS` table to determine the retention period to use to delete the old data. The GC PL/SQL procedures are managed by Oracle's `dba_scheduler_jobs`.

Data Retention Period

When GC runs, it looks at the values in the `AUDIT_PARAMS` table to determine what retention period to use when deleting objects.



The `AUDIT_PARAMS` table is not replicated, so there is a possibility that these retention periods may become unsynchronized, which can cause severe Multimaster conflict issues. You must ensure that the values in the `AUDIT_PARAMS` table are exactly the same for all the cores in a mesh.

```
# Sqlplus "/ as sysdba"
SQL> col name format a20;
SQL> col value format a20;
```

```
SQL> col AUDIT_PARAM_ID format a15;
SQL> select AUDIT_PARAM_ID, NAME, VALUE from audit_params;
```

The parameters from AUDIT_PARAMS table and their default values are:

AUDIT_PARAM_ID	NAME	VALUE	
2	DAYS_WAY	30	(These are the completed way sessions)
3	DAYS_CHANGE_LOG	180	(These are the server history events)
4	LAST_DATE_WAY	04-APR-09	
5	LAST_DATE_CHANGE_LOG	05-NOV-08	
6	DAYS_AUDIT_LOG	180	(These are the audit logs)
7	LAST_DATE_AUDIT_LOG	05-NOV-08	
8	DAYS_WLM	30	(These are completed WLM jobs)
9	LAST_DATE_WLM	05-NOV-08	



As of SA 9.10, the DAY_TRAN parameter that controlled retention time for transactions was removed. To control transaction retention time, instead use the system configuration parameter `vault.garbageCollector.daysToPreserve`.

Select the **Administration** tab in the SA Client, then select System Configuration in the navigation panel. Select Model Repository Multimaster Component. Locate and change the value of the above parameter. (The default value is 7.) Select the Save button to save your change.

Modifying the Retention Period Values

To update the data, run a SQL command similar to the following example as user LCREP:

```
# Su - oracle
# Sqlplus "/ as sysdba"
SQL> grant create session to lcrep;
SQL> connect lcrep/<password>
SQL> update AUDIT_PARAMS set value=30 where name = 'DAYS_AUDIT_LOG';
```



The values in the AUDIT_PARAMS table must be exactly the same for all the cores in a mesh.

Viewing GC DBA_SCHEDULER_JOBS

When the Model Repository is installed, the SA Installer sets up these jobs, which perform garbage collection.

GC jobs can be viewed by logging in to SQL*Plus and running the following SQL commands:

```
# Su - oracle
# Sqlplus "/ as sysdba"
SQL> set line 200
SQL> col job_name format a50
SQL> col owner format a14
SQL> col last_date format a17
SQL> col next_date format a17
SQL> col job_action format a50
```

```
SQL> select job_name, owner, to_char(LAST_START_DATE, 'MM/DD/YY HH:MI:SS')
last_date,to_char(next_run_date, 'MM/DD/YY HH:MI:SS') next_date, job_action
from dba_scheduler_jobs where owner='GCADMIN';
```

JOB_NAME	OWNER	LAST_DATE	NEXT_DATE	JOB_ACTION
WLPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	WLPURGE.GC_JOBS
STORAGEINITIATORPURGE_GC	GCADMIN	04/02/12 09:47:30	04/03/12 10:47:30	STORAGEINITIATORPURGE.GC_ STORAGEINITIATORS
AUDITPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	AUDITPURGE.GC_AUDITLOGS
CHANGELOGPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	CHANGELOGPURGE.GC_CHANGELOGS
WAYPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	WAYPURGE.GC_SESSIONS

where:

WAYPURGE.GC_SESSIONS - Performs a sessions garbage collection

CHANGELOGPURGE.GC_CHANGELOGS - Performs a changelogs garbage collection

AUDITPURGE.GC_AUDITLOGS - Performs auditlogs garbage collection

STORAGEINITIATORPURGE.GC_STORAGEINITIATORS - Performs storage data garbage collection

WLPURGE.GC_JOBS - Performs WLM garbage collection

Manually Running GC Jobs

You can run GC jobs by logging in to SQL*Plus and entering the following:

```
# Su - oracle
# Sqlplus "/ as sysdba"

SQL> grant create session to gcadmin
SQL> connect gcadmin/<password>
SQL> exec dbms_scheduler.run_job('<job_name_value>');
```

For example, this sample command runs the waypurge_gc job:

```
SQL> exec dbms_scheduler.run_job('WAYPURGE_GC');
```

Monitor the ERROR_INTERNAL_MSG Table

The garbage collection jobs write exceptions to the truth.ERROR_INTERNAL_MSG table. You can monitor this table for errors (checking daily is recommended). Data from this table should be purged by the DBA on a regular basis. For example:

```
# Su - oracle
# Sqlplus "/ as sysdba"

SQL> set line 200
SQL> col ERR_ID format 999999
SQL> col ERR_USER format a8
SQL> col ERR_TABLE format a25
SQL> col ERR_TABLE_PK_ID format a10
SQL> col ERR_CODE format 9999999
SQL> col ERR_TEXT format a20
SQL> col ERR_INFO format a30
```

```
SQL> select ERROR_INTERNAL_MSG_ID ERR_ID,
        ERR_DATE,
        ERR_USER,
        ERR_TABLE,
        ERR_TABLE_PK_ID,
        ERR_CODE,
        ERR_TEXT,
        DELETE_FLG,
        ERR_INFO
from ERROR_INTERNAL_MSG
order by ERR_DATE;
```

Database Monitoring Strategy

Because the Model Repository is a critical component of SA, the DBA should implement a monitoring strategy. The DBA can write custom monitoring scripts or use third-party products.

This section contains example commands for monitoring the Oracle database used by the Model Repository. When issuing the commands shown in this section, you must log on to the server as the user `oracle` and run SQLPlus:

```
# su - oracle
# sqlplus "/ as sysdba"
```

Verify that the Database Instances are Up and Responding

To verify that the Database Instances are up and running, perform the following steps:

- 1 Check to see if the Oracle processes are running by entering the following command:

```
# ps -ef | grep ora_
```

The `ps` command will generate 15-20 session information similar to the following:

```
oracle  14674      1  0 Apr18 ?          00:00:00 ora_*_truth
```

where `*` is `smon`, `dbw0`, `ckpt` etc.

- 2 Verify that the database status is `ACTIVE` by entering the following command in `sqlplus`:

```
SQL> select database_status from v$instance;
```

- 3 Verify that the open mode is `READ WRITE` by entering the following command in `sqlplus`:

```
SQL> select name, log_mode, open_mode from v$database;
```

Verify that the Data Files are Online

To verify that the data files are online, in `SQL*Plus`, enter the following commands:

```
SQL> Col file_name format a50
SQL> Col status format a10
SQL> Set line 200
SQL> Select file_id, status, bytes, file_name from dba_data_files order by
SQL> tablespace_name;
```

The status should be `AVAILABLE` for all the data files.

Verify That the Listener is Running

To verify that the listener is running, perform the following steps:

- 1 Check to see if the Oracle listener processes are running by entering the following command:

```
# ps -ef | grep tns
```

```
oracle 11664 1 0 Mar22 ? 00:08:05 /u01/app/oracle/product/
11.2.0/db_2/bin/tnslsnr LISTENER -inherit
oracle 22725 22706 0 22:44 pts/2 00:00:00 grep tns
```

- 2 Check the status of the listener with the `lsnrctl` command:

```
# lsnrctl status
```

The listener should be listening on port 1521 (default), or on the port that you have designated that the Oracle listener process use, with the TCP protocol, and should be handling the instance named `truth`. The `lsnrctl` command should generate output similar to the following lines:

```
. . . .
Connecting to (ADDRESS=(PROTOCOL=tcp)
(HOST=perl.performance.qa.example.com) (PORT=1521))
. . . .
Instance "truth", status READY, has 1 handler(s) for this service...
```

- 3 Test connectivity to the instance from the Data Access Engine (`spin`) and Web Services Data Access Engine (`twist`) hosts by running the `tnsping` utility:

```
# tnsping truth
```

The OK statement displayed by the `tnsping` utility confirms that the listener is up and can connect to the instance. The `tnsping` utility should generate output similar to the following lines:

```
. . . .
Used parameter files:

Used HOSTNAME adapter to resolve the alias
Attempting to contact
(DESCRIPTION=(CONNECT_DATA=(SERVICE_NAME=truth.performance.qa.example.com)
)) (ADDRESS=(PROTOCOL=TCP) (HOST=192.168.165.178) (PORT=1521)))
OK (0 msec)
```

```
Attempting to contact
(DESCRIPTION=(ADDRESS=(HOST=localhost) (PORT=1521) (PROTOCOL=tcp)) (CONNECT_
DATA=(SERVICE_NAME=truth)))
OK (0 msec)
```

As an alternative to running the `tnsping` utility in this step, you can check the connectivity by running `sqlplus` and connecting to the database instance with the service name (TNS alias), for example:

```
# sqlplus myuser/mypass@truth
```

Examine the Log Files

To examine the log files, perform the following steps:

- 1 Look for errors in the `alert.log` file.

For each instance, locate the `alert.log` file in the background dump destination directory:

```
$ORACLE_BASE/diag/rdbms/<SID>/<SID>/trace/
```

This is an example `bdump` directory for an instance with the `truth` SID:

```
/u01/app/oracle/diag/rdbms/truth/truth/trace/
```

- 2 Look for errors in the other log and trace files, located in various directories under:

```
$ORACLE_BASE/diag/rdbms/<SID>/<SID>
```

Check for Sufficient Free Disk Space in the Tablespaces

To check for sufficient disk space, perform the following steps:

- 1 Enter the following commands in `sqlplus`:

```
SQL> set line 200
SQL> column dummy noprint
SQL> column pct_used format 999.9 heading "Pct|Used"
SQL> column name format a16 heading "Tablespace Name"
SQL> column mbytes format 999,999,999 heading "Current|File Size|MB"
SQL> column used format 999,999,999 heading "Used MB "
SQL> column free format 999,999,999 heading "Free MB"
SQL> column largest format 999,999,999 heading "Largest|Contiguous|MB"
SQL> column max_size format 999,999,999 heading "Max Possible|MB"
SQL> column pct_max_used format 999.999 heading "Pct|Max|Used"
SQL> break on report
SQL> compute sum of Mbytes on report
SQL> compute sum of free on report
SQL> compute sum of used on report
```

```
SQL> SELECT
      nvl(df.tablespace_name,'UNKOWN') name, df.mbytes_alloc Mbytes,
      df.mbytes_alloc-nvl(fs.mbytes_free,0) used, nvl(fs.mbytes_free,0)
free,
      ((df.mbytes_alloc-nvl(fs.mbytes_free,0)) / df.mbytes_alloc) * 100
pct_used,
      nvl(df.largest,0) largest, nvl(df.mbytes_max,df.mbytes_alloc)
Max_Size,
      ((df.mbytes_alloc-nvl(fs.mbytes_free,0)) / df.mbytes_max) * 100
pct_max_used
FROM
      ( SELECT tablespace_name, sum(bytes)/1024/1024 Mbytes_alloc,
max(bytes)/1024/1024 largest,
```

```

sum(decode(autoextensible, 'YES', greatest(bytes, maxbytes), bytes)) / 1024 /
1024 Mbytes_max
      FROM
          dba_data_files GROUP BY tablespace_name
    ) df,
    ( SELECT tablespace_name, sum(bytes) / 1024 / 1024 Mbytes_free
      FROM dba_free_space GROUP BY tablespace_name
    ) fs
WHERE
    df.tablespace_name = fs.tablespace_name(+)
UNION
SQL> SELECT
    D.tablespace_name name, D.mbytes_alloc Mbytes, ((ss.used_blocks *
F.block_size) / 1024 / 1024) used,
    D.mbytes_alloc - ((ss.used_blocks * F.block_size) / 1024 / 1024) free,
    ((D.mbytes_alloc - nvl((D.mbytes_alloc - ((ss.used_blocks *
F.block_size) / 1024 / 1024)), 0)) / D.mbytes_alloc) * 100 pct_used,
    nvl((G.max_blocks * F.block_size) / 1024 / 1024), 0) largest,
Max_Mbytes Max_Size,
    ((D.mbytes_alloc - nvl((D.mbytes_alloc - ((ss.used_blocks *
F.block_size) / 1024 / 1024)), 0)) / D.Max_Mbytes) * 100 pct_pct_used
FROM
    ( SELECT tablespace_name, used_blocks, free_blocks, max_size
      FROM v$sort_segment
    ) ss,
    ( SELECT tablespace_name, sum(bytes) / 1024 / 1024 Mbytes_alloc,

sum(decode(autoextensible, 'YES', greatest(bytes, maxbytes), bytes)) / 1024 /
1024 Max_Mbytes
      FROM dba_temp_files GROUP BY tablespace_name
    ) D,
    ( SELECT B.name, C.block_size, SUM (C.bytes) / 1024 / 1024 mb_total
      FROM v$tablespace B, v$tempfile C
      WHERE B.ts# = C.ts# GROUP BY B.name, C.block_size
    ) F,
    ( SELECT B.name, max(blocks) max_blocks, sum(blocks) total_blocks
      FROM v$tablespace B, v$tempfile C
      WHERE B.ts# = C.ts# GROUP BY B.name
    ) G
WHERE ss.tablespace_name = D.tablespace_name and ss.tablespace_name =
F.name and ss.tablespace_name = G.name;

```

In the output generated by the preceding commands, compare the numbers under the Used and Free headings.

- 2 To list the existing data, index, and temporary files, enter the following commands in sqlplus:

```
SQL> Select file_id, bytes, file_name from dba_data_files;
```

- 3 If a tablespace has auto-extended to its maximum size and is running out of disk space, then add new data files by entering the ALTER TABLESPACE command in sqlplus.

The following example commands add data files to four of the tablespaces. For a full list of tablespaces and data files, see the output generated by the commands in the preceding two steps.

```

SQL> ALTER TABLESPACE "AAA_DATA"
SQL> ADD DATAFILE '/u01/oradata/truth/aaa_data10.dbf'
SQL> SIZE 32M AUTOEXTEND ON NEXT 128M MAXSIZE 4000M ;

SQL> ALTER TABLESPACE "AAA_INDX"
SQL> ADD DATAFILE '/u02/oradata/truth/aaa_indx11.dbf'
SQL> SIZE 32M AUTOEXTEND ON NEXT 128M MAXSIZE 4000M ;

SQL> ALTER TABLESPACE "UNDO"
SQL> ADD DATAFILE '/u03/oradata/truth/undo12.dbf' SIZE 32M AUTOEXTEND ON
NEXT 128M MAXSIZE 4000M ;

SQL> ALTER TABLESPACE "TEMP" ADD
SQL> TEMPFILE '/u04/oradata/truth/temp14.dbf' SIZE 32M AUTOEXTEND ON NEXT
128M MAXSIZE 4000M ;

```

Database Statistics Collection



As of SA 10.0, the schema and index statistics collection for SA database users AAA, TRUTH etc. have been moved from `dba_jobs` to Oracle's Automatic Optimizer Statistics Collection. Starting with SA 10 you will no longer see these jobs in `dba_jobs`.

SA relies on Oracle's Automatic Optimizer statistics collection to collect schema statistics on all objects that have stale or no statistics. Gathering frequent database statistics is recommended to avoid database performance degradation.

Enable Oracle Automatic Optimizer Statistics Collection

To verify that the Oracle Automatic optimizer statistics collection is enabled, perform the following steps:

- 1 Enter the following commands in SQL*Plus:

```

# su - oracle
# sqlplus "/" as sysdba"

```

```

SQL> set line 200
SQL> col status format a10
SQL> SELECT status FROM dba_autotask_client where client_name='auto
optimizer stats collection';

```

The output from the above statement should be as follows:

```

STATUS
-----
ENABLED

```

- 2 If the status is not `ENABLED`, execute the following statement to enable Oracle's Automatic Optimizer statistics collection.

```

SQL> EXEC DBMS_AUTO_TASK_ADMIN.ENABLE(client_name => 'auto optimizer stats
collection',operation => NULL, window_name => NULL);

```

Verifying the Interval and Frequency of the Default Windows

The SA-Supplied Oracle RDBMS has increased the frequency and time of the default windows. The windows are now run every 8 hours at 06, 14 and 21 (UTC) hours. By default Oracle runs the windows once a day at 10:00p.m. UTC. SA recommends collecting the statistics more frequently to avoid database performance degradation.

```
SQL> set line 200
SQL> col repeat_interval format a75
col duration format a25
select window_name, repeat_interval, duration from dba_scheduler_windows
order by window_name;
```

WINDOW_NAME	REPEAT_INTERVAL	DURATION
FRIDAY_WINDOW	freq=daily;byday=FRI;byhour=06,14,21;byminute=0;bysecond=0	+000 04:00:00
MONDAY_WINDOW	freq=daily;byday=MON;byhour=06,14,21;byminute=0;bysecond=0	+000 04:00:00
SATURDAY_WINDOW	freq=daily;byday=SAT;byhour=06,14,21;byminute=0;bysecond=0	+000 20:00:00
SUNDAY_WINDOW	freq=daily;byday=SUN;byhour=06,14,21;byminute=0;bysecond=0	+000 20:00:00
THURSDAY_WINDOW	freq=daily;byday=THU;byhour=06,14,21;byminute=0;bysecond=0	+000 04:00:00
TUESDAY_WINDOW	freq=daily;byday=TUE;byhour=06,14,21;byminute=0;bysecond=0	+000 04:00:00
WEDNESDAY_WINDOW	freq=daily;byday=WED;byhour=06,14,21;byminute=0;bysecond=0	+000 04:00:00
WEEKEND_WINDOW	freq=daily;byday=SAT;byhour=14;byminute=0;bysecond=0	+002 00:00:00
WEEKNIGHT_WINDOW	freq=daily;byday=MON,TUE,WED,THU,FRI;byhour=14;byminute=0;bysecond=0	+000 08:00:00

9 rows selected

Modifying the Default STALE_PERCENT Threshold

The value of `STALE_PERCENT` determines the percentage of rows in a table that must change before the statistics on that table are deemed stale and should be re-gathered. The default value is 10%.

You can determine a table's `stale_percent` by running the following SQL:

```
SQL> select DBMS_STATS.GET_PREFS('STALE_PERCENT', '<schema_owner', 'table_name')
table_stale_percnt from dual;
```

For example:

```
# sqlplus "/ as sysdba"
SQL> select DBMS_STATS.GET_PREFS('STALE_PERCENT', 'LCREP', 'DML_VALUES')
table_stale_percnt from dual;
exit;
```

When upgrading or installing SA 10, the installer sets the threshold for some of the larger tables. You can adjust the `stale_percent` threshold value in Oracle 11g by using the SQL:

```
DBMS_STATS.SET_TABLE_PREFS('LCREP', 'DML_VALUES', 'STALE_PERCENT', '5');
```

The threshold for the following `schema_owners,s table_names` is modified:

```
'LCREP', 'DML_VALUES'
'LCREP', 'TRANSACTION_DML'
'LCREP', 'TRANSACTIONS'
'LCREP', 'OMB_MESSAGE'
'TRUTH', 'UNITS'
'TRUTH', 'INSTALLED_UNITS'
'TRUTH', 'COMPLIANCE_DETAIL'
'AAA', 'AUDIT_LOG_ENTRY'
```

Manually Gathering Database Statistics

You can manually gather SA schema statistics by using the following SQL:

```
# Sqlplus "/ as sysdba"
SQL> exec dbms_stats.gather_schema_stats (ownname=>'AAA',options=>'GATHER AUTO');
SQL> exec dbms_stats.gather_schema_stats (ownname=>'TRUTH',options=>'GATHER AUTO');
SQL> exec dbms_stats.gather_schema_stats (ownname=>'LCREP',options=>'GATHER AUTO');
SQL> exec dbms_stats.gather_schema_stats (ownname=>'GCADMIN',options=>'GATHER AUTO');
```

SA Support for dba_scheduler_jobs

Starting with Oracle 10g, the DBA_JOB package was superceded by the improved Oracle Scheduler (dba_scheduler) package. Although Oracle still supports the DBA_JOB package for backward compatibility, Oracle will make no further enhancements to the package. Since the DBA_SCHEDULER provides better functionality, all the SA jobs that used the DBA_JOB package have been redesigned in this SA release to use the DBA_SCHEDULER package. The affected jobs can be found in the dba_scheduler_jobs table. These changes are only relevant to new SA 10.0 Cores and cores upgraded to SA 10.0.

To view the jobs and changes made, you can run the following from SQL*Plus:

```
# Su - oracle
# Sqlplus "/ as sysdba"
SQL> set line 200
SQL> col owner format a14
SQL> col job_action format a50
SQL> col job_name format a50
SQL> select job_name, owner, job_action from dba_scheduler_jobs where owner in
('OPSWARE_ADMIN', 'LCREP', 'GCADMIN');
```

Your output should be as follows:

JOB_NAME	OWNER	JOB_ACTION
WLMPURGE_GC	GCADMIN	WLMPURGE.GC_JOBS
STORAGEINITIATORPURGE_GC	GCADMIN	STORAGEINITIATORPURGE.GC_
		STORAGEINITIATORS
AUDITPURGE_GC	GCADMIN	AUDITPURGE.GC_AUDITLOGS
CHANGELOGPURGE_GC	GCADMIN	CHANGELOGPURGE.GC_
		CHANGELOGS
WAYPURGE_GC	GCADMIN	WAYPURGE.GC_SESSIONS
LCREP_INDEX_STATS	LCREP	gather_lcrep_stats
OPSWARE_ADMIN_SYSTEM_STATS	OPSWARE_ADMIN	gather_opsware_admin_
		sys_stats

7 rows selected.

Running dba_scheduler_jobs manually

If you need to run the System/Index Statistics and the Garbage Collection jobs manually, you must first grant the following privilege.

```
SQL> grant create session to lcrep, gcadmin;
```

To run the statistics collection jobs manually in SQL*Plus, use the commands shown below. If you copy and paste the following command examples, replace the variables like `schema_user_value` with the values of the **schema_user** column displayed by the preceding select statement. Substitute the variables such as `job_name_value` with the values of the job column displayed by the same select statement.

```
SQL> connect <schema_user_value>/<password>
SQL> exec dbms_scheduler.run_job('<job_name_value>');
```

After you are done running the jobs, you should revoke the privileges granted above. Log in to SQL*Plus and enter the following command:

```
SQL> revoke create session from lcrep, gcadmin;
```

Changing the Time Jobs are Run

`dba_scheduler_jobs` are run at UTC time. To change the time at which the jobs are run, follow these instructions:

```
sqlplus "/ as sysdba"
SQL> set line 300
SQL> col job_name format a30
SQL> col owner format a14
SQL> col last format a17
SQL> col next format a17
SQL> col repeat_interval format a40
SQL> col job_action format a30
```

```
SQL> select job_name, owner, to_char(LAST_START_DATE, 'MM/DD/YY HH:MI:SS') last,
to_char(next_run_date, 'MM/DD/YY HH:MI:SS') next, repeat_interval, job_action from
dba_scheduler_jobs where owner in ('OPSWARE_ADMIN', 'LCREP', 'GCADMIN');
```

The above statement provides information about a job. Note the job name and the owner that has the privilege to run this job.

The output of the above statement is similar to the following (formatting is compressed due to space limitations):

JOB_NAME	OWNER	LAST	NEXT	REPEAT_INTERVAL	JOB_ACTION
WLPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	TRUNC(SYSDATE+1)+9/24	WLPURGE_GC_JOBS STORAGE INITIATOR PURGE_GC_ STORAGE INITIATORS
STORAGEINITIATORPURGE_GC	GCADMIN	04/02/12 09:47:30	04/03/12 10:47:30	SYSDATE+1/24	AUDITPURGE_GC_ AUDITLOGS
AUDITPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	TRUNC(SYSDATE+1)+9/24	CHANGELOGPURGE. GC_CHANGELOGS
CHANGELOGPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	TRUNC(SYSDATE+1)+9/24	WAYPURGE_GC_ SESSIONS
WAYPURGE_GC	GCADMIN	04/02/12 09:00:02	04/04/12 09:00:00	TRUNC(SYSDATE+1)+9/24	gather_lcrep_ stats
LCREP_INDEX_STATS	LCREP	04/01/12 11:00:04	04/03/12 11:00:00	TRUNC(SYSDATE+2)+11/24	
OPSWARE_ADMIN_SYSTEM_STATS	OPSWARE_ADMIN	04/02/12 06:00:01	04/03/12 06:00:00	TRUNC(SYSDATE+1) + 18/24 + mod(abs(to_ number(to_char	gather_opsware admin_sys_stats

```
(sysdate + 1, 'D'))  
- 7) + 2, 7)
```

7 rows selected.

In this example the user lcrep changes the time/interval at which the job is run. Any other user can be substituted for the user lcrep.

```
3 sqlplus "connect / as sysdba"  
SQL> grant create session to lcrep;  
Grant succeeded.
```

In the example:

- job name=LCREP_INDEX_STATS
- owner = lcrep

In this example, the job LCREP_INDEX_STATS runs at 11:00 a.m. UTC. To change this to 9:00 a.m. UTC, the command is:

```
SQL> connect lcrep/<password_for_lcrep>  
Connected.
```

```
SQL> exec dbms_scheduler.set_attribute('LCREP_INDEX_STATS',  
attribute=>'REPEAT_INTERVAL', value=>'TRUNC(SYSDATE+2)+9/24');
```

Verifying SA Database Scheduler Jobs

When the Model Repository is installed, the SA Installer sets up the System/Index Statistics and the Garbage Collection jobs in Oracle's dba_scheduler_jobs which then runs these jobs at specified time-intervals. The jobs perform system/ index statistics collection and garbage collection. If the system/index statistics collection jobs do not run successfully, database performance degrades. If the garbage collection jobs do not run, old data accumulates and requires additional disk space. Performance can also be affected.

To verify that the Jobs in DBA_SCHEDULER_JOBS ran successfully, perform the following steps:

- 1 Enter the following commands in SQL*Plus:

```
SQL> set line 200  
SQL> col job_name format a50  
SQL> col owner format a14  
SQL> col last format a17  
SQL> col next format a17  
SQL> col state format a10  
SQL> col job_action format a50
```

```
SQL> select job_name, owner, to_char(LAST_START_DATE, 'MM/DD/YY HH:MI:SS')  
last, to_char(next_run_date, 'MM/DD/YY HH:MI:SS') next, state, job_action  
from dba_scheduler_jobs where owner in ('OPSWARE_ADMIN', 'LCREP', 'GCADMIN');
```

In the output generated from the preceding statement, the value of the JOB_ACTION column indicates the type of job. The jobs owned by GCADMIN perform the garbage collection. The job owned by LCREP performs index statistics collection and the job owned by OPSWARE_ADMIN performs system statistics collection. Sample output looks like this:

JOB_NAME	OWNER	LAST	NEXT	STATE	JOB_ACTION
WLPURGE_GC	GCADMIN	04/03/12 09:00:00	04/04/12 09:00:00	SCHEDULED	WLPURGE.GC_JOBS
STORAGEINITIATORPURGE_GC_PURGE_GC	GCADMIN	04/03/12 09:00:00	04/02/12 09:47:30	SCHEDULED	STORAGEINITIATORPURGE.GC_STORAGEINITIATORS
AUDITPURGE_GC	GCADMIN	04/03/12 09:00:00	04/04/12 09:00:00	SCHEDULED	AUDITPURGE.GC_AUDITLOGS
CHANGELOGPURGE_GC	GCADMIN	04/03/12 09:00:00	04/04/12 09:00:00	SCHEDULED	CHANGELOGPURGE.GC_CHANGELOGS
WAYPURGE_GC	GCADMIN	04/03/12 09:00:00	04/04/12 09:00:00	SCHEDULED	WAYPURGE.GC_SESSIONS
LCREP_INDEX_STATS	LCREP	04/02/12 11:00:00	04/03/12 11:00:00	SCHEDULED	gather_lcrep_stats

```
OPSWARE_ADMIN_SYSTEM OPSWARE_ADMIN 04/02/12 06:00:00 04/03/12 06:00:00 SCHEDULED gather_opsware_admin_sys_stats
_STATS
```

7 rows selected.

where:

JOB_NAME - name of the job

OWNER - the user who with permissions to run the job

LAST_DATE - last date-time when the job was run

NEXT_DATE - next date the job will run

STATE - The status of the scheduled job:

- disabled - The job is disabled
- scheduled - The job is scheduled to be executed
- running - The job is currently running
- completed - The job has completed, and is not scheduled to run again
- broken - The job is broken
- failed - The job was scheduled to run once and failed
- retry scheduled - The job has failed at least once and a retry has been scheduled to be executed
- succeeded - The job was scheduled to run once and completed successfully
- **JOB_ACTION** - the procedure that the job runs

Monitoring Database Users

To monitor database users, perform the following steps:

- 1 To check the database users, enter the following command in sqlplus:

```
SQL> Select username, account_status, default_tablespace,
temporary_tablespace from dba_users;
```

Monitoring the ERROR_INTERNAL_MSG Table

Various SA internal PL/SQL procedures write exceptions to the `truth.ERROR_INTERNAL_MSG` table. You should monitor this table for errors (daily checks are recommended) on all Model Repository (Oracle) databases.

Executing the SQL below lists the data in `error_internal_msg` from the last fifteen days.



You can remove the `WHERE` clause if you want to display all data in the `truth.ERROR_INTERNAL_MSG` table.

```
# Su - oracle
# Sqlplus "/" as sysdba"
SQL> set line 200
```

```

SQL> col ERR_ID format 999999
SQL> col ERR_USER format a8
SQL> col ERR_TABLE format a25
SQL> col ERR_TABLE_PK_ID format a10
SQL> col ERR_CODE format 9999999
SQL> col ERR_TEXT format a20
SQL> col ERR_INFO format a30

SQL> select ERROR_INTERNAL_MSG_ID ERR_ID,
ERR_DATE,
ERR_USER,
ERR_TABLE,
ERR_TABLE_PK_ID,
ERR_CODE,
ERR_TEXT,
DELETE_FLG,
ERR_INFO
from ERROR_INTERNAL_MSG
where ERR_DATE > sysdate - 15
order by ERR_DATE;

```

Rebuilding the SHADOW_FOLDER_UNIT Table

The procedure `SHADOW_FOLDER_UNIT_RELOAD` is provided in case the contents of `SHADOW_FOLDER_UNIT` table becomes out of synchronization or there are multiple records of the type (`shadow_folder_unit.folder_id = -1`).

The table can be rebuilt without stopping the system. Simply connect as user `TRUTH`, `TWIST`, `SPIN`, or `OPSWARE_ADMIN` and issue the command:

```
SQL> exec SHADOW_FOLDER_UNIT_UTIL.SHADOW_FOLDER_UNIT_RELOAD
```

Check the results from monitoring the `ERROR_INTERNAL_MSG` table. If the results contain:

```
'ERR_TABLE' = 'UNIT_RELATIONSHIPS'
```

do the following:

- 1 Check if there are records in `truth.SHADOW_FOLDER_UNIT` of the type (`folder_id = -1`).

```
# connect / as sysdba
SQL> select count(*) from shadow_folder_unit where folder_id = -1;
```
- 2 If the above SQL returns more than zero rows, then run the following during low database usage time:

```
SQL> grant create session to truth;
SQL> connect truth/<password>
SQL> exec SHADOW_FOLDER_UNIT_UTIL.SHADOW_FOLDER_UNIT_RELOAD;
```
- 3 Run the SQL from [Monitoring the ERROR_INTERNAL_MSG Table](#) on page 258 and check if the procedure has listed any faulty records. `SHADOW_FOLDER_UNIT_UTIL.SHADOW_FOLDER_UNIT_RELOAD` is idem potent therefore the faulty records can be fixed and you can rerun `SHADOW_FOLDER_UNIT_UTIL.SHADOW_FOLDER_UNIT_RELOAD`.

HP recommends that you gather table statistics after the data reload:

```
SQL> connect truth/<password>
SQL> exec dbms_stats.gather_table_stats (
           ownname=> 'TRUTH',
           tabname=> 'SHADOW_FOLDER_UNIT',
           estimate_percent=> DBMS_STATS.AUTO_SAMPLE_SIZE,
           cascade => true);
```

4 Revoke the permissions given to user truth:

```
SQL> connect / as sysdba
SQL> revoke create session to truth;
```

Oracle Database Backup Methods

It is important that you back up the database on a regular basis. Be sure to use more than one backup method and to test your recovery process.

You can use the following methods to back up the Oracle database:

- **Export-Import:** An export extracts logical definitions and data from the database and writes the information to a file. Export-import does not support point-in-time recoveries. Do not use Export-Import as your only backup and recovery strategy.

See the information on the `Export-Import` subdirectory in [Oracle/SA Installation Scripts, SQL Scripts, and Configuration Files](#) on page 233.

- **Cold or Off-Line Backups:** This procedure shuts the database down and backs up all data, index, log, and control files. Cold or off-line backups do not support point-in-time recoveries.
- **Hot or Online Backups:** During these backups, the database must be available and in ARCHIVELOG mode. The tablespaces are set to backup mode. This procedure backs up tablespace files, control files, and archived redo log files. Hot or online backups support point-in-time recoveries.
- **RMAN Backups:** While the database is either off-line or on-line, use the `rman` utility to back up the database.

Regardless of your backup strategy, remember to back up all required Oracle software libraries, parameter files, password files, and so forth. If your database is in ARCHIVELOG mode, you also need to back up the archived log files.

For more information on backing up Oracle databases, see the following Oracle documents:

- *Oracle Database 2 Day DBA*
- *Oracle Database Concepts*
- *Oracle Database Administrator's Guide*

These guides are on the Oracle web site at the following URL:

<http://www.oracle.com/technology/documentation/index.html>

Troubleshooting System Diagnosis Errors

If an additional privilege (permission) has been made manually to the database, when SA performs a system diagnosis on the Data Access Engine, an error message might be generated. For example, if an additional grant has been made to the `truth.facilities` table, the following error appears:

```
Test Information
Test Name: Model Repository Schema
Description: Verifies that the Data Access Engine's version of the schema
matches the Model Repository's version.
Component device: Data Access Engine (spin.blue.qa.example.com)
Test Results: The following tables differ between the Data Access Engine
and the Model Repository: facilities.
```

To fix this problem, revoke the grant. For example, if you need to revoke a grant on the `truth.facilities` table, log on to the server with the database and enter the following commands:

```
# su - oracle
# sqlplus "/" as sysdba"
SQL> grant create session to truth;
SQL> connect truth/<truth passwd>;
SQL> revoke select on truth.facilities from spin;
SQL> exit
# sqlplus "/" as sysdba"
SQL> revoke create session from truth;
```

Useful SQL

The following SQL commands help you manage information in the Oracle database.

BIN\$ Objects

If the SA Installer discovers the existence of BIN\$ objects in the database, enter the following SQL commands:

```
SQL> show parameter recyclebin;
SQL> SELECT owner,original_name,operation,type FROM dba_recyclebin;
connect <owner>/password purge recyclebin; or purge table BIN$xxx;
```

By default, `recyclebin` is set to OFF.

