

# TUNING THE LARGE POOL FOR RMAN

SUCCESSFUL CUSTOM BR/DR/BC MODELS

NYOUG MEETING, DECEMBER 2008



**ANTHONY D NORIEGA**  
anthony.noriega@adnmis.com

**ADN R & D** [www.adnmis.com](http://www.adnmis.com)

**ORACLE**

**CERTIFIED  
PROFESSIONAL**



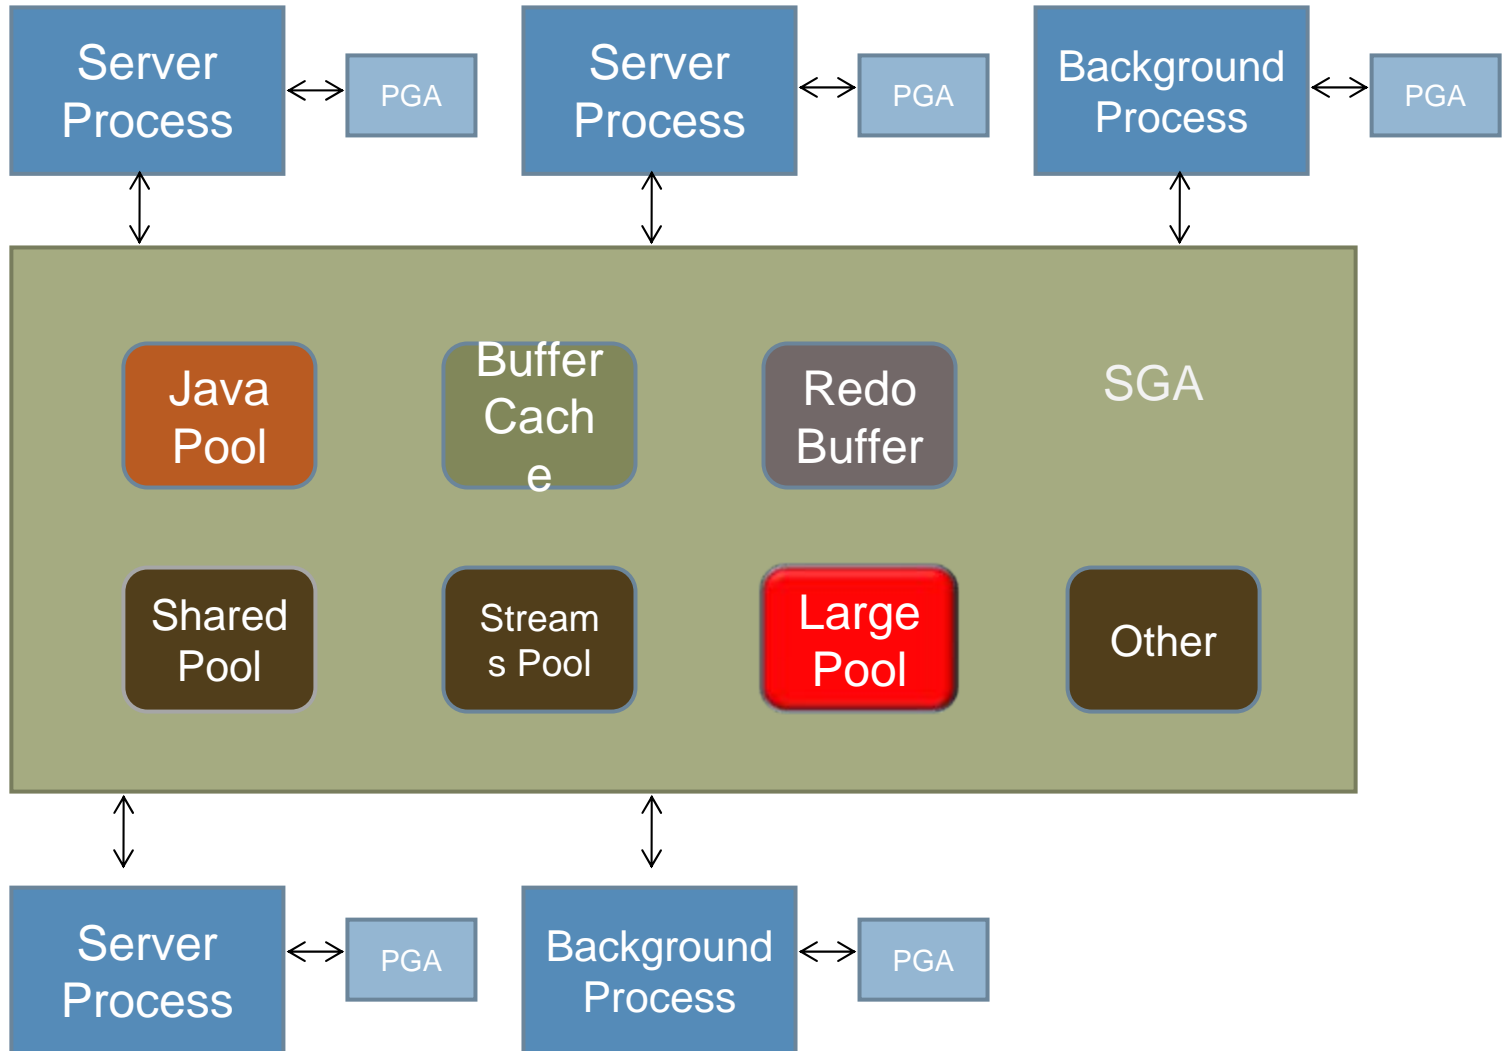
# Speaker Qualifications

- Independent Consultant, ADN
- Speaker at NYOUG meetings
- 24 years of IT experience
- 18 years of Oracle experience, 13 as a DBA
- RMAN experience with Oracle 8i, 9i, 10g, and 11g, since 1999.
- BS Systems Engineering, Universidad del Norte, 1987.
- MS Computer Science, NJIT, 1993
- PhD CIS candidate, NJIT, 1997
- MBA MIS, Montclair State University, 2006
- College Math Professor and former HS Math Teacher Principal.

# Objectives

- ❑ Provide a concise approach to tune the Large Pool for RMAN
- ❑ Analyze and compare Oracle recommendations and custom settings
- ❑ Present a robust approach to optimize RMAN time and storage space
- ❑ Derive a series of technical arguments and useful rules to attain a BR/BC/DR optimization methodology.

# Large Pool: Memory Structures



# Large Pool Purposes

- ❑ **Session memory for the shared server and the Oracle XA interface (used where transactions interact with more than one database)**
- ❑ **I/O server processes**
- ❑ **Oracle Database backup and restore operations**
- ❑ **Used by backup process for disk I/O buffers.**

# RMAN Reliability

- Based on experimental gathering RMAN tasks are more reliable with a Large pool sized with Oracle recommended settings.
- RMAN tasks tend to improve performance with higher Large Pool settings.

# Production Environment

- **Linux 4-node RACs, 9i and 10g**
- **SAN architecture using Symantec (Veritas) Netbackup**
- **MPP (Linux Intel 64-bit 8/16 dual-core processor RH Linux servers).**
- **Oracle RAC and Data Guard in place (MAA).**
- **Incremental backup level 0 mostly.**
- **Block Change Tracking (Oracle 10g/11g)**

# Oracle-Recommended Settings

- According to Oracle Corporation, the recommended setting for the Large Pool Size is 16M.
- So, your init.ora parameter should read:
- **LARGE\_POOL\_SIZE=16M**
- **fast\_start\_mttr\_target=900**



# Custom Settings

- Some of the custom settings are;
- **Third-Party**  
(Veritas Netbackup including both RAC databases and Standalone Databases)

<b>LARGE POOL SIZE</b>	<b>RELIABILITY</b>	<b>PERFORMANCE</b>
4M	Best	Decreased
8M	Better	Average
16M	Average	Improved
High-Value	Expected	Better

\* e.g., 32M-180M, 180M-2G +

# Custom Settings





- With ASM and File System Databases

□ <b>LARGE POOL SIZE</b>	<b>RELIABILITY</b>	<b>PERFORMANCE</b>
4M	Improved	Decreased
8M	Better	Average
16M	Average	Improved
High-Value	Expected	Better

\* e.g., 32M-180M, 180M-2G +

# Shared Segment and LP Size

- Analogy relating Large Pool and Shared segment sizes

<b>LARGE POOL SIZE</b>	<b>SHARED SEGMENT USED</b>
4M	
8M	
16M	
High-Value *	

\* e.g., 32M-180M, 180M-2G+

# Storage Space and Time

LARGE POOL SIZE	BACKUP PIECE SIZE	ESTIMATED BACKUP SIZE $\sum_{i=1}^n s_i$	DURATION
$2^j M$	$S_j = S_{j-1} + j \Delta s$	$\sum_{i=1}^n s_i = n(S_0 + j \Delta s)$	$t_0 = t_0 - \Delta t$
....	....	....	....
32M	$S_3 = S_2 + 3 \Delta s$	$\sum_{i=1}^n s_i = n(S_0 + 3 \Delta s)$	$t_1 = t_0 - \Delta t$
16M	$S_2 = S_1 + 2 \Delta s$	$\sum_{i=1}^n s_i = n(S_0 + 2 \Delta s)$	$t_0$
8M	$S_1 = S_0 + \Delta s$	$\sum_{i=1}^n s_i = n(S_0 + \Delta s)$	$t_{-1} = t_0 + \Delta t$
4M	$S_0$	$\sum_{i=1}^n s_i = n S_0$	$t_{-2} = t_0 + 2 \Delta t$

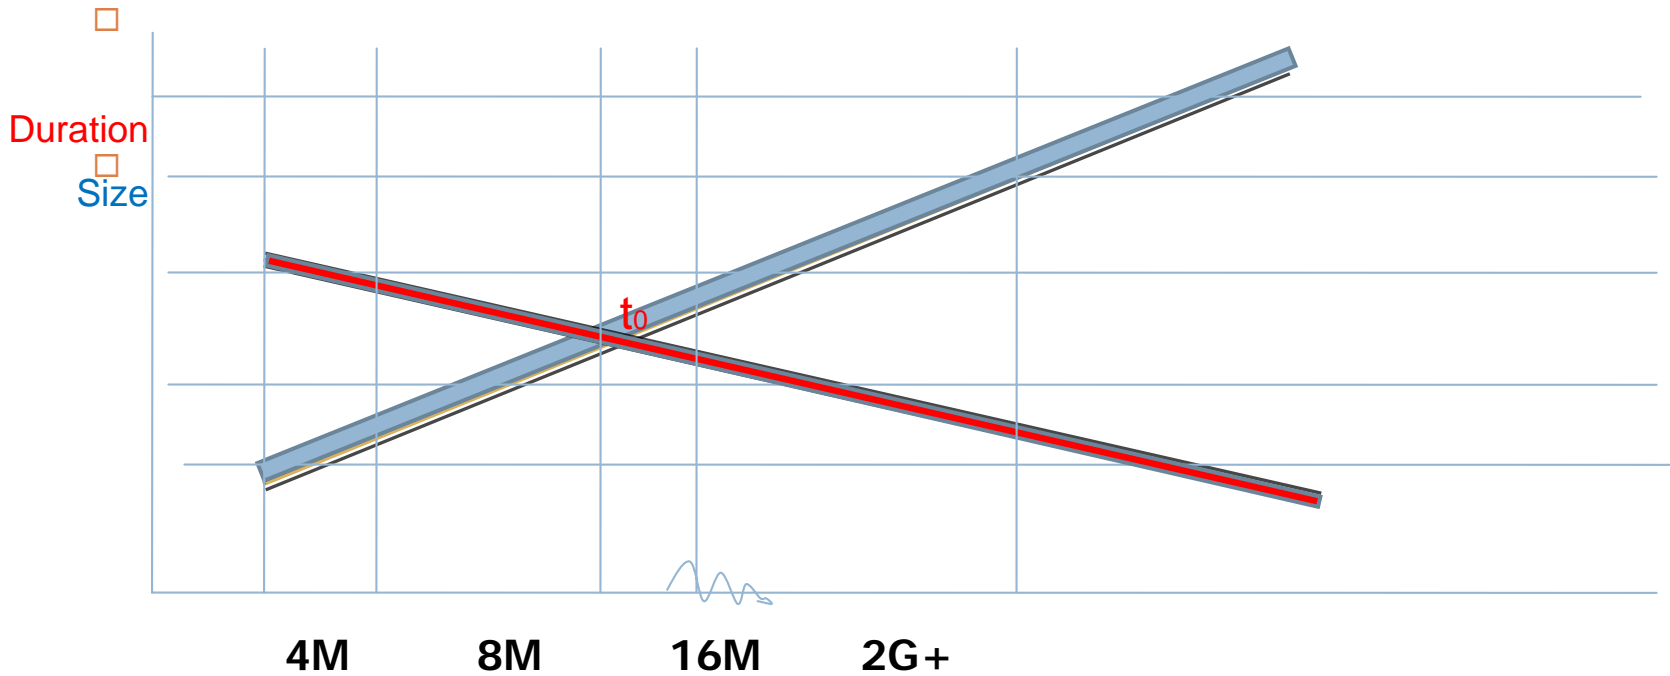
# Storage Space and Time

LARGE POOL SIZE	BACKUP PIECE SIZE	ESTIMATED BACKUP SIZE $\sum_{i=1}^n s_i$	DURATION
2jM	$S_j = S_{j-1} + j\Delta s$	$\sum_{i=1}^n s_i = n(S_0 + j\Delta s)$	$t_0 = t_0 - \Delta t$
...	...	...	...
32M	$S_3 = S_2 + 3\Delta s$	$\sum_{i=1}^n s_i = n(S_0 + 3\Delta s)$	$t_1 = t_0 - \Delta t$
16M	$S_2 = S_1 + 2\Delta s$	$\sum_{i=1}^n s_i = n(S_0 + 2\Delta s)$	$t_0$
8M	$S_1 = S_0 + \Delta s$	$\sum_{i=1}^n s_i = n(S_0 + \Delta s)$	$t_{-1} = t_0 + \Delta t$
4M	$S_0$	$\sum_{i=1}^n s_i = nS_0$	$t_{-2} = t_0 + 2\Delta t$

# Storage Space and Time

LARGE POOL SIZE	BACKUP PIECE SIZE	ESTIMATED BACKUP SIZE $\sum_{i=1}^n s_i$	DURATION
2iM	$S_j = S_{j-1} + j\Delta s$	$\sum_{i=1}^n s_i = n(S_0 + j\Delta s)$	$t_0 = t_j - \Delta t$
....	....	....	....
32M	$S_3 = S_2 + 3\Delta s$	$\sum_{i=1}^n s_i = n(S_0 + 3\Delta s)$	$t_1 = t_0 - \Delta t$
16M	$S_2 = S_1 + 2\Delta s$	$\sum_{i=1}^n s_i = n(S_0 + 2\Delta s)$	$t_0$
8M	$S_1 = S_0 + \Delta s$	$\sum_{i=1}^n s_i = n(S_0 + \Delta s)$	$t_{-1} = t_0 + \Delta t$
4M	$S_0$	$\sum_{i=1}^n s_i = nS_0$	$t_{-2} = t_0 + 2\Delta t$

# Backup Size and Duration



□ Relating the Large pool size and backup size and duration

Duration   
Size 

# RMAN Compatibility Table

Target/Auxiliary Database	RMAN client	Recovery Catalog Database	Recovery Catalog Schema
8.0.6	8.0.6	>=8.1.7	>=8.0.6
8.1.7	8.0.6.1	>=8.1.7	>=8.1.7
8.1.7	8.1.7	>=8.1.7	>=RMAN client
8.1.7.4	8.1.7.4	>=8.1.7	8.1.7.4
8.1.7.4	8.1.7.4	>=8.1.7	>= 9.0.1.4
9.0.1	9.0.1	>=8.1.7	>= RMAN client
9.2.0	>=9.0.1.3 and <= target database executable	>=8.1.7	>= RMAN client
10.1.0	>=9.0.1.3 and <= target database executable	>=9.0.1	>= RMAN client
10.2.0	>=9.0.1.3 and <= target database executable	>=9.0.1	>= RMAN client
11.1.0	>=9.0.1.3 and <= target database executable	>=9.0.1	>= RMAN client



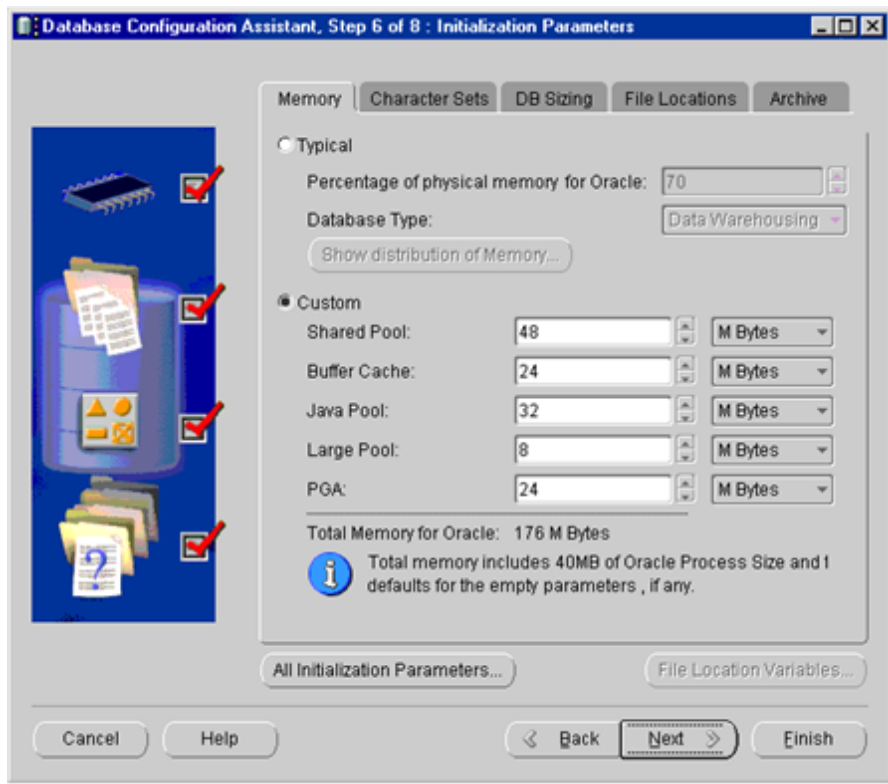
# Oracle By Version

<b>VERSION</b>	<b>LARGE POOL</b>	<b>REMARKS</b>
Oracle 8i	NO	N/A
Oracle9i	YES	Dynamic Parameter, BR disk buffers, Parallel Execution Message Buffer, No LRU, Optional. Defaults to 0. ksfqxc internal shared memory allocation error
Oracle10g	YES	Like 9i. Enhanced for Shared Server
Oracle11g	YES	Better satisfied segment requests than using shared pool. Faster parallel query execution. Automatic Memory Management.

# Oracle By Version

<b>VERSION</b>	<b>LARGE POOL</b>	<b>REMARKS</b>
Oracle9i	YES	<p>Dynamic Parameter, BR disk buffers, Parallel Execution Message Buffer, No LRU, Optional. Defaults to 0.</p> <p>ksfqxcrc internal shared memory allocation error</p> <p>Oracle recommended setting: 8M</p> <hr/> <p>300 KB to at least 2 GB (actual maximum is operating system-specific)</p>

# About Oracle9i Large Pool



```
SELECT POOL, NAME, SUM(BYTES)
FROM V$SGSTAT
WHERE POOL LIKE '%pool%'
GROUP BY ROLLUP (POOL, NAME);
```

POOL	NAME	SUM(BYTES)
large pool	PX msg pool	38092812
large pool	free memory	299988
large pool		38392800
....		
....		

# Oracle By Version

---

## □ **VERSION    REMARKS**

Oracle10g    Dynamic Parameter, BR disk buffers (Oracle allocates buffers that are a few hundred kilobytes)  
Parallel query, Parallel Execution Message Buffer  
No LRU and no aging out of pool  
Optional.

---

Specified **LARGE\_POOL\_SIZE** value a minimum value for the memory pool (allocation heap).

---

300 KB to at least 2 GB (actual maximum is operating system-specific)

---

# Oracle By Version

## □ **VERSION    REMARKS**

---

Oracle11g    Dynamic Parameter, BR disk buffers (Oracle allocates buffers that are a few hundred kilobytes)

Parallel query, Parallel Execution Message Buffer

No LRU and no aging out of pool

Optional.

Defaults to 0. Minimum Value 300K.

30M recommended for about 100 sessions (shared server, about 200-300k/session).

For effective sort usage of LP, set `sort_area_size` and `sort_area_retained_size` to same value.

~~Practical in Automatic Memory Management.~~

---

# Oracle11g Dynamic Sizing

- **V\$MEMORY\_CURRENT\_RESIZE\_OPS**: current memory resize operations (both automatic and manual) in progress.
- **V\$MEMORY\_DYNAMIC\_COMPONENTS**: current sizes of all dynamically tuned memory components, including SGA and PGA.
- **V\$MEMORY\_RESIZE\_OPS**: last 800 completed memory resize operations (both automatic and manual), no current operations.
- **V\$MEMORY\_TARGET\_ADVICE**: tuning advice for the MEMORY\_TARGET initialization parameter.
- **V\$SGA\_CURRENT\_RESIZE\_OPS**: current SGA resize operations in progress, i.e., a grow or a shrink of a dynamic SGA component.
- **V\$SGA\_RESIZE\_OPS**: last 800 completed SGA resize operations. No current operations.
- **V\$SGA\_DYNAMIC\_COMPONENTS**: dynamic components in SGA, summary of all completed SGA resize operations since startup.
- **V\$SGA\_DYNAMIC\_FREE\_MEMORY**: amount of SGA memory available for future dynamic SGA resize operations.

# Sample RMAN Shell

**RMAN  
Backup  
shell for  
Oracle  
Instances  
Studied  
(Without  
a  
Recovery  
Catalog)**

```
run
{
set until time =
"trunc(sysdate,'dd')-$NRETRODAYS+$HR/24+$MN/1440+$SC/86400";
crosscheck backupset of controlfile;
alter database mount ;
crosscheck backupset of database;
crosscheck backupset of archivelog all;
crosscheck archivelog all;
}
run
{
set until time =
"trunc(sysdate,'dd')-$NRETRODAYS+$HR/24+$MN/1440+$SC/86400";
restore controlfile to 'xx' until time="trunc(sysdate,'dd')-$NRETRODAYS
+ $HR/24 + $MN/1440 + $SC/86400" validate;
restore database until time="trunc(sysdate,'dd')-$NRETRODAYS + $HR/24 +
$MN/1440 + $SC/86400" validate;
restore archivelog time between "trunc(SYSDATE-$NRETRODAYS,'dd')" and
"trunc(SYSDATE-$NRETRODAYS,'dd') + $HR/24+$MN/1440 + $SC/86400"
validate;
}
exit;
EORMANV
#) | tee $LOG
```

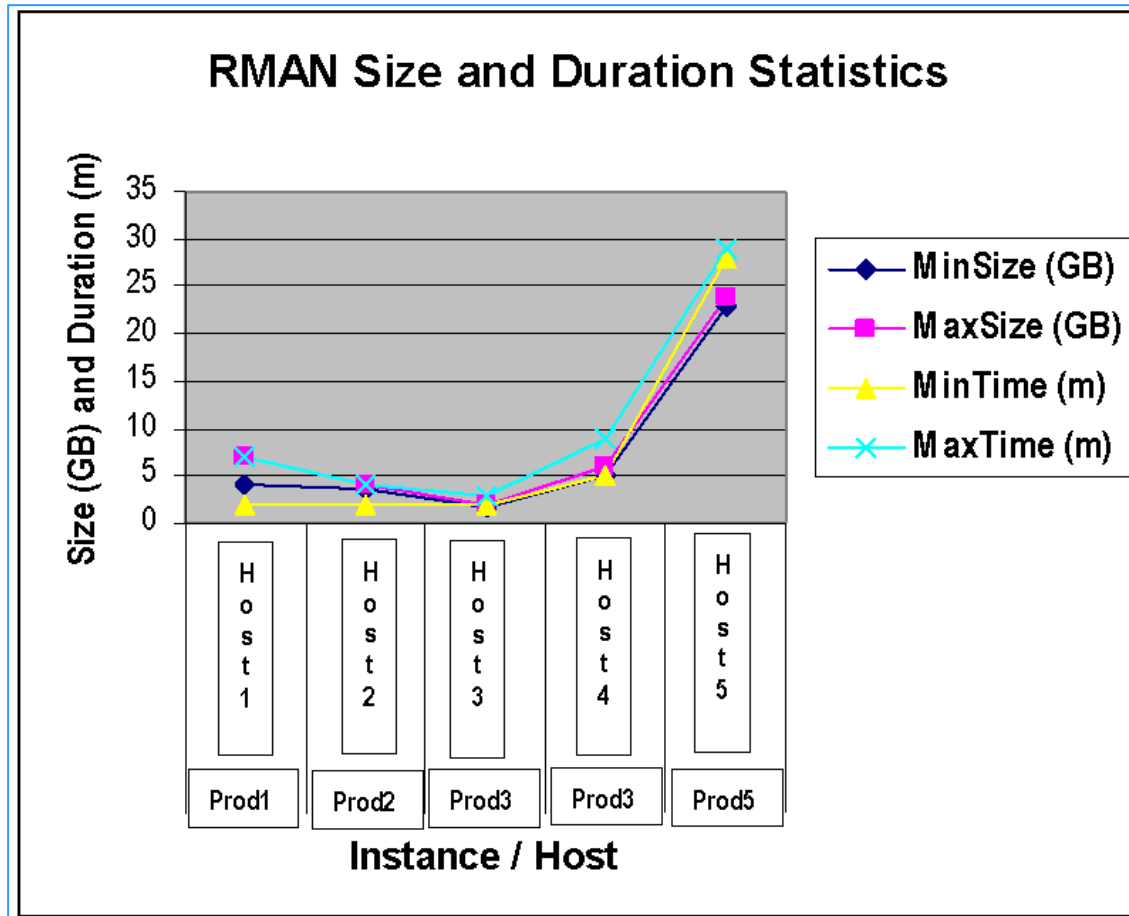
# Backup Size and Duration

Instance Name	Host Name	Average Database Size	Rman Backup Min Size (GB)	Rman Backup Max Size (GB)	Min Duration (Minutes)	Max Duration (Minutes)
Prd1	Host1	512GB	21	81	55	88
Dv1	Host2	512GB	22	99	65	90
Risk1	Host3	256GB	12	31	45	68
Uat1	Host4	300GB	30	105	45	77
Risk2	Host5	628GB	23.0	121	65	116

- **RMAN Backup Duration for Oracle Instances Studied (With a Recovery Catalog)**



# Backup Size Duration Line Chart



# Sample Shell

- **RMAN Backup shell for Oracle Instances Studied (With a Recovery Catalog)**

```
75 rman <<EORMAN
76 connect target $$SCHEMA1/$PASS1@eisqa
77 connect rcvcat $$SCHEMA2/$PASS2@oem
78   run {
79     allocate channel m1 type disk;
80     sql 'alter system archive log current';
81     resync catalog;
82     change archivelog all crosscheck;
83     release channel m1;
84     allocate channel c1 type disk;
85     setlimit channel c1 kbytes ${FILESIZE};
86     backup
87       incremental level ${LEVEL}
88       format '${DFILES}_%U.rman'
89       tag = '${BACKUPTAG}'
90       (database);
91     backup
92       format '${CFILES}_%U.rman'
93       (current controlfile);
94     sql 'alter system archive log current';
95     backup
96       format '${AFILES}_%U.rman'
97       (archivelog all);
98     release channel c1;
99   }
100   list backupset of database;
101   list backupset of controlfile;
102   list backupset of archivelog all;
103 exit;
104 EORMAN
```

# Backup Size and Duration

Instance Name	Host Name	Average Database Size	Rman Backup Min Size (GB)	Rman Backup Max Size (GB)	Min Duration (Minutes)	Max Duration (Minutes)
Prd1	Host1	512GB	21	81	55	88
Dv1	Host2	512GB	22	99	65	90
Risk1	Host3	256GB	12	31	45	68
Uat1	Host4	300GB	30	105	45	77
Risk2	Host5	628GB	23.0	121	65	116

- **RMAN Backup Duration for Oracle Instances Studied (With a Recovery Catalog)**

# A Simple RMAN Regression Model

```
SQL> get d:\batch\sqlAn.sql
 1 SELECT db_name,avg_db_size,max_duration,
 2 REGR_SLOPE(avg_db_size, max_duration)
 3     OVER (PARTITION BY backup_type) slope,
 4 REGR_INTERCEPT(avg_db_size, max_duration)
 5     OVER (PARTITION BY backup_type) intcpt,
 6 REGR_R2(avg_db_size, max_duration)
 7     OVER (PARTITION BY backup_type) rsqr,
 8 REGR_COUNT(avg_db_size, max_duration)
 9     OVER (PARTITION BY backup_type) count,
10 REGR_AVGX(avg_db_size, max_duration)
11     OVER (PARTITION BY backup_type) avgx,
12 REGR_AVGY(avg_db_size, max_duration)
13     OVER (PARTITION BY backup_type) avgy
14 FROM rman_records
15* ORDER BY rsqr desc
SQL> /
```

DB_NAME	AUG_DB_SIZE	MAX_DURATION	SLOPE	INTCPT	RSQR	COUNT	AUGX	AUGY
Uat1	300	77	8.12163814	-271.47983	.870652924	5	87.8	441.6
Risk1	256	68	8.12163814	-271.47983	.870652924	5	87.8	441.6
Dv1	512	90	8.12163814	-271.47983	.870652924	5	87.8	441.6
Prd1	512	88	8.12163814	-271.47983	.870652924	5	87.8	441.6
Risk2	628	116	8.12163814	-271.47983	.870652924	5	87.8	441.6
dev2	248	9	4.71880492	176.924429	.58766193	5	10.4	226
dev1	128	3	4.71880492	176.924429	.58766193	5	10.4	226
prod2	198	4	4.71880492	176.924429	.58766193	5	10.4	226
prod1	256	7	4.71880492	176.924429	.58766193	5	10.4	226
QA1	300	29	4.71880492	176.924429	.58766193	5	10.4	226

```
10 rows selected.
```

# RMAN Configuration

```
Copyright (c) 1995, 2002, Oracle Corporation. All rights reserved.
```

```
RMAN> connect target /
```

```
connected to target database: DEUTWBG (DBID=1885319376)  
using target database controlfile instead of recovery catalog
```

```
RMAN> show all;
```

```
RMAN configuration parameters are:
```

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 30 DAYS;
```

```
CONFIGURE BACKUP OPTIMIZATION OFF;
```

```
CONFIGURE DEFAULT DEVICE TYPE TO DISK; # default
```

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;
```

```
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK TO '/opt/oracle/admin/deutwbg/ctl/deutwbg_%F';
```

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 1; # default
```

```
CONFIGURE DATAFILE BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
```

```
CONFIGURE ARCHIVELOG BACKUP COPIES FOR DEVICE TYPE DISK TO 1; # default
```

```
CONFIGURE MAXSETSIZE TO UNLIMITED; # default
```

```
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/opt/oracle/product/9.2.0/dbs/snapcf_deutwbg.f'; # default
```

```
RMAN> █
```

# Session Long Operations

```
1 select target,target_desc,message,last_update_time
2 from v$session_longops
3* where last_update_time > sysdate-15/1440
anoriega@adndb1>/
```

TARGET	TARGET_DESC	MESSAGE	LAST_UPDATE_TIME
1303	Set Count	RMAN: incremental datafile backup: Set Count 1303: 25992 out of 25992 Blocks done	26-jan-2006 15:44:51
1303	Set Count	RMAN: incremental datafile backup: Set Count 1303: 10363 out of 10363 Blocks done	26-jan-2006 15:44:51
20	backup	RMAN: aggregate output: backup 20: 19944 out of 19 944 Blocks done	26-jan-2006 15:44:53
51	copy	RMAN: aggregate input: copy 51: 500 out of 500 Blo cks done	26-jan-2006 15:44:23

# Business Considerations

- Overall Preparedness for any RMAN process.
  - Disaster Recovery (DR)
  - Business Continuity (BC)
  - Backup and Recovery (BR)
- SDLC
- Quality Assurance
  - Backup set approach
  - Archived redo log backup policy
  - Retention policy strategy.

# Concluding Remarks

- ❑ **The Large Pool Size Setting could affect BR/DR/BC Processes, involving both reliability and performance.**
- ❑ **Large Pool Size is a factor when Correlating database size/backup size and duration.**
- ❑ **The large pool size is directly proportional to the backup size and inversely proportional to the backup duration.**