

Oracle 10.2 for z/OS

Outsmarting the 31 Bit Line




Agenda

- Agenda
 - Oracle for z/OS Architecture
 - Sizing the SGA
 - Sizing the PGA
 - Additional Memory used in case of multiple Address Spaces
 - RAC in one LPAR

Oracle for z/OS 31 Bit Architecture

- SGA is shared across Address Spaces
- PGA is distributed across Address Spaces
- Result:
 - SGA Size is max 1.5 GB
 - PGA Size can be increased by adding Address Spaces





Oracle for z/OS 31 Bit Architecture

- How many Address Spaces do I need ?
 - Each AS can Address 2048 MB
 - minus private Area (SQA,CSA,LPA)
 - minus z/OS nucleus and related data
- The remaining memory contains:
 - SGA
 - User Context
 - The Oracle Kernel
 - Load Modules (ORARASC, ORARSSRB, ORADIE)



Oracle for z/OS 31 Bit Architecture

- Sizing the SGA
 - SGA
 - is shared accross Address Spaces
 - implemented with IARVSERV
 - SGA belongs to the first AS and is shared by all auxiliary Address Spaces
 - SGA_MAX_SIZE and SGA_TARGET
 - MAX Size is initially reserved in all Address Spaces



Oracle for z/OS 31 Bit Architecture

- Sizing the SGA
 - Recommendation for Oracle on z/OS only
 - **Do not use SGA_MAX_Size and SGA_TARGET**
 - The Reason: Memory is a restricted resource in a 31 Address Space. Memory has to be fine tuned
- Use the following parameters on z/OS
 - DB_CACHE_SIZE
 - SHARED_POOL_SIZE
 - LARGE_POOL_SIZE
 - JAVA_POOL_SIZE
 - STREAMS_POOL_SIZE



Oracle for z/OS 31 Bit Architecture

- Sizing the SGA
 - How can the SGA Size be optimized ?
 - Statspack / AWR
 - Shared Pool Advisory
 - Buffer Pool Advisory

Oracle for z/OS 31 Bit Architecture

Shared Pool Advisory DB/Inst: MYDB/MYINST Snap: 12501

-> SP: Shared Pool Est LC: Estimated Library Cache Factr: Factor

-> Note there is often a 1:Many correlation between a single logical object in the Library Cache, and the physical number of memory objects associated with it. Therefore comparing the number of Lib Cache objects (e.g. in v\$librarycache), with the number of Lib Cache Memory Objects is invalid.

Shared Pool Size (M)	SP Size	Est LC Size (M)	Est LC Mem Obj	Est LC Time Saved (s)	Est LC Time Saved Factr	Est LC Load Time (s)	Est LC Load Time Factr	Est LC Mem Hits
1,584	.6	294	9,632	#####	1.0	#####	#####	#####
1,872	.7	577	21,673	#####	1.0	#####	153.8	#####
2,160	.8	871	31,547	#####	1.0	#####	61.9	#####
2,448	.9	1,158	44,238	#####	1.0	#####	34.8	#####
2,736	1.0	1,445	63,228	#####	1.0	7,417	1.0	#####
3,024	1.1	1,732	76,193	#####	1.0	1	.0	#####
3,312	1.2	1,990	88,413	#####	1.0	1	.0	#####
3,600	1.3	1,990	88,413	#####	1.0	1	.0	#####
3,888	1.4	1,990	88,413	#####	1.0	1	.0	#####
4,176	1.5	1,990	88,413	#####	1.0	1	.0	#####
4,464	1.6	1,990	88,413	#####	1.0	1	.0	#####
4,752	1.7	1,990	88,413	#####	1.0	1	.0	#####

Oracle for z/OS 31 Bit Architecture

Shared Pool Advisory

- ◆ SP: Shared Pool Est. LC: Estimated Library Cache Factor: Factor
- ◆ Note there is often a 1:Many correlation between a single logical object in the Library Cache, and the physical number of memory objects associated with it. Therefore comparing the number of Lib Cache objects (e.g. in v\$librarycache), with the number of Lib Cache Memory Objects is invalid.

Shared Pool Size(M)	SP Size Factor	Est LC Size (M)	Est LC Mem Obj	Est LC Time Saved (s)	Est LC Time Saved Factor	Est LC Load Time (s)	Est LC Load Time Factor	Est LC Mem Obj Hits
1,040	0.30	421	5,395	65,311,766	1.00	650,632	1.90	566,255,501
1,392	0.40	765	66,092	65,483,669	1.00	473,639	1.38	567,483,196
1,744	0.50	1,110	80,234	65,537,550	1.00	424,818	1.24	567,803,733
2,096	0.60	1,461	94,778	65,571,987	1.00	390,331	1.14	567,892,505
2,448	0.70	1,812	109,673	65,582,572	1.00	379,736	1.11	567,936,821
2,800	0.80	2,163	124,856	65,592,675	1.00	369,633	1.08	567,963,231
3,152	0.90	2,514	139,892	65,613,019	1.00	346,349	1.01	568,048,277
3,504	1.00	2,865	154,366	65,623,123	1.00	342,245	1.00	568,149,276
3,856	1.10	3,216	169,736	65,624,004	1.00	338,354	0.99	568,259,541
4,208	1.20	3,567	185,409	65,625,300	1.00	337,058	0.98	568,293,582
4,560	1.30	3,917	199,992	65,634,759	1.00	327,639	0.96	568,304,203
4,912	1.40	4,268	215,512	65,644,537	1.00	317,831	0.93	568,311,167
5,264	1.50	4,619	231,474	65,663,366	1.00	294,032	0.86	568,322,847
5,616	1.60	4,970	246,965	65,677,670	1.00	284,638	0.83	568,384,815
5,968	1.70	5,321	263,435	65,683,057	1.00	279,311	0.82	568,393,519
6,320	1.80	5,672	276,024	65,683,166	1.00	279,232	0.82	568,398,268
6,672	1.90	6,023	292,051	65,683,276	1.00	279,032	0.82	568,402,448
7,024	2.00	6,374	307,892	65,683,338	1.00	279,030	0.82	568,406,185

Oracle for z/OS 31 Bit Architecture

Buffer Pool Advisory

DB/Inst: MYDB/MYINST Snap: 12501

-> Only rows with estimated physical reads >0 are displayed

-> ordered by Block Size, Buffers For Estimate

				Est	
	Size for	Size	Buffers for	Read	Estimated
P	Est (M)	Factor	Estimate	Factor	Physical Reads
D	1,712	.1	205,333	15.5	24,153,515
D	3,424	.2	410,666	3.4	5,354,652
D	5,136	.3	615,999	2.1	3,242,758
D	6,848	.4	821,332	1.7	2,616,904
D	8,560	.5	1,026,665	1.4	2,120,834
D	10,272	.6	1,231,998	1.1	1,789,339
D	11,984	.7	1,437,331	1.1	1,672,908
D	13,696	.8	1,642,664	1.1	1,643,114
D	15,408	.9	1,847,997	1.0	1,599,954
D	17,120	1.0	2,053,330	1.0	1,562,209
D	17,152	1.0	2,057,168	1.0	1,561,518
D	18,832	1.1	2,258,663	1.0	1,549,496
D	20,544	1.2	2,463,996	1.0	1,546,034
D	22,256	1.3	2,669,329	1.0	1,543,644
D	23,968	1.4	2,874,662	1.0	1,536,359



Oracle for z/OS 31 Bit Architecture

- Sizing the PGA
- The PGA is distributed across Address Spaces
 - PGA Size
 - Depending on behavior of the application
 - Open Cursors, SQL Statements, PL/SQL Procedures etc.
 - Database Region Parameters
 - INIT_STACK_SIZE



Oracle for z/OS 31 Bit Architecture

- Sizing the PGA
- Recommendation for Oracle on z/OS only
 - PGA_AGGREGATE_TARGET
 - Can be used if different workload characteristics are processed. Does not make sense with a static workload
 - If the workload characteristics do not vary
 - use
 - SORT_AREA_SIZE
 - HASH_AREA_SIZE
 -



Oracle for z/OS 31 Bit Architecture

- Sizing the PGA
 - How to find an optimum PGA Size
 - An undersized PGA results in High I/O activity on the temporary Tablespace.
 - PGA size can be increased unlimited by adding more Address Spaces



Oracle for z/OS 31 Bit Architecture

- Sizing the PGA
 - Additional Memory used by adding more Address Spaces
 - IARVSERV needs 32 Bytes of ESQA for each 4K Page in SGA
(primary and auxiliary Address Space)
 - Example SGA Size = 512 MB, 10 AS are used
$$32 * 10 * (512 * 1024 * 1024) / 4096 = 40M$$



Oracle for z/OS 31 Bit Architecture

- Sizing the PGA
 - Monitoring PGA usage with AWR/Statspack

	Total	per Second	per Trans
sorts (disk)	0.0	0	0.0
sorts (memory)	1.6	1,570	3.2
sorts (rows)	298,214	606.4	302.8

Oracle for z/OS 31 Bit Architecture

- **PGA Memory Advisory**

PGA Memory Advisory

- ◆ When using Auto Memory Mgmt, minimally choose a pga_aggregate_target value where Estd PGA Overalloc Count is 0

PGA Target Est (MB)	Size Factr	W/A MB Processed	Estd Extra W/A MB Read/Written to Disk	Estd PGA Cache Hit %	Estd PGA Overalloc Count
138	0.13	12,975,504.88	5,147,642.41	72.00	41,968
275	0.25	12,975,504.88	1,760,587.08	88.00	7,102
550	0.50	12,975,504.88	84,307.41	99.00	0
825	0.75	12,975,504.88	83,612.23	99.00	0
1,100	1.00	12,975,504.88	48,890.84	100.00	0
1,320	1.20	12,975,504.88	25,689.78	100.00	0
1,540	1.40	12,975,504.88	25,689.78	100.00	0
1,760	1.60	12,975,504.88	25,689.78	100.00	0
1,980	1.80	12,975,504.88	25,689.78	100.00	0
2,200	2.00	12,975,504.88	25,689.78	100.00	0
3,300	3.00	12,975,504.88	25,689.78	100.00	0
4,400	4.00	12,975,504.88	25,689.78	100.00	0
6,600	6.00	12,975,504.88	25,689.78	100.00	0
8,800	8.00	12,975,504.88	25,689.78	100.00	0

[Back to Advisory Statistics](#)



Oracle for z/OS 31 Bit Architecture

- Avg PGA Storage used per session has to be calculated
 - Total PGA Storage used can be determined from
 - Statspack
 - AWR
 - V\$PGASTAT
 - SMF-data



Oracle for z/OS 31 Bit Architecture

- PGA Memory used (AWR Report)

<code>statistic</code>	<code>Total</code>	<code>per Second</code>	<code>per Trans</code>
<code>session pga memory</code>	<code>14,326,472</code>	<code>29,130.3</code>	<code>14,544.6</code>
<code>session pga memory max</code>	<code>278,371,016</code>	<code>566,016.8</code>	<code>282,610.2</code>
<code>session uga memory</code>	<code>8,589,462,336</code>	<code>17,465,107.9</code>	<code>8,720,266.3</code>
<code>session uga memory max</code>	<code>12,675,376</code>	<code>25,773.1</code>	<code>12,868.4</code>

useful only in V\$SESSTAT; it has no meaning in V\$SYSSTAT.



Oracle for z/OS 31 Bit Architecture

PGA Memory used

Select * from v\$pgastat

NAME	VALUE	UNIT
-----	-----	-----
aggregate PGA target parameter	26214400	bytes
aggregate PGA auto target	15989760	bytes
global memory bound	5242880	bytes
total PGA inuse	8444928	bytes
total PGA allocated	8444928	bytes
maximum PGA allocated	18275328	bytes
maxium PGA used for manual workareas	0	bytes
over allocation count	0	
bytes processed	3907584	bytes
extra bytes read/written	0	bytes
cache hit percentage	100	percent
recompute count (total)	196	



Oracle for z/OS 31 Bit Architecture

PGA Memory used

- Session Memory usage can also be monitored with SMF Records
 - HI STG is the session memory used. Memory usage can be calculated based on that value

Oracle for z/OS 31 Bit Architecture

SMF Report

SSN	SERVICE	SMFAUTH	SMFCONN	ORACLE ID	DATE	TIME	CPU SECONDS	LOG READS	PHY READS	LOG WRITES	DMC	DMR	DED	HI	STG
OSDI	ORAS	MAIER	TSO		00.327	11:45:54.77	.029900	56	3	0	0	0	0	0	607K
OSDI	ORAS		TCP/IP		00.327	11:46:39.12	.216741	525	8	0	0	0	0	0	1440K
OSDI	ORAS		TCP/IP		00.327	11:56:01.20	.034338	21	0	0	0	0	0	0	704K
OSDI	ORAS		TCP/IP		00.327	11:30:34.31	2.746457	23768	179	1601	68	0	0	0	1094K
OSDI	ORAS	MAIER	BATCH		00.327	11:57:03.93	.199057	0	0	0	0	0	0	0	891K
OSDI	ORAS	SCHULZE	BATCH		00.327	11:58:17.73	.555580	3981	154	81	0	0	0	0	623K
OSDI	ORAS	MAIER	BATCH		00.327	07:09:20.66	.610526	4050	162	120	0	0	0	0	8755K
OSDI	ORAS		TCP/IP		00.327	07:10:56.53	.322252	4294967046	4294967283	4294967291	0	0	0	0	766K
OSDI	ORAS	SCHULZE	TSO		00.327	09:20:17.77	.017690	10	0	0	0	0	0	0	640K
OSDI	ORAS	MAIER	BATCH		00.327	11:29:21.12	.180678	2984	67	1	0	0	0	0	821K
OSDI	ORAS		TCP/IP		00.327	11:29:56.76	.304108	824	61	27	0	0	0	0	962K
OSDI	ORAS	MAIER	BATCH		00.327	11:34:43.92	1.034874	5806	66	1083	9	0	0	0	1430K



Oracle for z/OS 31 Bit Architecture

- `INIT_STACK_SIZE`
 - Default is 128k
 - Storage is acquired, in an Oracle Region, when a user session starts. The storage is acquired in „pieces“ of the size `INIT_STACK_SIZE`
 - Parameter can be modified and Storage usage can be monitored (e.g. with SMF-records)



Oracle for z/OS 31 Bit Architecture

- Miscellaneous Parameters
 - MAX_SESSIONS
 - Limitation of Sessions per Address Space
 - MAX_SESSION_MEM
 - Defines the maximum memory used by a session
 - ORA-04030 is returned if more memory is requested by a session



Oracle for z/OS 31 Bit Architecture


- How many Address Spaces are used
 - Memory available = 2GB – SGA – z/OS Private Storage
 - # of sessions per AS = $\frac{\text{Memory available}}{\text{Avg Storage used per session}}$
 - Avg Storage used per session = AVG(HIGHWTM SMF Record (202) or Statspack/AWR Report information)
- # of Address Spaces = $\frac{\text{max \# of sessions}}{\text{\# of sessions per AS}}$



Oracle for z/OS 31 Bit Architecture

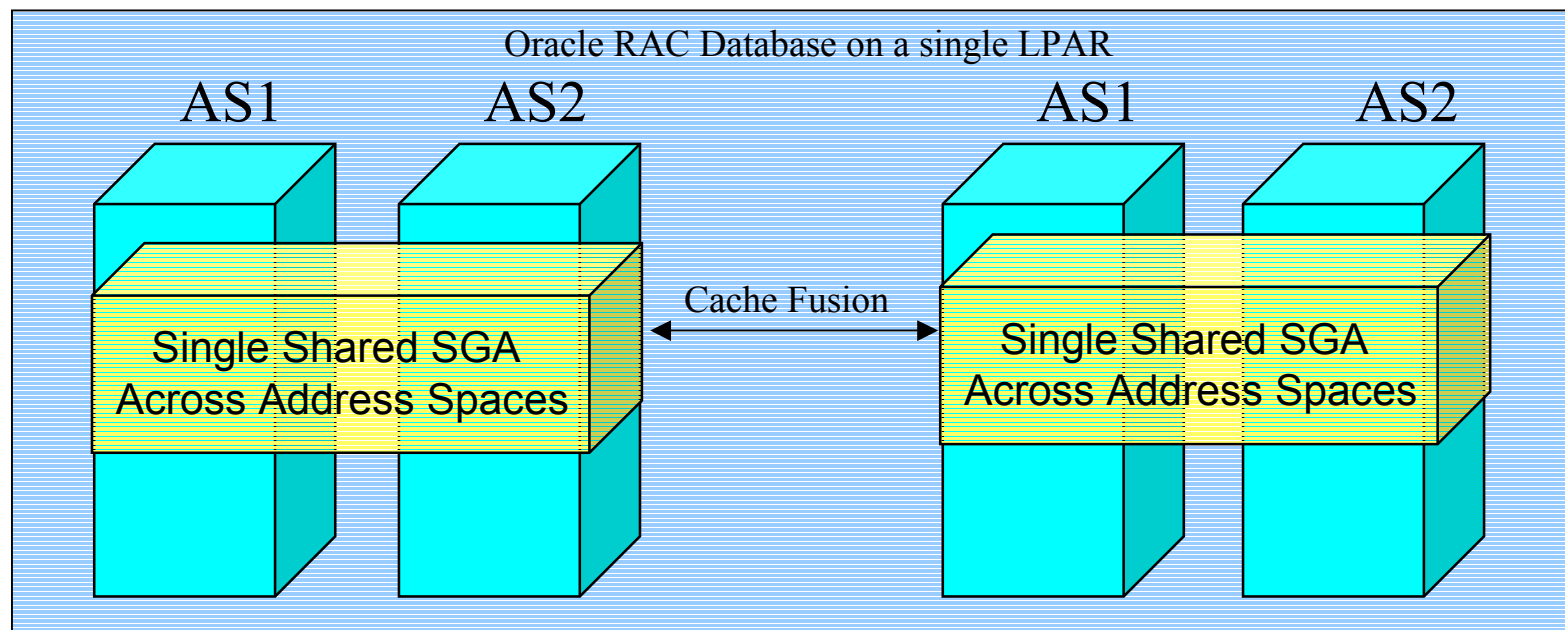
- How can the SGA be increased
 - SGA can be increased by using RAC on one LPAR
 - Max. 1,5 GB SGA can be defined for each member of the RAC database
 - Cache Fusion is performed in XM Mode

Oracle 10g Real Application Clusters



Increasing SGA and PGA size

- Oracle for z/OS allows an unlimited sizing of the PGA
- SGA has to be sized in the possible ranges of an 31 bit Address Space per Instance in a RAC Database





ORACLE IS THE INFORMATION COMPANY