ORACLE®

Oracle memory Managment

Bryan Grenn
Principal Engineered systems sales consultant

PARTNERNETWORK

Specialized. Recognized by Oracle.
Preferred by Customers.

Meltdown

Topics for Today

- SGA/PGA settings
- Huge pages
- How to size your machine
- Cursor sharing
- Result cache

- Extra Credit -- NUMA vs UMA

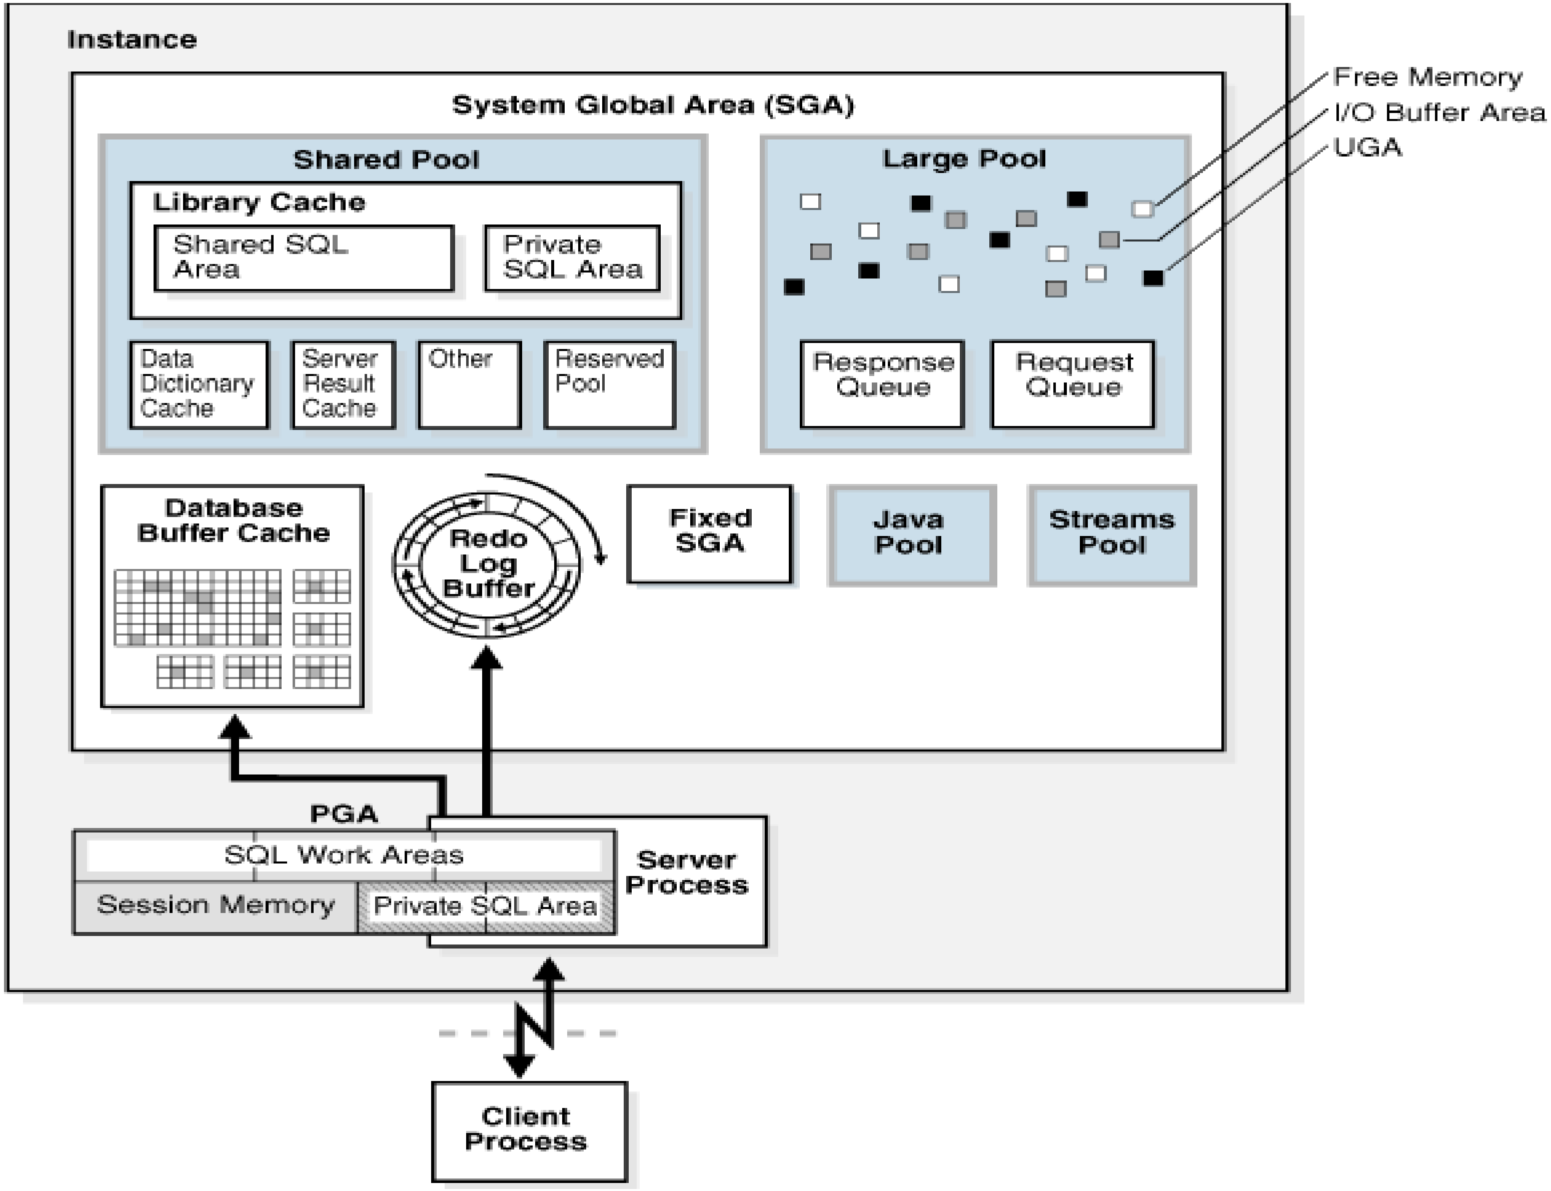
Instance
System Global Area (SGA)
Shared Pool
Library Cache
Shared SQL Area
Private SQL Area
Data Dictionary Cache
Server Result Cache
Other
Reserved Pool
Large Pool
Free Memory
I/O Buffer Area
UGA
Response Queue
Request Queue
Database Buffer Cache
Redo Log Buffer
Fixed SGA
Java Pool
Streams Pool
PGA
SQL Work Areas
Session Memory
Private SQL Area
Server Process
Client Process

PGA

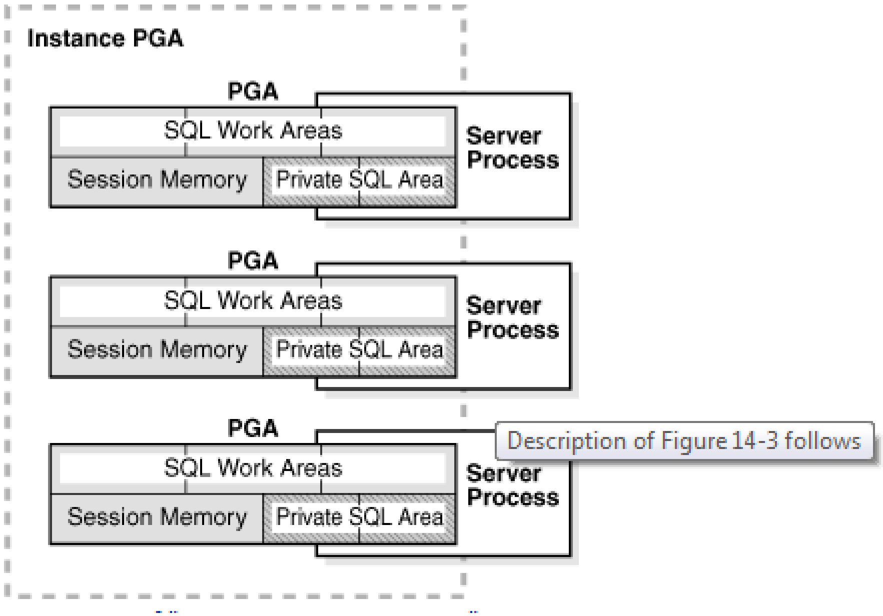
Instance PGA
PGA
SQL Work Areas
Session Memory
Private SQL Area
Server Process
PGA
SQL Work Areas
Session Memory
Private SQL Area
Server Process
PGA
SQL Work Areas
Session Memory
Private SQL Area
Server Process
Description of Figure 14-3 follows

PGA

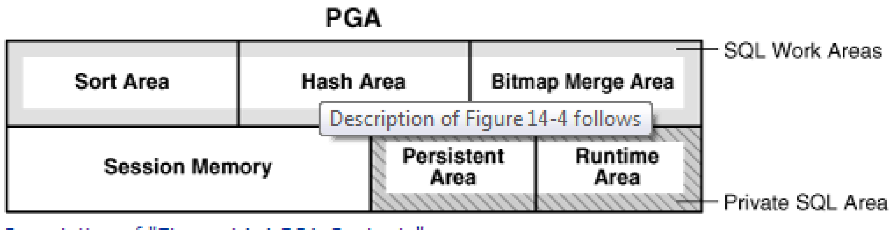
PGA
Sort Area
Hash Area
Bitmap Merge Area
SQL Work Areas
Description of Figure 14-4 follows
Session Memory
Persistent Area
Runtime Area
Private SQL Area

Sga

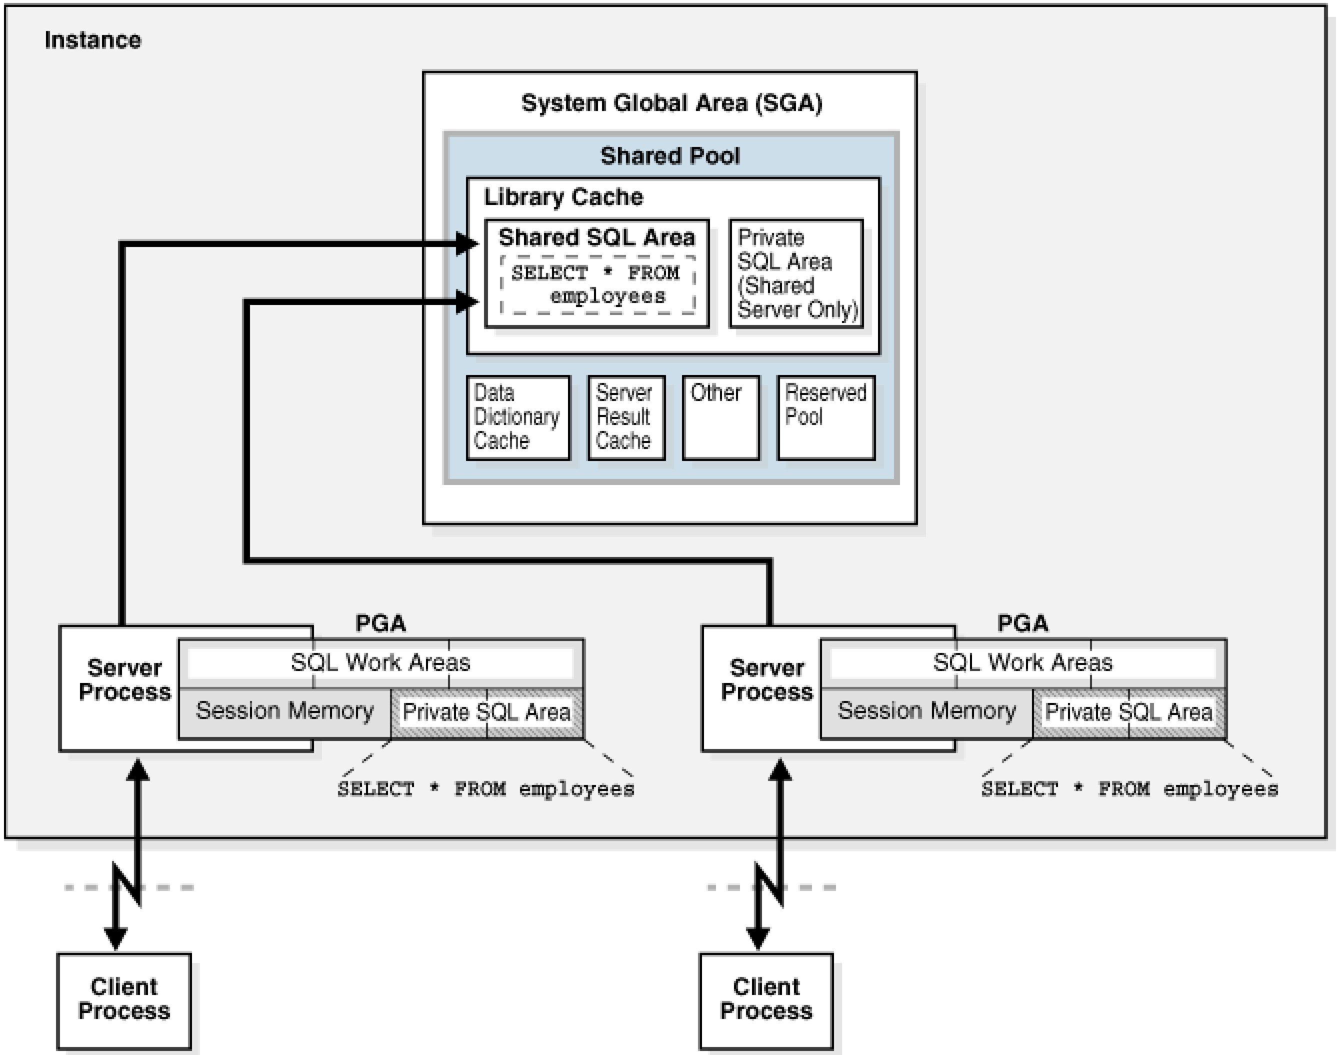
Instance
System Global Area (SGA)
Shared Pool
Library Cache
Shared SQL Area
SELECT * FROM employees
Private SQL Area (Shared Server Only)
Data Dictionary Cache
Server Result Cache
Other
Reserved Pool
PGA
Server Process
SQL Work Areas
Session Memory
Private SQL Area
SELECT * FROM employees
PGA
Server Process
SQL Work Areas
Session Memory
Private SQL Area
SELECT * FROM employees
Client Process
Client Process

SGA/PGA settings

- Automatic Memory Management
- SGA_TARGET/PGA_TARGET
- Manual Memory managment

Automatic Memory Management

- SGA and PGA are set by the system

SGA_TARGET/PGA_TARGET

- SGA_target/SGA_MAX_SIZE is set
- PGA_TARGET is set

Manual Memory management

- SGA_TARGET=0,Memory_TARGET=0

Different block sizes

- **Example of Setting Block and Cache Sizes**
- DB_BLOCK_SIZE=4096
- DB_CACHE_SIZE=1024M
  DB_2K_CACHE_SIZE=256M
  DB_8K_CACHE_SIZE=512M

DB block buffers

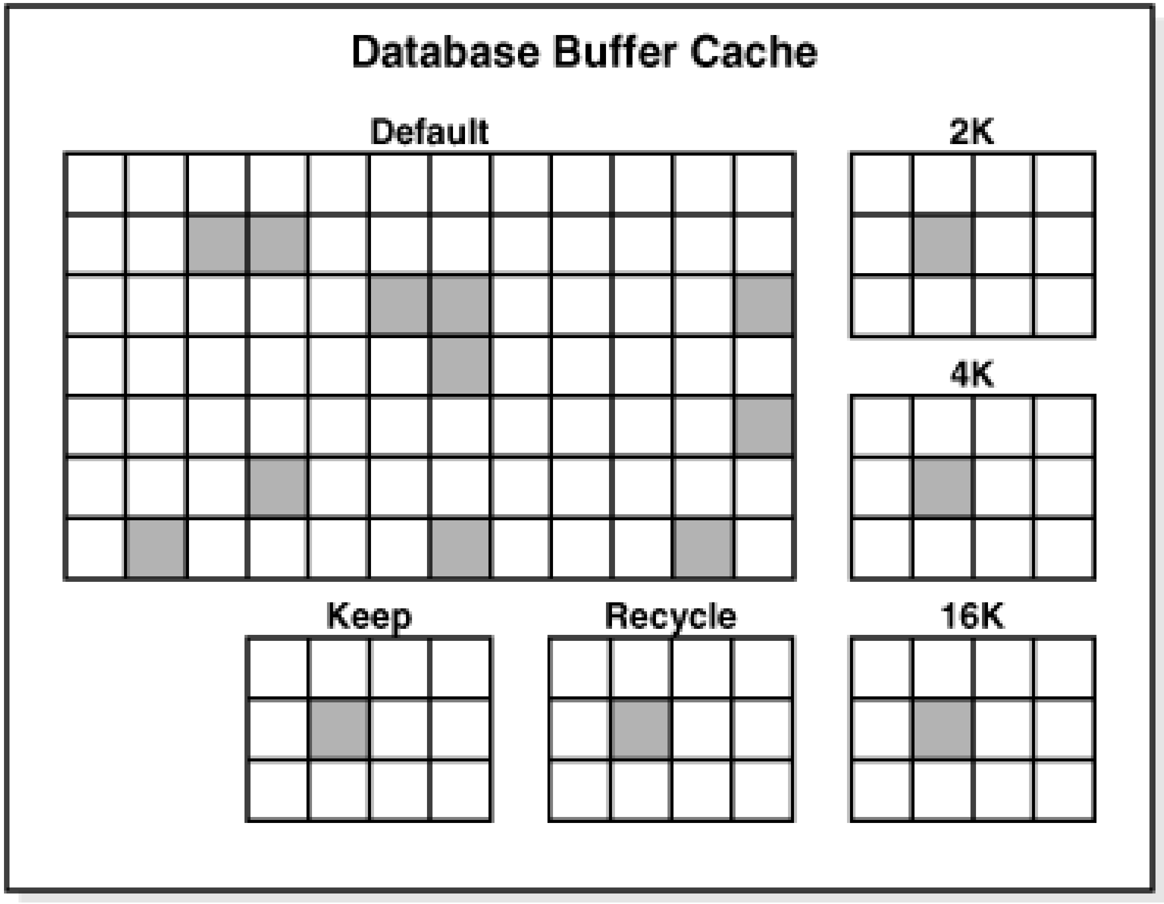

Redo log buffers

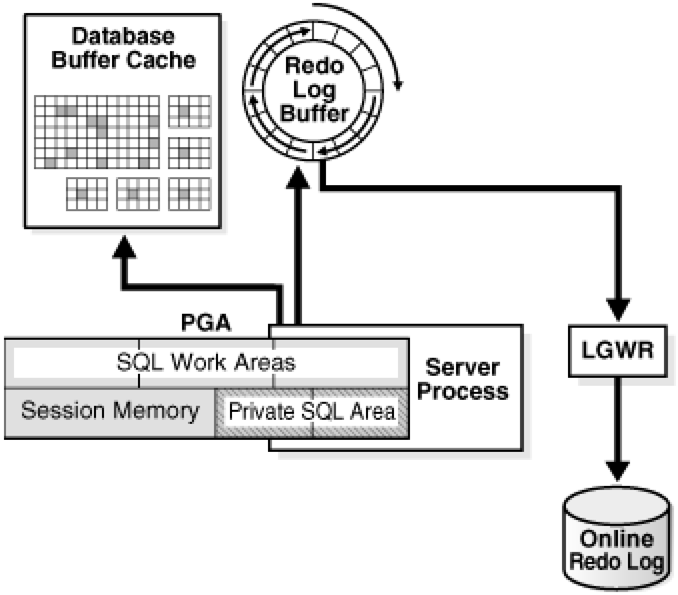

PGA sizing

- For OLTP: PGA_AGGREGATE_TARGET = (*total_mem* * 80%) * 20%
- For DSS: PGA_AGGREGATE_TARGET = (*total_mem* * 80%) * 50%
- where *total_mem* is the total amount of physical memory available on the system

SGA sizing

- Shared pool
- Buffer cache
- Redo log buffer size

Sizing Examples

	Count	Size per	Total (g)
Server sizing			128
Available (size * 80%)			102.4
SGA			60
pga			10
Processes	2000	0.015	30
			100

SGA (how am I doing)

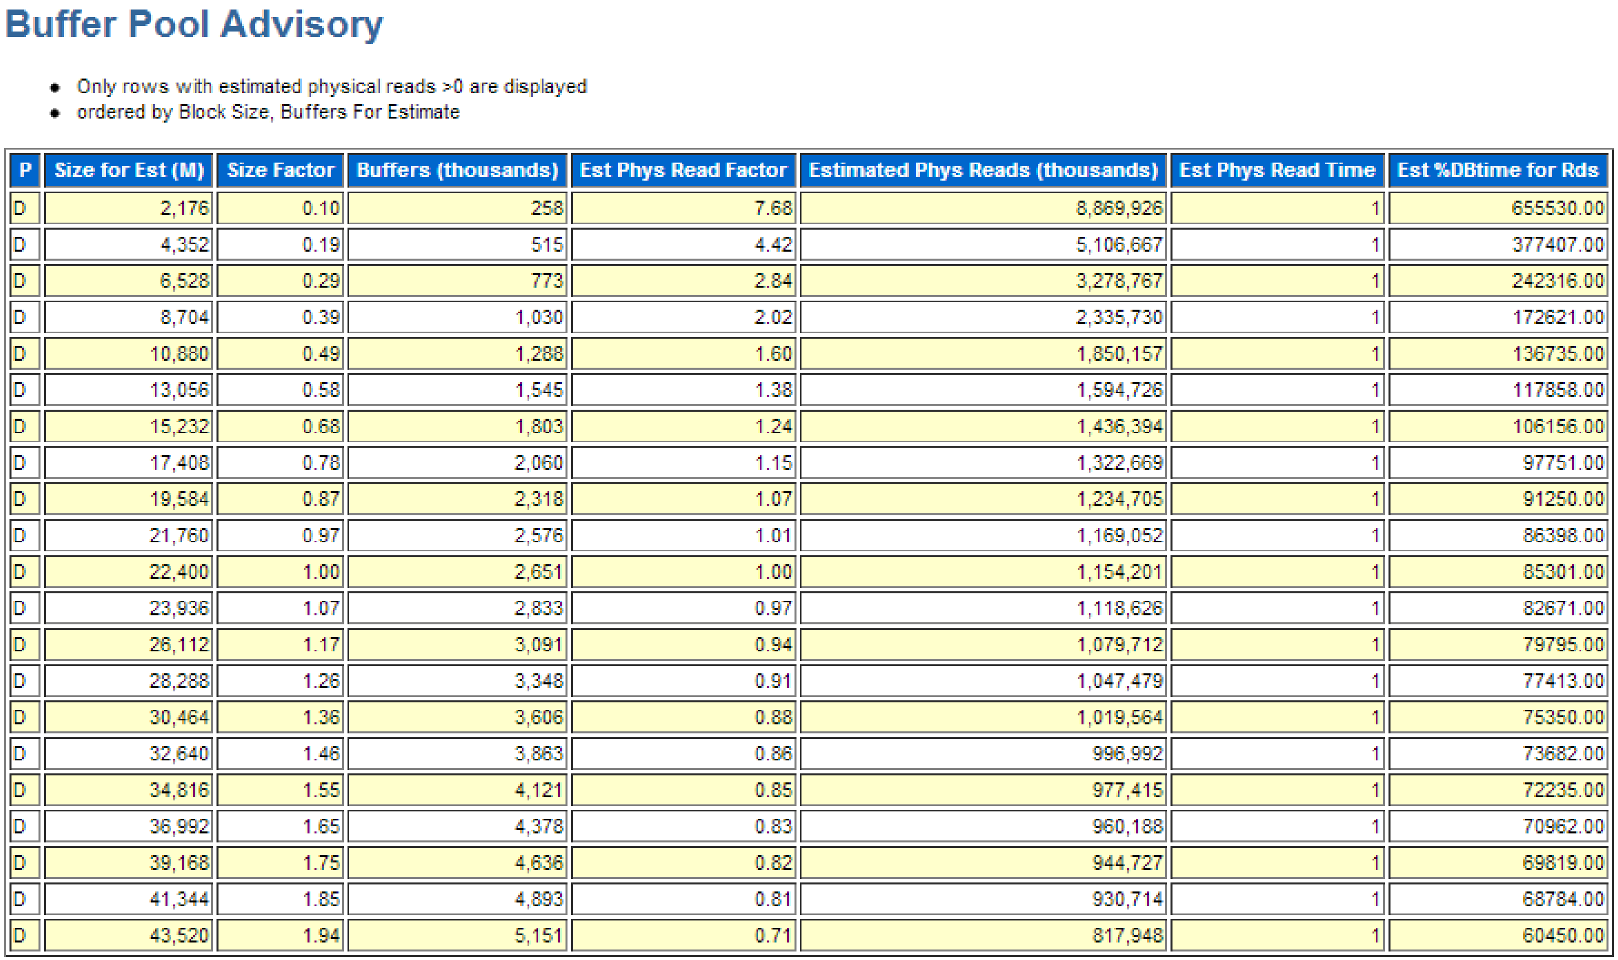

Buffer Pool Advisory

- Only rows with estimated physical reads >0 are displayed
- ordered by Block Size, Buffers For Estimate

P	Size for Est (M)	Size Factor	Buffers (thousands)	Est Phys Read Factor	Estimated Phys Reads (thousands)	Est Phys Read Time	Est %DBtime for Rds
D	2,176	0.10	258	7.68	8,869,926	1	655530.00
D	4,352	0.19	515	4.42	5,106,667	1	377407.00
D	6,528	0.29	773	2.84	3,278,767	1	242316.00
D	8,704	0.39	1,030	2.02	2,335,730	1	172621.00
D	10,880	0.49	1,288	1.60	1,850,157	1	136735.00
D	13,056	0.58	1,545	1.38	1,594,726	1	117858.00
D	15,232	0.68	1,803	1.24	1,436,394	1	106156.00
D	17,408	0.78	2,060	1.15	1,322,669	1	97751.00
D	19,584	0.87	2,318	1.07	1,234,705	1	91250.00
D	21,760	0.97	2,576	1.01	1,169,052	1	86398.00
D	22,400	1.00	2,651	1.00	1,154,201	1	85301.00
D	23,936	1.07	2,833	0.97	1,118,626	1	82671.00
D	26,112	1.17	3,091	0.94	1,079,712	1	79795.00
D	28,288	1.26	3,348	0.91	1,047,479	1	77413.00
D	30,464	1.36	3,606	0.88	1,019,564	1	75350.00
D	32,640	1.46	3,863	0.86	996,992	1	73682.00
D	34,816	1.55	4,121	0.85	977,415	1	72235.00
D	36,992	1.65	4,378	0.83	960,188	1	70962.00
D	39,168	1.75	4,636	0.82	944,727	1	69819.00
D	41,344	1.85	4,893	0.81	930,714	1	68784.00
D	43,520	1.94	5,151	0.71	817,948	1	60450.00

PGA (how am I doing)

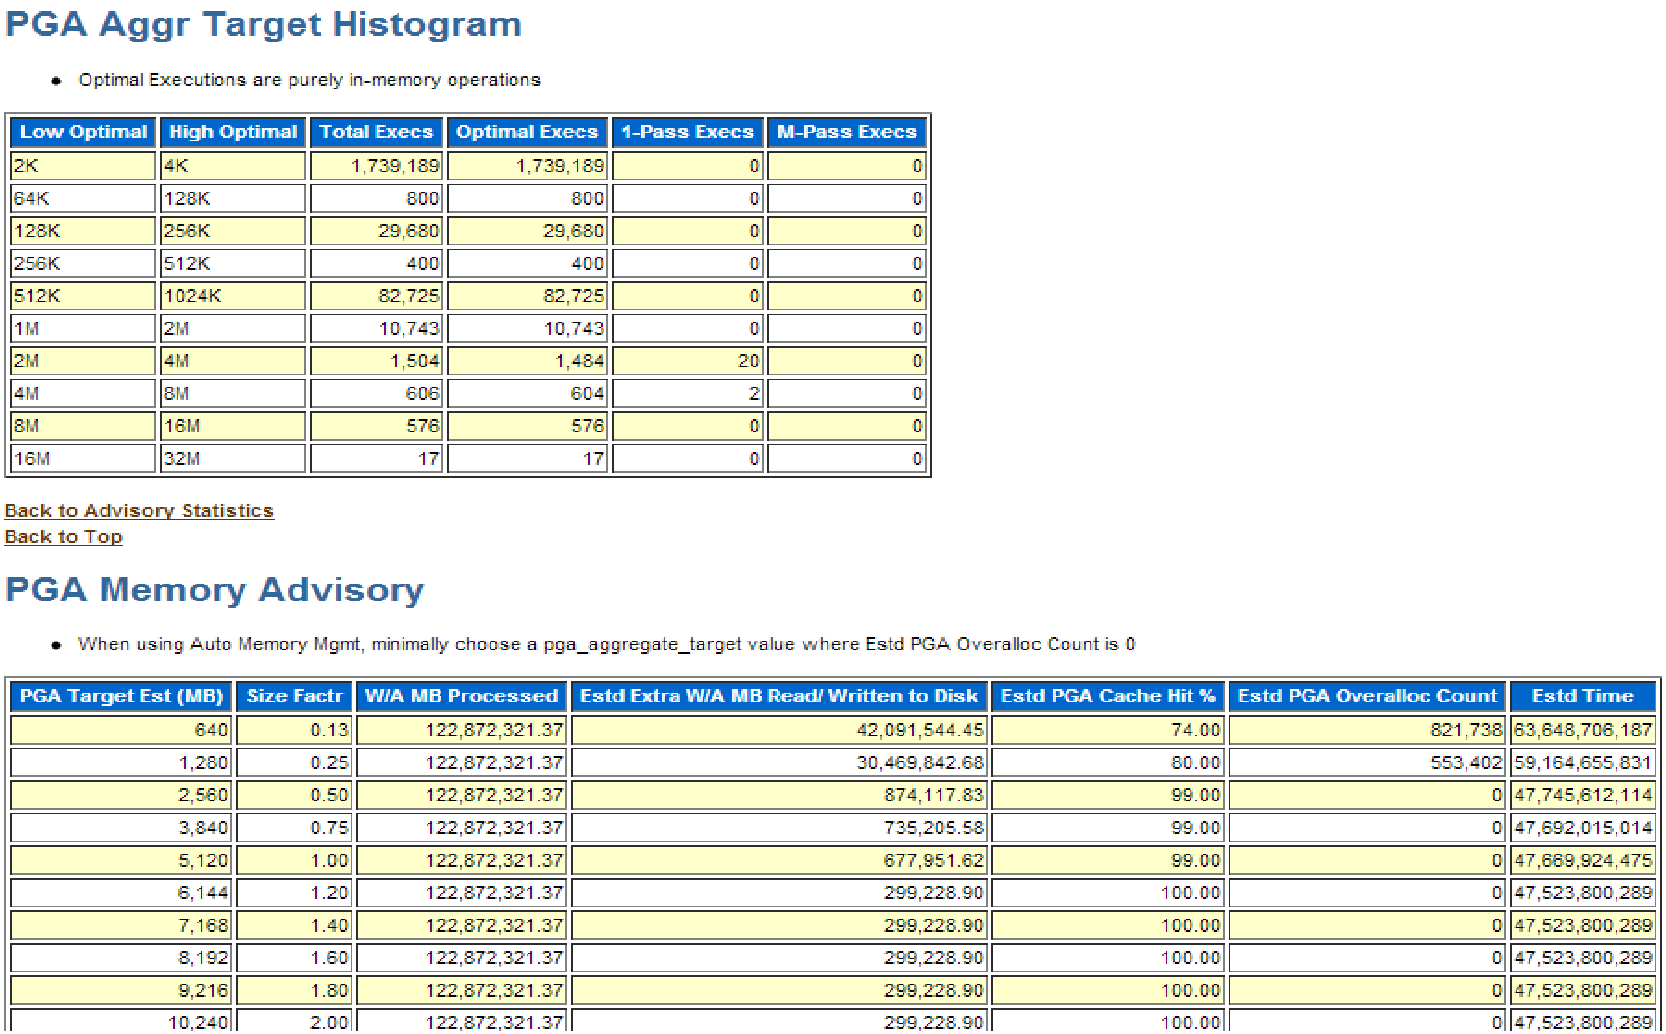

PGA Aggr Target Histogram

- Optimal Executions are purely in-memory operations

Low Optimal	High Optimal	Total Execs	Optimal Execs	1-Pass Execs	M-Pass Execs
2K	4K	1,739,189	1,739,189	0	0
64K	128K	800	800	0	0
128K	256K	29,680	29,680	0	0
256K	512K	400	400	0	0
512K	1024K	82,725	82,725	0	0
1M	2M	10,743	10,743	0	0
2M	4M	1,504	1,484	20	0
4M	8M	606	604	2	0
8M	16M	576	576	0	0
16M	32M	17	17	0	0

Back to Advisory Statistics
Back to Top

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	Estd Time
640	0.13	122,872,321.37	42,091,544.45	74.00	821,738	63,648,706,187
1,280	0.25	122,872,321.37	30,469,842.68	80.00	553,402	59,164,655,831
2,560	0.50	122,872,321.37	874,117.83	99.00	0	47,745,612,114
3,840	0.75	122,872,321.37	735,205.58	99.00	0	47,692,015,014
5,120	1.00	122,872,321.37	677,951.62	99.00	0	47,669,924,475
6,144	1.20	122,872,321.37	299,228.90	100.00	0	47,523,800,289
7,168	1.40	122,872,321.37	299,228.90	100.00	0	47,523,800,289
8,192	1.60	122,872,321.37	299,228.90	100.00	0	47,523,800,289
9,216	1.80	122,872,321.37	299,228.90	100.00	0	47,523,800,289
10,240	2.00	122,872,321.37	299,228.90	100.00	0	47,523,800,289

Pga advice

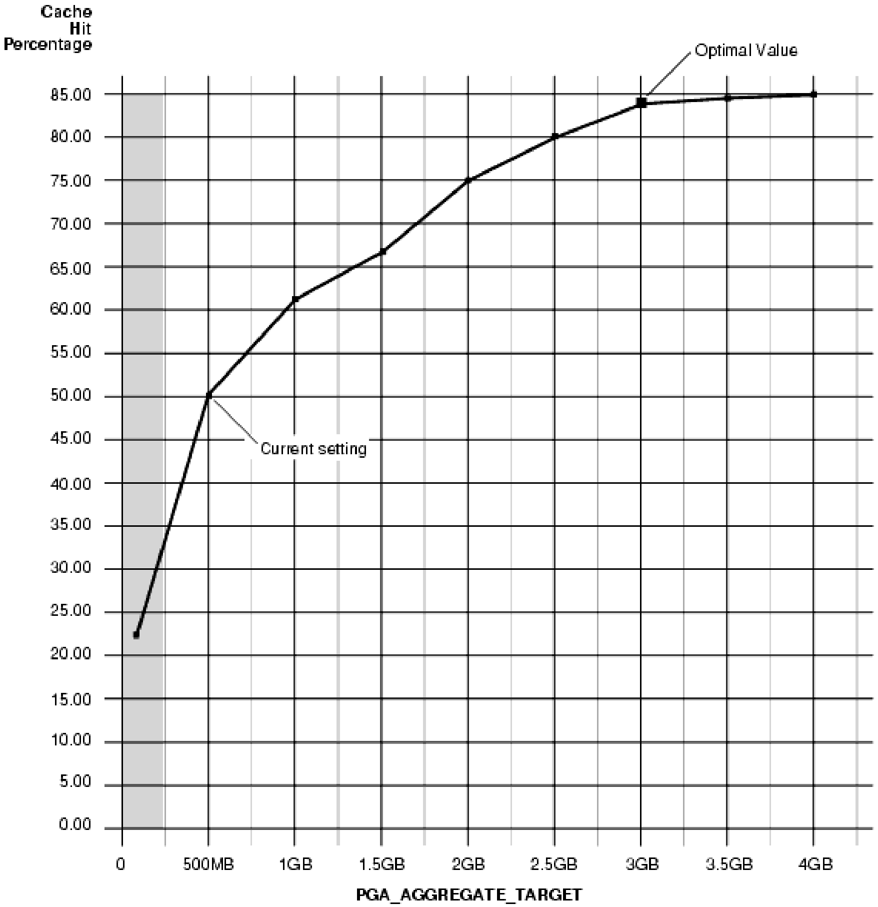
Cache
Hit
Percentage
Optimal Value
85.00
80.00
75.00
70.00
65.00
60.00
55.00
50.00
45.00
40.00
35.00
30.00
25.00
20.00
15.00
10.00
5.00
0.00
Current setting
0
500MB
1GB
1.5GB
2GB
2.5GB
3GB
3.5GB
4GB
PGA_AGGREGATE_TARGET

Huge pages

- To huge page or not ?
- vm.nr_hugepages = 0 -- set to 0 or set to correct values

- Pros and cons of Huge pages

Using large pages

- cat /proc/sys/vm/nr_hugepages 0 To view the current setting using the sysctl command:
- # sysctl vm.nr_hugepages vm.nr_hugepages = 0 To set the number of huge pages using /proc entry:

- !grep -i huge /proc/meminfo
- HugePages_Total: 800
- HugePages_Free: 800
- HugePages_Rsvd: 0
- Hugepagesize: 2048 kB

use_large_pages

- Only –won’t start
- True – eithor (default)

Cursor sharing

- FORCE -Forces statements that may differ in some literals, but are otherwise identical, to share a cursor, unless the literals affect the meaning of the statement.
- SIMILAR - Causes statements that may differ in some literals, but are otherwise identical, to share a cursor, unless the literals affect either the meaning of the statement or the degree to which the plan is optimized.
- EXACT - Only allows statements with identical text to share the same cursor.

Result cache

- Anyone use this ?

Contact me

Email - Bryan.grenn@oracle.com

Twitter - @bryangrenn

Blog http://bryangrenn.blogspot.com

ORACLE®

Uma or NUMA … what does this mean

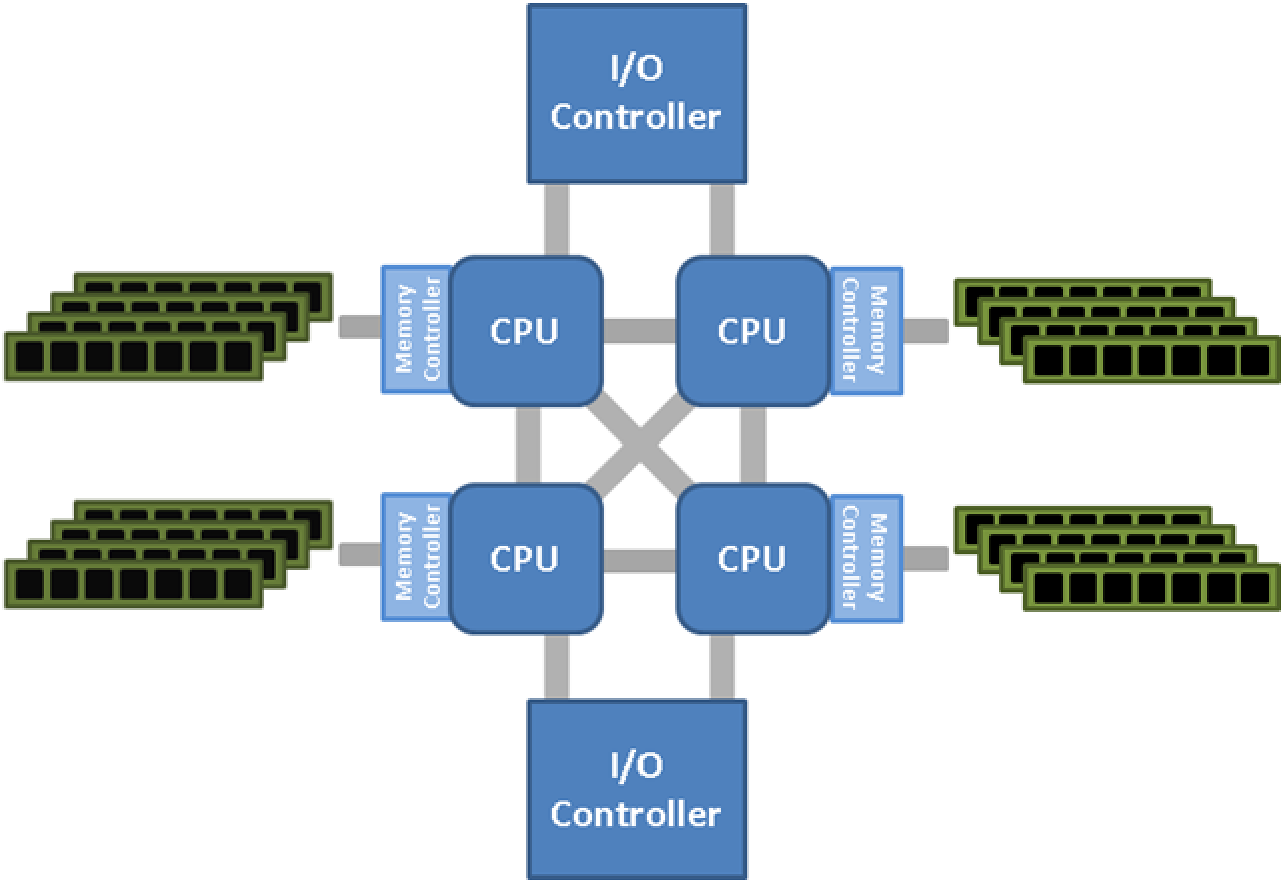
I/O Controller
Memory Controller
CPU
CPU
Memory Controller
Memory Controller
CPU
CPU
Memory Controller
I/O Controller

Assign CPU’s to specific processes with Numactl

- Boot the database under numactl as follows:

  ```
  $ numactl --cpunodebind=0,1,2,3 --interleave=0,1,2,3 sqlplus '/ as sysdba' blah blah blah
  ```

- http://linux.die.net/man/8/numactl