Avoid the confusion surrounding Automatic Memory Management

In Oracle Database 11g, the Automatic Memory Management feature requires more shared memory (/dev/shm) and file descriptors.
Metalink Note 169706.1 states that shared memory should be sized to be at least the greater of the database initialisation parameters; MEMORY_MAX_TARGET and MEMORY_TARGET for each Oracle instance on the server.

However, I have discovered in most cases, that the shared memory needs to be at least 2G greater than MEMORY_TARGET else the dynamic memory allocation becomes inappropriate and can crash the server.

When setting the MEMORY_TARGET parameter, the Oracle RDBMS will dynamically size the SGA and PGA and write "underscore-underscore" parameters to the spfile. This is done to protect the database initialisation memory settings following an instance crash.

On Linux, if the shared memory is not sized (and mounted) to a value greater than MEMORY_TARGET you will receive the following error when attempting to start the database instance.

ORA-00845: MEMORY_TARGET not supported on this system

Below are some examples of the dynamic settings made by the Oracle RDBMS when setting the MEMORY_TARGET parameter.

(Note the small SGA and huge PGA size when MEMORY_TARGET 1GB less than shared memory size)

With 20G MEMORY_TARGET and 25G shared memory (stable)

```
# Oracle init.ora parameter file generated by instance ORCL1 on
08/26/2009 11:39:02
*.__db_cache_size=8G
*.__java_pool_size=256M
*.__large_pool_size=256M
*.__pga_aggregate_target=8G
*.__sga_target=12G
*.__shared_io_pool_size=0
*.__shared_pool_size=2816M
*.__streams_pool_size=512M
```

With 24G MEMORY_TARGET and 25G shared memory (unstable)

```
# Oracle init.ora parameter file generated by instance ORCL1 on
08/17/2009 14:30:08
ORCL1.__db_cache_size=512M
```

```
ORCL2.__db_cache_size=1G
*.__java_pool_size=256M
*.__large_pool_size=256M
*.__pga_aggregate_target=20736M
*.__sga_target=3840M
*.__shared_io_pool_size=0
ORCL1.__shared_pool_size=2G
ORCL2.__shared_pool_size=2304M
*.__streams_pool_size=512M
```

Configuring Shared Memory

The shared memory can be configured using a swap partition or swap file. Best practice is to create an appropriately sized swap partition. However, if the existing swap partition exists and is too small, you can add a swap file by following the example below (as root user):

1. To determine the amount of shared memory available, enter the following command:

```
df -h /dev/shm/
Filesystem            Size  Used Avail Use% Mounted on
shmfs                  8G     0   8G   0% /dev/shm
```

Create a 16G swap file:

```
dd if=/dev/zero of=/swapfile bs=1024 count=16777216
```

2. Setup the swap file with the command:

```
mkswap /swapfile
```

3. To enable the swap file immediately but not automatically at boot time:

```
swapon /swapfile
```

4. To enable it at boot time, edit /etc/fstab to include:

```
/swapfile   swap  swap  defaults    0 0
```

5. To view swap space:

```
cat /proc/swaps
Filename                Type            Size    Used    Priority
/dev/cciss/c0d0p5       partition       8193108 0       -1
/swapfile               file            16777208        0       -2
```

vi /etc/fstab to set shmfs size at bootup:

```
shmfs       /dev/shm    tmpfs size=24g    0
```

6. To re-read /etc/fstab and set shmfs size on the fly:

```
mount -t tmpfs shmfs -o size=24g /dev/shm
```

7. View /etc/fstab

```
cat /etc/fstab
LABEL=/                 /          ext3    defaults        1 1
LABEL=/u02              /u02       ext3    defaults        1 2
LABEL=/u01              /u01       ext3    defaults        1 2
LABEL=/boot             /boot      ext3    defaults        1 2
tmpfs                   /dev/shm   tmpfs   defaults        0 0
shmfs                   /dev/shm   tmpfs   size=24g        0
devpts                  /dev/pts   devpts  gid=5,mode=620  0 0
sysfs                   /sys       sysfs   defaults        0 0
proc                    /proc      proc    defaults        0 0
LABEL=SW-cciss/c0d0p5   swap       swap    defaults        0 0
/swapfile               swap       swap    defaults        0 0
```

N.B.
MEMORY_MAX_TARGET and MEMORY_TARGET cannot be used when LOCK_SGA is enabled or with huge pages on Linux

Conclusion

From experience I suggest sizing the shared memory partition (/dev/shm) to at least 2 GB greater than the MEMORY_TARGET or MEMORY_MAX_TARGET for your database instance. This will enable Oracle to grow instance memory based on workload.

This is particularly important when using Automatic Storage Management. Oracle will set the MEMORY_TARGET parameter by default for an ASM instance, which will of course also use the available shared memory on the server.

More Gotchas

Setting tmpfs size

Another Oracle 11g gotcha to be aware of is; receiving the ORA-00845 error even though you have created a swap partition/file greater than your database MEMORY_TARGET. This is because the default size of tmpfs on Linux (2.4 kernel and above) is approximately half the physical system memory.

This limitation can be dynamically changed by editing the `/etc/fstab` file, removing the "defaults" attribute and specifying a maximum size for tmpfs as shown in example below:

```
tmpfs                   /dev/shm   tmpfs   size=24g        0 0
```

Hitting Bug 7272646

If your RDBMS Oracle Home has newly been upgraded from 11.1.0.6 to 11.1.0.7 with no additional interim patches applied, you are likely to hit the following bug if the MEMORY_TARGET database initialisation parameter exceeds 3G in size.

BUG 7272646 - ORA-27103 WHEN MEMORY_TARGET > 3G

The bug causes the Oracle instance to crash when opening the database. The following errors are seen in the instance's alert log:

```
ORA-27103: internal error
Linux-x86_64 Error: 2: No such file or directory
Additional information: -1
Additional information: 1
MMAN (ospid: 12211): terminating the instance due to error 27103
Tue May 25 12:46:46 2010
ORA-1092 : opitsk aborting process
Tue May 25 12:46:46 2010
System state dump is made for local instance
System State dumped to trace file
/u01/app/oracle/diag/rdbms/orcl/ORCL1/trace/ORCL1_diag_12185.trc
Trace dumping is performing id=[cdmp_20100525124646]
Instance terminated by MMAN, pid = 12211
```

To resolve this issue:

1. Download patch 7272646 from Oracle Metalink
2. Shutdown the database instance
3. Apply the patch as described in the README.txt file
4. Startup the Oracle instance