

Deconstructing Startup

a.k.a

Probably Useless, Maybe Interesting

Neil Johnson

o·su·mo

Who Am I

Neil Johnson

o·su·mo

Independent Oracle DBA



@neiljdba



neil@osumo.co.uk



<http://oraganism.wordpress.com/>



@yodbuk

Deconstructing Startup

a.k.a


Probably Useless, Maybe Interesting

Neil Johnson

o·su·mo

Sysadmin

I, UM, MESSED UP
MY SERVER AGAIN.



DBA Skill Set

- *shutdown*
- *startup*
- *SELECT * FROM database;*

What The Docs Say...

When Oracle Database starts an instance, it performs the following basic steps:

- Searches for a server parameter file in a platform-specific default location and, if not found, for a text initialization parameter file
- Reads the parameter file to determine the values of initialization parameters
- Allocates the SGA based on the initialization parameter settings
- Starts the Oracle background processes
- Opens the alert log and trace files and writes all explicit parameter settings to the alert log in valid parameter syntax

The Plan

- Sounds pretty boring
- But there are some interesting techniques employed behind the scenes
- Plan to track the process on Oracle Linux using OS tools
- Test system:
 - Oracle Database 12.1.0.2.0 - 64bit
 - Oracle Linux Server release 6.6
 - No ASM, no RAC, no PDB, nomount – keep it simple
 - `THREADED_EXECUTION = FALSE`
- Might be interesting, may or may not be useful

Let's Start By Connecting

conn / as sysdba

```
$ id
uid=54322(neil) gid=54322(dba) groups=54322(dba)
$ echo $$
7712
```

```
$ sqlplus /nolog
SQL*Plus: Release 12.1.0.2.0 Production
on Mon Dec 29 16:49:12 2014

Copyright (c) 1982, 2014, Oracle. All
rights reserved.

SQL>
```

```
$ pstree -p 7712
bash(7712)---sqlplus(7805)

$ ps -fp 7712 -p 7805
UID          PID    PPID  C  STIME TTY          TIME CMD
neil         7712   7699  0  21:07 pts/1        00:00:00 -bash
neil         7805   7712  0  21:09 pts/1        00:00:00 sqlplus
```


conn / as sysdba

```
SQL> conn / as sysdba
Connected to an idle instance.
SQL>
```

```
$ pstree -p 7712
bash(7712) ---sqlplus(7805) ---oracle(7816)
```

```
$ ps -fp 7712 -p 7805 -p 7816
UID          PID    PPID    C  STIME TTY          TIME CMD
neil         7712   7699    0  21:07 pts/1        00:00:00 -bash
neil         7805   7712    0  21:09 pts/1        00:00:00 sqlplus
oracle       7816   7805    0  21:10 ?            00:00:00 oracleSLOB
```

conn / as sysdba

```
$ echo $$  
7712  
$ sqlplus /nolog
```

```
# id  
uid=0(root) gid=0(root) groups=0(root)  
  
# pstree -p 7712  
bash(7712)---sqlplus(7822)  
  
# ps -fp 7712 -p 7822  
UID          PID    PPID    C  STIME TTY          TIME CMD  
neil         7712   7699    0  21:07 pts/1        00:00:00 -bash  
neil         7822   7712    0  21:11 pts/1        00:00:00 sqlplus  
  
# strace -o /tmp/sqlplus-idle.trc -f -p 7822
```

```
SQL> conn / as sysdba  
Connected to an idle instance.
```

```
Process 7822 attached - interrupt to quit  
Process 7853 attached
```

```
SQL> exit  
Disconnected
```

```
Process 7822 detached  
Process 7853 detached
```

conn / as sysdba

```
$ ls -l /tmp/sqlplus-idle.trc  
-rw-r--r--  1 root root 10244871 Nov 13 22:03 /tmp/sqlplus-idle.trc
```

```
...  
7822 read(0, "conn / as sysdba\n", 1024) = 17  
...
```

```
neil          7822  7712  0 21:11 pts/1      00:00:00 sqlplus
```

```
...  
7822 clone(child_stack=0, flags=CLONE_CHILD_CLEARTID|  
CLONE_CHILD_SETTID|SIGCHLD, child_tidptr=0x7f5f0dd04890) = 7853  
...
```

```
7853 execve("/u01/app/oracle/product/12.1.0/dbhome_1/bin/oracle",  
["oracleSLOB", "(DESCRIPTION=(LOCAL=YES) (ADDRESS"...],  
[/* 28 vars */]) = 0  
...
```

conn / as sysdba

```
$ cd /u01/app/oracle/product/12.1.0/dbhome_1/bin  
$ ls -l oracle  
-rwsr-s--x. 1 oracle oinstall 323762270 Oct 30 20:52 oracle
```



SETUID flags

Allow a process to effectively “become” a different user/group

```
$ ps -p 2252 -p 2253 -o pid,ruser,rgroup,user,group,cmd  
PID RUSER RGROUP USER GROUP CMD  
2252 neil dba neil dba sqlplus  
2253 neil dba oracle oinstall oracleSLOB (DESCRIPTION=...
```

demo1.sh

Oracle Executable Ownership

While talking about “oracle” binary ownership...



Previous output without ASM

```
$ cd $ORACLE_HOME/bin
$ ls -l oracle
-rwsr-s--x. 1 oracle oinstall 323762270 Oct 30 20:52 oracle
```

What about with ASM

```
$ cd $ORACLE_HOME/bin
$ ls -l oracle
-rwsr-s--x 1 oracle asmadmin 294728401 Oct 4 2013 oracle
```

Oracle Executable Ownership

- If GI running under a “grid” user - this is how “oracle” can read ASM disks

```
$ cd $ORACLE_HOME/bin
$ ls -l oracle
-rwsr-s--x 1 oracle asmadmin 294728401 Oct  4 2013 oracle

$ ls -l /dev/asm*
brw-rw---- 1 grid asmadmin 8, 17 Mar 25 10:40 /dev/asm1
brw-rw---- 1 grid asmadmin 8, 33 Mar 25 10:40 /dev/asm2
brw-rw---- 1 grid asmadmin 8, 49 Mar 25 10:40 /dev/asm3
brw-rw---- 1 grid asmadmin 8, 65 Mar 25 10:40 /dev/asm4
brw-rw---- 1 grid asmadmin 8, 81 Mar 25 10:40 /dev/asm5
brw-rw---- 1 grid asmadmin 8, 97 Mar 25 10:40 /dev/asm6
brw-rw---- 1 grid asmadmin 8, 113 Mar 25 10:40 /dev/asm7
brw-rw---- 1 grid asmadmin 8, 129 Mar 25 10:40 /dev/asm8
brw-rw---- 1 grid asmadmin 8, 145 Mar 25 10:36 /dev/asm9
```

- Isn't set automatically at software install time
- It is “initially” set when you create a database with DBCA and if using ASM
- A re-link will reset it to “oracle:oinstall”, as will a home clone

Oracle Executable Ownership

Sample DBCA script

```
$ cd /u01/app/oracle/admin/cdb/scripts

$ cat cdb_1.sql
set verify off
ACCEPT sysPassword CHAR PROMPT 'Enter new password for SYS: ' HIDE
ACCEPT systemPassword CHAR PROMPT 'Enter new password for SYSTEM: ' H
host /u01/app/oracle/product/12.1.0/dbhome_1/bin/orapwd file=+DATA/c
host /u01/app/12.1.0/grid_1/bin/setasmgidwrap o=/u01/app/oracle/produ
host /u01/app/oracle/product/12.1.0/dbhome_1/bin/srvctl add srvpool
host /u01/app/oracle/product/12.1.0/dbhome_1/bin/srvctl add database
host /u01/app/oracle/product/12.1.0/dbhome_1/bin/srvctl disable datab
@/u01/app/oracle/admin/cdb/scripts/CreateDB.sql
@/u01/app/oracle/admin/cdb/scripts/CreateDBFiles.sql
@/u01/app/oracle/admin/cdb/scripts/CreateDBCatalog.sql
@/u01/app/oracle/admin/cdb/scripts/JServer.sql
...
```

```
host /u01/app/12.1.0/grid_1/bin/setasmgidwrap \  
o=/u01/app/oracle/product/12.1.0/dbhome_1/bin/oracle
```


Oracle Executable Ownership

A re-link will reset “oracle” binary ownership

```
$ cd $ORACLE_HOME/bin
$ ls -l oracle
-rwsr-s--x 1 oracle asmadmin 294728401 Oct  4 2013 oracle

$ relink all
writing relink log to:
/u01/app/oracle/product/12.1.0/dbhome_1/install/relink.log

$ ls -l oracle
-rwsr-s--x 1 oracle oinstall 294728401 Mar 24 22:40 oracle
```

CRS will run “setasmgidwrap” when DB started using “srvctl”

```
$ srvctl start database -d cdb

$ ls -l oracle
-rwsr-s--x 1 oracle asmadmin 294728401 Mar 24 22:40 oracle
```

Oracle Executable Ownership

Change made by CRS “oraagent” process

Extract from “agent/crsd/oraagent_oracle/oraagent_oracle.log”...

```
... [start] crsHome = /u01/app/12.1.0/grid_1
... [start] oracleHome = /u01/app/oracle/product/12.1.0/dbhome_1
... [start] command = '/u01/app/12.1.0/grid_1/bin/setasmgidwrap oracle_b
... [start] start dependency = hard(ora.DATA.dg,ora.FRA.dg) weak(type:or
... [start] ASM disk group dependency found
... [start] Utils:execCmd action = 1 flags = 6 ohome = /u01/app/12.1.0/g
```

● This means care needs to be taken when

- Not using DBCA (the norm?)
- Cloning homes
- Installing standby database servers

● Otherwise you could start DB via SQL*Plus, add to CRS and get...

```
ORA-27300: OS system dependent operation:invalid_egid failed with status: 1
ORA-27301: OS failure message: Operation not permitted
ORA-27302: failure occurred at: skgpwinit6
ORA-27303: additional information: startup egid = 1000 (oinstall),
                                     current egid = 1020 (asmadmin)
```

● Blog post:

<http://oraganism.wordpress.com/2010/02/23/setasmgidwrap-an-illegitimate-google-whack/>



Saw earlier with “ps”

```
$ ps -fp 7712 -p 7805 -p 7816
UID          PID    PPID  C  STIME TTY          TIME CMD
neil         7712   7699  0   21:07 pts/1        00:00:00 -bash
neil         7805   7712  0   21:09 pts/1        00:00:00 sqlplus
oracle       7816   7805  0   21:10 ?           00:00:00 oracleSLOB
```

Saw earlier with “strace”

```
...
7853  execve("/u01/app/oracle/product/12.1.0/dbhome_1/bin/oracle",
...

```

“oracleSLOB” <> “oracle”

conn / as sysdba

Another trick

```
$ man execve
```

```
int execve(const char *filename, char *const argv[],  
           char *const envp[]);
```

```
...  
7853  execve('/u01/app/oracle/product/12.1.0/dbhome_1/bin/oracle',  
         ['oracleSLOB', "(DESCRIPTION=(LOCAL=YES)(ADDRESS"...],  
         [/* 28 vars */) = 0  
...
```

```
$ man execve
```

```
argv is an array of argument strings passed to  
the new program.
```

Argv[0] = program name = "oracleSLOB"

Argv[1] = first parameter to program = "(DESCRIPTION=(LOCAL=YES)(ADDRESS"

demo2.sh

conn / as sysdba

What else can we see in strace output...

access arch bind brk connect create epoll
fcntl fstat futex getcwd getdents
geteuid getppid getrlimit getuid ioctl list
lseek lstat mempolicy mmap
mprotect munmap open
pipe poll read readlink recvfrom recvmsg
robust rt sched sendto setrlimit setsid shmat
shmctl shmdt shmget
sigaction sigprocmask
socket stat statfs times umask
uname write

Eventually audit file created

```
7853 stat("/u01/app/oracle/admin/SLOB/adump", ...) = -1 ENOENT (No such
file or directory)
7853 open("/u01/app/oracle/product/12.1.0/dbhome_1/rdbms/audit/
SLOB_ora_7853_20141102211327049759143795.aud", ...) = 12
7853 write(12, "Audit file ", 11) = 11
...
7853 write(12, "Unix process pid: 7853, image: ", 31) = 31
```

If **no instance present** to query then check Oracle Base (from orabasetab) and try that

```
$ cat /u01/app/oracle/product/12.1.0/dbhome_1/install/orabasetab
#orabasetab file is used to track Oracle Home associated with Oracle Base
/u01/app/oracle/product/12.1.0/dbhome_1:/u01/app/oracle:OraDB12Home1:N:
```

If `{orabase}/admin/${ORACLE_SID}/adump` exists then use that for “.aud” files

If above does not exist then fall back on `${ORACLE_HOME}/rdbms/audit`

Finally control returns to SQL*Plus

```
7822 write(1, "Connected to an idle instance.\n", 31) = 31
7822 stat("login.sql", 0x7fff21a205c0) = -1 ENOENT (No such file or direct
7822 stat("/u01/app/oracle/product/12.1.0/dbhome_1/sqlplus/admin/glogin.sc
7822 access("/u01/app/oracle/product/12.1.0/dbhome_1/sqlplus/admin/glogin.
7822 statfs("/u01/app/oracle/product/12.1.0/dbhome_1/sqlplus/admin/glogin.
7822 open("/u01/app/oracle/product/12.1.0/dbhome_1/sqlplus/admin/glogin.sc
7822 read(10, "--\n-- Copyright (c) 1988, 2005, "...", 57344) = 342
7822 read(10, "", 57344) = 0
7822 close(10) = 0
7822 write(1, "SQL> ", 5) = 5
7822 read(0, <unfinished ...>
```

Action Ticklist

When Oracle Database starts an instance, it performs the following basic steps:

- Searches for a server parameter file in a predefined default location and, if not found, for a local parameter file
- Reads the parameter file to determine the values of initialization parameters
- Allocates the SGA and sets the initialization parameter settings
- Starts the background processes
- Opens the alert log and trace files and writes all explicit parameter settings to the alert log in valid parameter syntax

Erm... 25 slides in and haven't reached the tick list yet!

Action Ticklist

- Connect as sysdba
- Searches for a server parameter file in a platform-specific default location and, if not found, for a text initialization parameter file
- Reads the parameter file to determine the values of initialization parameters
- Allocates the SGA based on the initialization parameter settings
- Starts the Oracle background processes
- Opens the alert log and trace files and writes all explicit parameter settings to the alert log in valid parameter syntax

startup nomount

Same setup as before to start tracing and connect as sysdba

```
$ echo $$  
7712  
$ sqlplus /nolog
```

```
# pstree -p 7712  
bash(7712)---sqlplus(3873)  
  
# strace -o /tmp/sqlplus-startup-nomount.trc -f -p 3873  
  
Process 3873 attached - interrupt to quit
```

```
SQL> conn / as sysdba  
Connected to an idle instance.
```

```
Process 3882 attached
```

```
# pstree -p 7712  
bash(7712)---sqlplus(3873)---oracle(3882)
```

startup nomount

```
SQL> startup nomount
```

Process 3873 attached - interrupt to quit

Process 3882 attached	Process 3904 detached	Process 3919 attached	Process 3933 attached
Process 3893 attached	Process 3906 attached	Process 3918 detached	Process 3932 detached
Process 3893 detached	Process 3907 attached	Process 3920 attached	Process 3934 attached
Process 3894 attached	Process 3906 detached	Process 3921 attached	Process 3882 detached
Process 3895 attached	Process 3908 attached	Process 3920 detached	Process 3935 attached
Process 3894 detached	Process 3909 attached	Process 3922 attached	Process 3936 attached
Process 3896 attached	Process 3908 detached	Process 3923 attached	Process 3935 detached
Process 3897 attached	Process 3910 attached	Process 3922 detached	Process 3937 attached
Process 3896 detached	Process 3911 attached	Process 3924 attached	Process 3938 attached
Process 3898 attached	Process 3910 detached	Process 3925 attached	Process 3937 detached
Process 3899 attached	Process 3912 attached	Process 3924 detached	Process 3936 detached
Process 3898 detached	Process 3913 attached	Process 3926 attached	Process 3939 attached
Process 3900 attached	Process 3912 detached	Process 3927 attached	Process 3940 attached
Process 3901 attached	Process 3914 attached	Process 3926 detached	Process 3939 detached
Process 3900 detached	Process 3915 attached	Process 3928 attached	Process 3938 detached
Process 3901 detached	Process 3914 detached	Process 3929 attached	Process 3941 attached
Process 3902 attached	Process 3907 detached	Process 3928 detached	Process 3942 attached
Process 3903 attached	Process 3916 attached	Process 3930 attached	Process 3941 detached
Process 3902 detached	Process 3917 attached	Process 3931 attached	Process 3940 detached
Process 3904 attached	Process 3916 detached	Process 3930 detached	Process 3942 detached
Process 3905 attached	Process 3918 attached	Process 3932 attached	

startup nomount

52 * Process *nnnn* attached

32 * Process *nnnn* detached

```
UID      PID  PPID  C  STIME TTY          TIME CMD
neil     3873 7712  0 12:59 pts/1      00:00:00 sqlplus
oracle   3895   1    0 13:00 ?          00:00:00 ora_pmon_SLOB
oracle   3897   1    0 13:00 ?          00:00:00 ora_psp0_SLOB
oracle   3899   1   10 13:00 ?          00:00:22 ora_vktm_SLOB
oracle   3903   1    0 13:00 ?          00:00:00 ora_gen0_SLOB
oracle   3905   1    0 13:00 ?          00:00:00 ora_mman_SLOB
oracle   3909   1    0 13:00 ?          00:00:00 ora_diag_SLOB
oracle   3911   1    0 13:00 ?          00:00:00 ora_dbrm_SLOB
oracle   3913   1    0 13:00 ?          00:00:00 ora_vkrm_SLOB
oracle   3915   1    0 13:00 ?          00:00:00 ora_dia0_SLOB
oracle   3917   1    0 13:00 ?          00:00:00 ora_dbw0_SLOB
oracle   3919   1    0 13:00 ?          00:00:00 ora_lgwr_SLOB
oracle   3921   1    0 13:00 ?          00:00:00 ora_ckpt_SLOB
oracle   3923   1    0 13:00 ?          00:00:00 ora_smon_SLOB
oracle   3925   1    0 13:00 ?          00:00:00 ora_reco_SLOB
oracle   3927   1    0 13:00 ?          00:00:00 ora_lreg_SLOB
oracle   3929   1    0 13:00 ?          00:00:00 ora_pxmn_SLOB
oracle   3931   1    0 13:00 ?          00:00:00 ora_mmon_SLOB
oracle   3933   1    0 13:00 ?          00:00:00 ora_mmln_SLOB
oracle   3934  3873  0 13:00 ?          00:00:00 oracleSLOB (DESCRIPTION=
```

Open Parameter File

In STRACE output - /tmp/sqlplus-startup-nomount.trc

Shadow process (PID 3882) opens the parameter file

```
3882  open("/u01/app/oracle/product/12.1.0/dbhome_1/dbs/spfileSLOB.ora",
        O_RDWR|O_DSYNC) = 14
3882  fcntl(14, F_SETFD, FD_CLOEXEC) = 0
3882  getrlimit(RLIMIT_NOFILE, {rlim_cur=4*1024, rlim_max=4*1024}) = 0
3882  fcntl(14, F_DUPFD, 256) = 256
3882  fcntl(256, F_SETFD, FD_CLOEXEC) = 0
3882  close(14) = 0
3882  fcntl(256, F_GETFL) = 0x9002 (flags
        O_RDWR|O_DSYNC|O_LARGEFILE)
3882  fcntl(256, F_SETLK, {type=F_RDLCK, whence=SEEK_SET, start=0,
        len=0}) = 0
3882  pread(256, "\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\0\5"...
        , 512, 0) = 512
3882  pread(256, "SLOB. data transfer cache size="..., 1024, 512) = 1024
3882  fcntl(256, F_SETLK, {type=F_UNLCK, whence=SEEK_SET, start=0, len=0})
3882  close(256) = 0
```


Action Ticklist

- Connect as sysdba
- Searches for a server parameter file in a platform-specific default location and, if not found, for a text initialization parameter file
- Reads the parameter file to determine the values of initialization parameters
- Allocates the SGA based on the initialization parameter settings
- Starts the Oracle background processes
- Opens the alert log and trace files and writes all explicit parameter settings to the alert log in valid parameter syntax

Allocate SGA

Allocate SGA

Total System Global Area	369098752	bytes
Fixed Size	2924544	bytes
Variable Size	322961408	bytes
Database Buffers	37748736	bytes
Redo Buffers	5464064	bytes

View shared memory segments

```
# ipcs -m

----- Shared Memory Segments -----
key          shmid      owner      perms      bytes      nattch
0x00000000   8880128    oracle     640        2924544    38
0x00000000   8912897    oracle     640        360710144  19
0x00000000   8945666    oracle     640        5464064    19
0x12e16844   8978435    oracle     640        8192       19
```

Allocate SGA

```
# ipcs -mp

----- Shared Memory Creator/Last-op -----
shmid      owner      cpid      lpid
8880128    oracle     3882      3942
8912897    oracle     3882      3942
8945666    oracle     3882      3942
8978435    oracle     3882      3942
```

```
3873 sqlplus
-- 3882 (oracleSLOB)
```

Allocate SGA

```
# ipcs -mp

----- Shared Memory Creator/Last-op -----
shmid      owner      cpid      lpid
8880128    oracle     3882     3942
8912897    oracle     3882     3942
8945666    oracle     3882     3942
8978435    oracle     3882     3942
```

STRACE for PID 3882

```
3882 shmget(IPC_PRIVATE, 5464064, IPC_CREAT|IPC_EXCL|0640) = 8945666
3882 shmat(8945666, 0x75c00000, 0) = ?
3882 times(NULL) = 431796997
3882 write(4, " Shared memory segment allocated"..., 105) = 105
3882 write(4, "\n", 1) = 1
```

SLOB_ora_3882.trc

Shared memory segment allocated: shmid: 8945666 size: 5464064 bytes

Allocate SGA

Previous test using ASMM

NAME	VALUE
memory_max_target	big integer 0
memory_target	big integer 0
sga_max_size	big integer 352M
sga_target	big integer 352M
pga_aggregate_target	big integer 100M

What about AMM?

NAME	VALUE
memory_max_target	big integer 452M
memory_target	big integer 452M
sga_max_size	big integer 352M
sga_target	big integer 352M
pga_aggregate_target	big integer 100M


Allocate SGA

```
Total System Global Area 473956352 bytes
Fixed Size                 2925744 bytes
Variable Size             427821904 bytes
Database Buffers         37748736 bytes
Redo Buffers              5459968 bytes
```

View shared memory segments with AMM

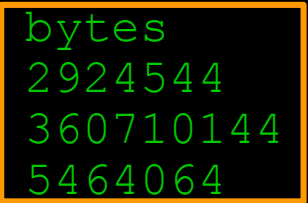
```
$ ipcs -m
```

```
----- Shared Memory Segments -----
key          shmid      owner      perms      bytes      nattch
0x00000000   2359296    oracle     640        4096       0
0x00000000   2392065    oracle     640        4096       0
0x00000000   2424834    oracle     640        4096       0
0x12e16844   2457603    oracle     640        12288      19
```



Earlier saw this with ASMM

```
key          shmid      owner      perms      bytes      nattch
0x00000000   8880128    oracle     640        2924544    38
0x00000000   8912897    oracle     640        360710144 19
0x00000000   8945666    oracle     640        5464064    19
0x12e16844   8978435    oracle     640        8192       19
```



Allocate SGA

<http://blog.tanelpoder.com/2007/08/21/oracle-11g-internals-part-1-automatic-memory-management/>

“ Oracle 11g likes to use /dev/shm for shared memory implementation

This is the Linux's POSIX-oriented SHM implementation, where everything, including shared memory segments, is a file ”

```
----- Shared Memory Segments -----
key          shmid      owner      perms      bytes      nattch
0x00000000  2359296   oracle     640        4096       0
0x00000000  2392065   oracle     640        4096       0
0x00000000  2424834   oracle     640        4096       0
0x12e16844  2457603   oracle     640        12288      19
```

```
$ ls -lo /dev/shm/*2392065* | head -5
```

```
-rw-r-----. 1 oracle 0 Apr 12 11:05 /dev/shm/ora_SLOB_2392065_0
-rw-r-----. 1 oracle 0 Apr 12 11:05 /dev/shm/ora_SLOB_2392065_1
-rw-r-----. 1 oracle 0 Apr 12 11:05 /dev/shm/ora_SLOB_2392065_10
-rw-r-----. 1 oracle 4194304 Apr 12 11:05 /dev/shm/ora_SLOB_2392065_100
-rw-r-----. 1 oracle 4194304 Apr 12 11:05 /dev/shm/ora_SLOB_2392065_101
```


Allocate SGA

```
----- Shared Memory Segments -----  
key          shmids  owner    perms    bytes    nattch  
0x000000000 2359296 oracle   640      4096     0  
0x000000000 2392065 oracle   640      4096     0  
0x000000000 2424834 oracle   640      4096     0  
0x12e16844 2457603 oracle   640     12288    19
```

```
$ du -k /dev/shm/*2359296* | awk '{a+=$1}END{print a/1024"MB"}'  
2.75391MB
```

```
$ du -k /dev/shm/*2392065* | awk '{a+=$1}END{print a/1024"MB"}'  
344MB
```

```
$ du -k /dev/shm/*2424834* | awk '{a+=$1}END{print a/1024"MB"}'  
0MB
```

Allocate SGA

Redo Buffer memory allocated lazily by Linux

```
$ du -k /dev/shm/*2424834*
```

```
0      /dev/shm/ora_SLOB_2424834_0  
0      /dev/shm/ora_SLOB_2424834_1
```

```
SQL> alter database open;
```

```
Database altered.
```

```
$ du -k /dev/shm/*2424834*
```

```
4096   /dev/shm/ora_SLOB_5177346_0  
1136   /dev/shm/ora_SLOB_5177346_1
```

Why is the SGA not allocated lazily too?

PRE_PAGE_SGA

Property	Description
Parameter type	Boolean
Default value	false

12c Oracle docs state:

PRE_PAGE_SGA Behaviour Change in Oracle Release 12c (Doc ID 1987975.1)

“The default value for the parameter PRE_PAGE_SGA has changed in 12c to TRUE. It used to be FALSE in prior releases.”

SGA size when PRE_PAGE_SGA = **false**

```
$ du -k /dev/shm/*8159233*|awk '{a+=$1}END{print a/1024"MB"}'  
200.312MB
```

SGA size when PRE_PAGE_SGA = **true**

```
$ du -k /dev/shm/*2392065*|awk '{a+=$1}END{print a/1024"MB"}'  
344MB
```

Action Ticklist

- Connect as sysdba
- Searches for a server parameter file in a platform-specific default location and, if not found, for a text initialization parameter file
- Reads the parameter file to determine the values of initialization parameters
- Allocates the SGA based on the initialization parameter settings
- Starts the Oracle background processes
- Opens the alert log and trace files and writes all explicit parameter settings to the alert log in valid parameter syntax

Background Processes

startup nomount

```
UID      PID  PPID  C  STIME TTY          TIME CMD
neil    3873  7712   0 12:59 pts/1      00:00:00 sqlplus
oracle  3895    1    0 13:00 ?          00:00:00 ora_pmon_SLOB
oracle  3897    1    0 13:00 ?          00:00:00 ora_psp0_SLOB
oracle  3899    1   10 13:00 ?          00:00:22 ora_vktm_SLOB
oracle  3903    1    0 13:00 ?          00:00:00 ora_gen0_SLOB
oracle  3905    1    0 13:00 ?          00:00:00 ora_mman_SLOB
oracle  3909    1    0 13:00 ?          00:00:00 ora_diag_SLOB
oracle  3911    1    0 13:00 ?          00:00:00 ora_dbrm_SLOB
oracle  3913    1    0 13:00 ?          00:00:00 ora_vkrm_SLOB
oracle  3915    1    0 13:00 ?          00:00:00 ora_dia0_SLOB
oracle  3917    1    0 13:00 ?          00:00:00 ora_dbw0_SLOB
oracle  3919    1    0 13:00 ?          00:00:00 ora_lgwr_SLOB
oracle  3921    1    0 13:00 ?          00:00:00 ora_ckpt_SLOB
oracle  3923    1    0 13:00 ?          00:00:00 ora_smon_SLOB
oracle  3925    1    0 13:00 ?          00:00:00 ora_reco_SLOB
oracle  3927    1    0 13:00 ?          00:00:00 ora_lreg_SLOB
oracle  3929    1    0 13:00 ?          00:00:00 ora_pxmn_SLOB
oracle  3931    1    0 13:00 ?          00:00:00 ora_mmon_SLOB
oracle  3933    1    0 13:00 ?          00:00:00 ora_mmln_SLOB
oracle  3934  3873   0 13:00 ?          00:00:00 oracleSLOB (DESCRIPTION=
```

startup nomount

```
3873 sqlplus
-- 3882 (oracleSLOB)
|   -- 3893 (/bin/df -k)
|   -- 3894 -- 3895 (PMON)
|   -- 3896 -- 3897 (PSP)
|   -- 3898 -- 3899 (VKTM)
|   |   -- 3900 -- 3901 (DISM)
|   -- 3902 -- 3903 (GEN0)
|   -- 3904 -- 3905 (MMAN)
|   -- 3906 -- 3907 (SA00)
|   -- 3908 -- 3909 (DIAG)
|   -- 3910 -- 3911 (DBRM)
|   -- 3912 -- 3913 (VKRM)
|   -- 3914 -- 3915 (DIA0)
|   -- 3916 -- 3917 (DBW0)
|   -- 3918 -- 3919 (LGWR)
|   -- 3920 -- 3921 (CKPT)
|   -- 3922 -- 3923 (SMON)
|   -- 3924 -- 3925 (RECO)
|   -- 3926 -- 3927 (LREG)
|   -- 3928 -- 3929 (PXMN)
|   -- 3930 -- 3931 (MMON)
|   -- 3932 -- 3933 (MMNL)
|   -- 3935 -- 3936 (M000)
|   -- 3937 -- 3938 (M000)
|   -- 3939 -- 3940 (M000)
|   -- 3941 -- 3942 (M000)
-- 3934 (oracleSLOB)
```

PSP0
Process SPawner Process
Spawns Oracle background
processes after initial instance
startup

NB: startup nomount

startup nomount

```
3873 sqlplus
-- 3882 (oracleSLOB)
|  -- 3893 (/bin/df -k)
|  -- 3894 -- 3895 (PMON)
|  -- 3896 -- 3897 (PSP)
|  -- 3898 -- 3899 (VKTM)
|  |  -- 3900 -- 3901 (DISM)
|  -- 3902 -- 3903 (GEN0)
|  -- 3904 -- 3905 (MMAN)
|  -- 3906 -- 3907 (SA00)
|  -- 3908 -- 3909 (DIAG)
|  -- 3910 -- 3911 (DBRM)
|  -- 3912 -- 3913 (VKRM)
|  -- 3914 -- 3915 (DIA0)
|  -- 3916 -- 3917 (DBW0)
|  -- 3918 -- 3919 (LGWR)
|  -- 3920 -- 3921 (CKPT)
|  -- 3922 -- 3923 (SMON)
|  -- 3924 -- 3925 (RECO)
|  -- 3926 -- 3927 (LREG)
|  -- 3928 -- 3929 (PXMN)
|  -- 3930 -- 3931 (MMON)
|  -- 3932 -- 3933 (MMNL)
|  -- 3935 -- 3936 (M000)
|  -- 3937 -- 3938 (M000)
|  -- 3939 -- 3940 (M000)
|  -- 3941 -- 3942 (M000)
-- 3934 (oracleSLOB)
```

NB: startup nomount

startup nomount

PSP0

```

3897 clone(child_stack=0, flags=..., child_tidptr=...) = 3898
3897 close(8) = 0
3897 read(7, <unfinished ...>
3898 set_robust_list(0x7fdf5e0ed540, 0x18) = 0
3898 close(7) = 0
3898 clone(child_stack=0, flags=..., child_tidptr=...) = 3899
3898 close(8) = 0
3898 exit_group(0) = ?

```

VKTM

Maybe just there to ensure all background processes are orphaned to PID 1?

UID	PID	PPID	C	STIME	TTY	TIME	CMD
oracle	3895	1	0	13:00	?	00:00:00	ora_pmon_SLOB
oracle	3897	1	0	13:00	?	00:00:00	ora_psp0_SLOB
oracle	3899	1	10	13:00	?	00:00:22	ora_vktn_SLOB

Or maybe not, their raison d'être is unknown to me

startup nomount

```
3873 sqlplus
-- 3882 (oracleSLOB)   Replaced by 3934
|  -- 3893 (/bin/df -k)
|  -- 3894 -- 3895 (PMON)
|  -- 3896 -- 3897 (PSP)
|
|      -- 3898 -- 3899 (VKTM)
|      |      -- 3900 -- 3901 (DISM)
|      -- 3902 -- 3903 (GEN0)
|      -- 3904 -- 3905 (MMAN)
|      -- 3906 -- 3907 (SA00)   SGA Allocator
|      -- 3908 -- 3909 (DIAG)
|      -- 3910 -- 3911 (DBRM)
|      -- 3912 -- 3913 (VKRM)
|      -- 3914 -- 3915 (DIA0)
|      -- 3916 -- 3917 (DBW0)
|      -- 3918 -- 3919 (LGWR)
|      -- 3920 -- 3921 (CKPT)
|      -- 3922 -- 3923 (SMON)
|      -- 3924 -- 3925 (RECO)
|      -- 3926 -- 3927 (LREG)
|      -- 3928 -- 3929 (PXMN)
|      -- 3930 -- 3931 (MMON)
|      -- 3932 -- 3933 (MMNL)
|      -- 3935 -- 3936 (M000)
|      -- 3937 -- 3938 (M000)
|      -- 3939 -- 3940 (M000)
|      -- 3941 -- 3942 (M000)
-- 3934 (oracleSLOB)
```

NB: startup nomount

Real time priority

Real time priority

```
# ps -fu oracle
UID      PID  PPID    C  STIME  TTY      TIME  CMD
oracle   3895    1    0  13:00  ?        00:00:00 ora_pmon_SLOB
oracle   3897    1    0  13:00  ?        00:00:00 ora_psp0_SLOB
oracle   3899    1   10  13:00  ?        00:00:22 ora_vktm_SLOB
oracle   3903    1    0  13:00  ?        00:00:00 ora_gen0_SLOB
oracle   3905    1    0  13:00  ?        00:00:00 ora_mman_SLOB
oracle   3909    1    0  13:00  ?        00:00:00 ora_diag_SLOB
oracle   3911    1    0  13:00  ?        00:00:00 ora_dbrm_SLOB
oracle   3913    1    0  13:00  ?        00:00:00 ora_vkrm_SLOB
oracle   3915    1    0  13:00  ?        00:00:00 ora_dia0_SLOB
oracle   3917    1    0  13:00  ?        00:00:00 ora_dbw0_SLOB
oracle   3919    1    0  13:00  ?        00:00:00 ora_lgwr_SLOB
oracle   3921    1    0  13:00  ?        00:00:00 ora_ckpt_SLOB
oracle   3923    1    0  13:00  ?        00:00:00 ora_smon_SLOB
oracle   3925    1    0  13:00  ?        00:00:00 ora_reco_SLOB
oracle   3927    1    0  13:00  ?        00:00:00 ora_lreg_SLOB
oracle   3929    1    0  13:00  ?        00:00:00 ora_pxmn_SLOB
oracle   3931    1    0  13:00  ?        00:00:00 ora_mmon_SLOB
oracle   3933    1    0  13:00  ?        00:00:00 ora_mmln_SLOB
oracle   3934  3873    0  13:00  ?        00:00:00 oracleSLOB (DESCRIPTION=
```

Real time priority

Real time priority

```
# ps -u oracle -o euser,pid,ppid,c,pri,rtprio,stime,TTY,time,cmd
UID      PID  PPID  C  PRI  RTPRIO  STIME  TTY      TIME  CMD
oracle   3895   1    0   19     - 13:00  ?       00:00:00 ora_pmon_SLOB
oracle   3897   1    0   19     - 13:00  ?       00:00:00 ora_psp0_SLOB
oracle   3899   1   10   41     1 13:00  ?       00:00:22 ora_vktm_SLOB
oracle   3903   1    0   19     - 13:00  ?       00:00:00 ora_gen0_SLOB
oracle   3905   1    0   19     - 13:00  ?       00:00:00 ora_mman_SLOB
oracle   3909   1    0   19     - 13:00  ?       00:00:00 ora_diag_SLOB
oracle   3911   1    0   19     - 13:00  ?       00:00:00 ora_dbrm_SLOB
oracle   3913   1    0   19     - 13:00  ?       00:00:00 ora_vkrm_SLOB
oracle   3915   1    0   19     - 13:00  ?       00:00:00 ora_dia0_SLOB
oracle   3917   1    0   19     - 13:00  ?       00:00:00 ora_dbw0_SLOB
oracle   3919   1    0   19     - 13:00  ?       00:00:00 ora_lgwr_SLOB
oracle   3921   1    0   19     - 13:00  ?       00:00:00 ora_ckpt_SLOB
oracle   3923   1    0   19     - 13:00  ?       00:00:00 ora_smon_SLOB
oracle   3925   1    0   19     - 13:00  ?       00:00:00 ora_reco_SLOB
oracle   3927   1    0   19     - 13:00  ?       00:00:00 ora_lreg_SLOB
oracle   3929   1    0   19     - 13:00  ?       00:00:00 ora_pxmn_SLOB
oracle   3931   1    0   19     - 13:00  ?       00:00:00 ora_mmon_SLOB
oracle   3933   1    0   19     - 13:00  ?       00:00:00 ora_mmln_SLOB
oracle   3934  3873   0   19     - 13:00  ?       00:00:00 oracle_SLOB (DESCR
```

Real time priority

```

3873 sqlplus
  -- 3882 (oracleSLOB)
  |   -- 3893 (/bin/df -k)
  |   -- 3894 -- 3895 (PMON)
  |   -- 3896 -- 3897 (PSP)
  |           -- 3898 -- 3899 (VKTM)
  |           |           -- 3900 -- 3901 (DISM)
  |           -- 3902 -- 3903 (GEN0)
  |           -- 3904 -- 3905 (MMAN)

```

```
3899 clone(child_stack=0, flags=..., child_tidptr=...) = 3900
```

```
3900 clone(child_stack=0, flags=..., child_tidptr=...) = 3901
```

```
3901 execve("/u01/app/oracle/product/12.1.0/dbhome 1/bin/oradism",
["ora_dism_SLOB"], [/* 34 vars */) = 0
```

```
3901 sched_setscheduler(3899, SCHED_RR, { 1 }) = 0
```

```
$ ls -l oradism
-rwsr-x---. 1 root oinstall 109247 Jul  7 2014 oradism
```



SETUID flag and root owned

Real time priority

NAME	VALUE	DESCRIPTION
<code>_high_priority_processes</code>	LMS*	High Priority Process Name Mask
<code>_highest_priority_processes</code>	VKTM	Highest Priority Process Name Mask

```
SQL> alter system set "_highest_priority_processes"='VKTM|LREG'
      2 scope=spfile sid='*';
```

```
SQL> startup force nomount
```

```
PID PRI RTPRIO CMD
2381 19 - ora_pmon_SLOB
2385 41 1 ora_vktm_SLOB
2415 41 1 ora_lreg_SLOB
```

I am NOT condoning the use of these parameters! Only showing you stuff

Action Ticklist

- Connect as sysdba
- Searches for a server parameter file in a platform-specific default location and, if not found, for a text initialization parameter file
- Reads the parameter file to determine the values of initialization parameters
- Allocates the SGA based on the initialization parameter settings
- Starts the Oracle background processes
- Opens the alert log and trace files and writes all explicit parameter settings to the alert log in valid parameter syntax

Write Parameters to Alert Log

```
3882 open("/u01/app/oracle/diag/rdbms/slob/SLOB/trace/alert_SLOB.log",
        O_WRONLY|O_CREAT|O_APPEND, 0660) = 18
3882 write(18, "System parameters with non-defau"..., 43) = 43
3882 close(18) = 0

3882 open("/u01/app/oracle/diag/rdbms/slob/SLOB/trace/alert_SLOB.log",
        O_WRONLY|O_CREAT|O_APPEND, 0660) = 18
3882 write(18, " processes = 100"..., 34) = 34
3882 close(18) = 0

3882 open("/u01/app/oracle/diag/rdbms/slob/SLOB/trace/alert_SLOB.log",
        O_WRONLY|O_CREAT|O_APPEND, 0660) = 18
3882 write(18, " shared_pool_size = 300"..., 34) = 34
3882 close(18) = 0

3882 open("/u01/app/oracle/diag/rdbms/slob/SLOB/trace/alert_SLOB.log",
        O_WRONLY|O_CREAT|O_APPEND, 0660) = 18
3882 write(18, " filesystemio_options = \"se\"..., 38) = 38
3882 close(18)
```

Action Ticklist

- Connect as sysdba
- Searches for a server parameter file in a platform-specific default location and, if not found, for a text initialization parameter file
- Reads the parameter file to determine the values of initialization parameters
- Allocates the SGA based on the initialization parameter settings
- Starts the Oracle background processes
- Opens the alert log and trace files and writes all explicit parameter settings to the alert log in valid parameter syntax

In summary...

You probably didn't learn much of use

Hopefully it was interesting

The End



@neiljdba



neil@osumo.co.uk



@yodbuk

