

RAC parameter configuration overview

Miguel Anjo – CERN Physics Databases

Distributed Databases Operations Workshop

November 2008



Objectives

- Raise the awareness
 - how important are the DB parameters
 - depend on the workload
 - depend on the hardware
- Tier-1 DBAs to re-check their DBs
 - Ask themselves if parameters are correct
 - Most of the cases they will be fine



Parameters depend on HW



- PhyDB Installation procedure
 - https://twiki.cern.ch/twiki/bin/view/PSSGroup/Installation_verbose
- PhyDB Streams Configuration Checklist
 - <https://twiki.cern.ch/twiki/bin/view/PSSGroup/StreamsConfigurationChecklist>
- Oracle Documentation (10.2)
 - Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide
 - Chapter 5 – Administering Database Instances and Cluster Databases
 - http://download.oracle.com/docs/cd/B19306_01/rac.102/b14197/dbinstmgt.htm
- ~~Google, forums~~
 - Badly set parameters can change whole DB behavior



We (CERN) are not “masters” ...

- ...neither we intend to be
- We use almost only default parameter values
- Values presented are for CERN PhyDB sizes, workload and hardware type
- Before setting any value, make sure it suits your needs
- We are not responsible for any damaged caused 😊



Why? / How to start?

- Why set parameters?
 - Describe database
 - Fully use resources available
 - Get best possible performance
- How to start?
 - Hardware specification (memory, disk, CPU, cores)
 - Type of DB usage (OLTP, DSS, mixed, streams)
 - Number of expected users, concurrent connections
 - Service Level Agreement / MoU



PFILE / SPFILE

- ~~OCR, Voting disk, ASM spfile~~
- Database parameter file: PFILE / SPFILE
 - PFILE – text based file
 - SPFILE – binary file
 - Backups with RMAN
 - Changes are persistent and allows auto-tuning
 - Location `$ORACLE_HOME/dbs/initlcgr1.ora`
 - Needs the same content on every instance
 - Attention not to have more than one `initSID.ora`, `init.ora`, `spfile.ora`, `spfileSID.ora` in `$ORACLE_HOME/dbs/`
 - Location `+DATADG1/LCGR/spfileLCGR.ora`
 - Shared disk accessible by all instances
 - Using spfile in ASM, init file (above) should point to it



See current parameter values

- select name, value, inst_id from GV\$PARAMETER
- show parameter xyz
- V\$PARAMETER shows current value for session
- V\$SPPARAMETER shows spfile value
- V\$PARAMETER2 shows current value for session in different format (better?)



Set parameters

- `alter system set processes=800
sid='*' scope=spfile;`

OR

1. Create `pfile='/tmp/init.ora'` from `spfile`
2. *Change the text file*
3. Shutdown
4. Create `spfile='+DATADG1/lcgr/spfile.ora'`
from `pfile='/tmp/init.ora'`;
5. Startup

Contents of `init.ora` – attention to RAC parameters

- `*.processes=800`
- `lcgr1.processes=500`
- `lcgr2.processes=600`

Reset parameter

- `Alter system reset processes sid='*' scope=spfile;`



- Static parameters
 - cannot be changed without restart DB instance
 - Should be set carefully
 - expecting the growth and or spikes
 - Changed only on SPFILE
- Dynamic parameters
 - Can be changed at any time
 - Either just in memory or memory + SPFILE
 - Few can also be changed at SESSION level



Documentation says...



DB_CACHE_SIZE

Property	Description
Parameter type	Big integer
Syntax	DB_CACHE_SIZE = integer [K M G]
Default value	<p>If SGA_TARGET is set: If the parameter is not specified, then the default is 0 (internally determined by the Oracle Database). If the parameter is specified, then the user-specified value indicates a minimum value for the memory pool.</p> <p>If SGA_TARGET is not set, then the default is either 48 MB or 4MB * number of CPUs * granule size, whichever is greater</p>
Modifiable	ALTER SYSTEM
Basic	No

DB_CACHE_SIZE specifies the size of the DEFAULT buffer pool for buffers with the primary block size (the block size defined by the DB_BLOCK_SIZE initialization parameter).

Describe DB

- *.cluster_database=TRUE
- *.db_domain='cern.ch'
- *.db_name=[DBNAME]
- *.global_names=TRUE
- Automatic set by DBCA (one per instance)
 - instance_number
 - local_listener
 - thread
 - undo_tablespace



Memory parameters

	RAM	16G	8G	4G
SGA_TARGET	9904	4900	2200	
SGA_MAX_SIZE	9904	4900	2200	
PGA_AGGREGATE_TARGET	3072	2000	1400	
PROCESSES	2000			800

- Not using ASMM (Automatic Shared Memory Management)

- Example for 16G RAM instances

```

*.db_cache_advice=OFF # (optional)
*.db_cache_size=6900000000 # 6580M – v$db_cache_advice
*.shared_pool_size=2569011200 # 2450M – v$shared_pool_advice
*.java_pool_size=133554432 # 127M – v$java_pool_advice
*.large_pool_size=133554432 # 127M – tuned for parallel query
*.sga_target=0
    
```



Workload sizing parameters

- PROCESSES=2000 # 16G or 800 - 4G
 - Number of maximum connections + 20. *Not modifiable*
- UNDO_RETENTION=36000 # 10h
 - Time it is possible to flashback query, depends on space on undo tablespace
- DB_FILES=2000
 - Number datafiles on DB. Set default TS type BIGFILE. *Not modifiable*
- OPEN_CURSORS=300
 - Number opened cursors per session
- ARCHIVE_LAG_TARGET=4000
 - Maximum time before log switch
- LOG_BUFFER=10485760 # 10M
 - Max(0.5M, (128K*N_CPU)), extra space does not harm



Checking Atlas Tier-1s



	CERN	BNL	TRIUMPH	CERN	ASCG	PIC
processes	2000	4000	800	800	150	2000
open_cursors	300	300	300	300	300	300
RAM (GB)	16	16	10	8	8	8
sga_max_size (MB)	9984	9008	5120	4900	3072	4912
sga_target (MB)		8000				
pga_aggregate_target (MB)	3072	4256	1024	2000	384	2000
log_buffer (MB)	14	15	14		14	14
db_files	2000	200	200		200	2000
undo_retention (sec)	36000	3600	7200		3600	36000
parallel_max_servers	50	40	80		40	80
event (number, level)		26749,2	26749,2		26749,2	26749,2
default tablespace type	BIGFILE BIGFILE					

	CERN	NDGF	INFN	IN2P3	GRIDKA	RAL	SARA
processes	800	800	500	300	500	350	1000
open_cursors	300	300	300	300	300	300	2000
RAM (GB)	4	4	4	4	4	4	4
sga_max_size (MB)	2200	1712	2512	1536	2000	1520	1536
sga_target (MB)						(2G)	160
pga_aggregate_target (MB)	1400	394	1024	1594	389	800	394
log_buffer (MB)		15	14	14	15	15	15
db_files		200	200	200	200	200	200
undo_retention (sec)		900	900	900	7200	3600	7200
parallel_max_servers		20	20	160	80	40	20
event (number, level)				26749,2			26749,2
default tablespace type							



Streams – important parameters

- GLOBAL_NAMES=TRUE
- PARALLEL_MAX_SERVERS=20 # 3*capture_proc + 3*apply_proc
- STREAMS_POOL_SIZE=600M # Minimum 200M
- _JOB_QUEUE_INTERVAL=1 # Recheck after 10.2.0.3 upgrade
- JOB_QUEUE_PROCESSES=10 # Minimum 4; max simultaneous jobs + 2
- DB_DOMAIN='cern.ch'
- LOCAL_LISTENER=listener_lcgRX
- REMOTE_LISTENER=listeners_lcgRX
- *.recyclebin=OFF
- aq_tm_processes must not be explicitly set to 0 (OK if not specified in spfile)
- On Capture and Apply side
 - *.event='26749 trace name context forever, level 2'
- On Apply side
 - *."_buffered_publisher_flow_control_threshold"=30000
- On Capture side
 - *."_capture_publisher_flow_control_threshold"=80000



Changes 10.2.0.3 → 10.2.0.4

- When upgrading DB / changing HW
 - check validity of underscore parameters
 - Check need of not basic parameters
- **remove** `*.db_file_multiblock_read_count`
- **remove** `*._high_priority_processes=''`
 - set only for systems with 2 cores
- **remove** `*.event='10867 trace name context forever, level 30000'`



```

*.archive_lag_target=4000
*.cluster_database_instances=4
*.cluster_database=TRUE
*.compatible='10.2.0.3'
*.db_block_size=8192
*.streams_pool_size=600m
*.sga_target=9904M # 16G RAM
*.sga_max_size=9904M # 16G RAC
*.db_create_file_dest='+ [DBNAME] _DATADG1'
*.db_files=2000
*.db_domain='cern.ch'
*.db_name=[DBNAME]
*.db_recovery_file_dest='+ [DBNAME] _RECODG1'
*.db_recovery_file_dest_size=6000g
*.global_names=TRUE
*.job_queue_processes=10
*.log_archive_dest_1='LOCATION=USE_DB_RECOVERY
    _FILE_DEST'
*.log_archive_format='log_%t_%s_%r.arc'
*.log_buffer=10485760
*.open_cursors=300
*.parallel_max_servers=20
*.pga_aggregate_target=3g
*.processes=2000

*.recyclebin=OFF
*.remote_listener='listener_alias_here'
*.remote_login_passwordfile='exclusive'
*.resource_limit=TRUE
*.undo_management='AUTO'
*.undo_retention=36000
*.audit_file_dest='/ORA/dbs00/oracle/admin/[DB_
    _NAME]/adump'
*.core_dump_dest='/ORA/dbs00/oracle/admin/[DB_
    NAME]/cdump'
*.background_dump_dest='/ORA/dbs00/oracle/admi
    n/[DB_NAME]/bdump'
*.user_dump_dest='/ORA/dbs00/oracle/admin/[DB_
    NAME]/udump'
*.audit_trail='db'

```

ASM parameters

```
*.asm_diskgroups='DATADG1'  
*.asm_diskstring='/dev/mpath/itstor??_??p?'  
*.db_cache_size=80M  
*.cluster_database=true  
*.cluster_database_instances=4  
*.instance_type='asm'  
*.large_pool_size=20M  
*.asm_power_limit=5  
*.processes=100  
*.remote_login_passwordfile='exclusive'  
*.sga_max_size=200M  
*.shared_pool_size=90M  
*.audit_file_dest='/ORA/dbs00/oracle/admin/+ASM/adump'  
*.user_dump_dest='/ORA/dbs00/oracle/admin/+ASM/udump'  
*.background_dump_dest='/ORA/dbs00/oracle/admin/+ASM/bdump'  
*.core_dump_dest='/ORA/dbs00/oracle/admin/+ASM/cdump'  
+ASMx.instance_number=x  
+ASMx.local_listener='LISTENER_+ASMx'
```



- PhyDB Installation procedure
 - https://twiki.cern.ch/twiki/bin/view/PSSGroup/Installation_verbose
- PhyDB Streams Configuration Checklist
 - <https://twiki.cern.ch/twiki/bin/view/PSSGroup/StreamsConfigurationChecklist>
- Oracle Documentation (10.2)
 - Oracle Clusterware and Oracle Real Application Clusters Administration and Deployment Guide
 - Chapter 5 – Administering Database Instances and Cluster Databases
 - http://download.oracle.com/docs/cd/B19306_01/rac.102/b14197/dbinstmgt.htm





Lisbon, Portugal



Thanks.
Questions?

miguel.anjo@cern.ch