Understanding and Leveraging the Oracle9*i* Advisories

Azeem Mohamed
Product Marketing Manager
Quest Software

QUEST SOFTWARE

Agenda

- Overview of the Oracle9i advisories
  - Buffer cache advisory
  - Shared Pool advisory
  - PGA advisory
- Optimizing a specific advisory
- Optimizing overall memory configuration
- Support for advisories in Quest Central

Overview of the advisories

- The advisories show you how much resource you would save if you changed the size of various areas of the SGA or the PGA
  - Shared pool
  - Buffer pools
  - PGA
- New dynamic parameters in Oracle9i allow you to resize these areas dynamically.

Dynamic SGA

- Not available in 8i or earlier
- Can be used to modify buffer cache, shared pool and large pool size parameters online in Oracle 9i
  - Uses Alter System Command
- Oracle provides the advisories to assist the DBA in managing these settings

Dynamic SGA Parameters

- Buffer cache has new parameters
  - DB_CACHE_SIZE, DB_KEEP_SIZE and DB_RECYCLE_SIZE which will replace the old parameters DB_BLOCK_BUFFERS, BUFFER_POOL_KEEP, BUFER_POOL_RECYCLE, which are still supported, bur aren't dynamic.
  - DB_nK_CACHE_SIZE parameter is used to handle non-standard block size caches.
  - Parameters for shared and large pool are remaining unchanged.

SGA_MAX_SIZE

- The total sum of the SGA cannot be greater than SGA_MAX_SIZE.
- This parameter is not dynamic.
- Consider setting SGA_MAX_SIZE to a larger value than your actual SGA when creating a database.
- Note that Oracle may allocate virtual memory for SGA_MAX_SIZE even if your actual SGA is smaller
- There is no limit on the PGA size.

Granules

- In order to do better and dynamic memory management, Oracle uses a new unit of memory allocation called Granule.
- This means that memory is allocated and freed in terms of granules.
- Granule size is determined by SGA_MAX_SIZE value and is platform dependent.

Granule Measurement

- On most UNIX systems, granule size is 4 MB if SGA_MAX_SIZE is less than 128 MB, otherwise it would be 16 MB
- On Windows, the largest granule size is 8 MB if SGA_MAX_SIZE is greater than 128 MB, otherwise it would be 4 MB

Granule Notes

- A granule cannot be shared among multiple SGA components
- All component sizes are rounded up to multiples of granule size
- Out of main four divisions of memory:
  - Database buffers and variable size components are always allocated in granules.
  - Fixed size, redo buffers are still allocated in bytes.

Oracle9*i* - Advisories

- In Release 2, Oracle provided the following two advisories for dynamic SGA tuning:
- 1. Buffer cache advisory
- 2. Shared pool advisory

Advisory Initialization

- All the advisories in Oracle9i Release 2 including the Buffer Cache Advisor are controlled by a newly introduced parameter STATISTICS_LEVEL.
- By default, the STATISTICS_LEVEL parameter is set to TYPICAL thereby enabling all the advisories.
- V$STATISTICS_LEVEL lists the status of the statistics or advisories controlled by the STATISTICS_LEVEL initialization parameter.
- Each row of V$STATISTICS_LEVEL represents one of these statistics or advisories.

The buffer pool in Oracle9i

- In Oracle9i, blocks may be allocated to one of 7 pools (depending on your default block size):
  - Default
  - Keep
  - Recycle
  - 2k block cache
  - 4k block cache
  - 16k block cache
  - 32k block cache
- Each of these caches can be dynamically resized

Quest Central For Oracle - [SQL Editor (1) Connected to MYORACLE.WORLD:MEL601301 as qco ('C:\sql\db_cache_advice.sql')]

File Edit View Actions Tools Shortcuts Configure Window Help

<No Filter>

SQL Spool

```
1 select name, size_for_estimate,estd_physical_reads
2   from v$db_cache_advice order by name, size_factor
```

[1]: Statement processed in 0.01 sec

Live Update

NAME	SIZE_FOR_ESTIMATE	ESTD_PHYSICAL_READS
DEFAULT	8	48363
DEFAULT	16	17092
DEFAULT	24	17092
DEFAULT	32	17092
DEFAULT	40	17092
DEFAULT	48	17092
DEFAULT	56	17092
DEFAULT	64	17092
DEFAULT	72	17092
DEFAULT	80	17092
DEFAULT	88	17092
DEFAULT	96	17092
DEFAULT	104	17092
DEFAULT	112	17092
DEFAULT	120	17092
DEFAULT	128	17092
DEFAULT	136	17092
DEFAULT	144	17092
DEFAULT	152	17092
DEFAULT	160	17092
KEEP	8	10783
KEEP	16	10783
KEEP	24	3230
KEEP	32	3230
KEEP	40	3230

db_cache_advice.sq 1:1 Lines: 2 Top Line: 1 Modified Insert

Quest Ce... Database... SQL Edito... SQLTUNE... SQL Edito...

Buffer Cache Advisory

- This advisory has V$DB_CACHE_ADVICE as its advisory view and it contains 20 rows that predict the number of physical reads for the cache size corresponding to the row.
- This advisory predicts the change in I/Os as the buffer cache size is changed from 10% to 200% of the current size.
- Cache size is determined by DB_CACHE_SIZE parameter at the instance startup. Its value is adjusted to the nearest multiples of granule size.

Shared Pool Advisory

- This advisory advices on about library cache memory and predicts the effect of altering the shared pool size on the total amount parsing activities in the system.
- This advisory has V$SHARED_POOL_ADVICE as its advisory table.
- Parse time saved refers to the amount of time saved by keeping library cache memory objects in the shared pool, as opposed to having to reload these objects.

Shared Pool Advisory

- Note: changing the shared pool size may cause latch contention and/or fragmentation and the advisory does not take this into account.
- Reducing the size of the shared pool can take some time, and at Quest we have seen some minor corruption in internal memory structures result

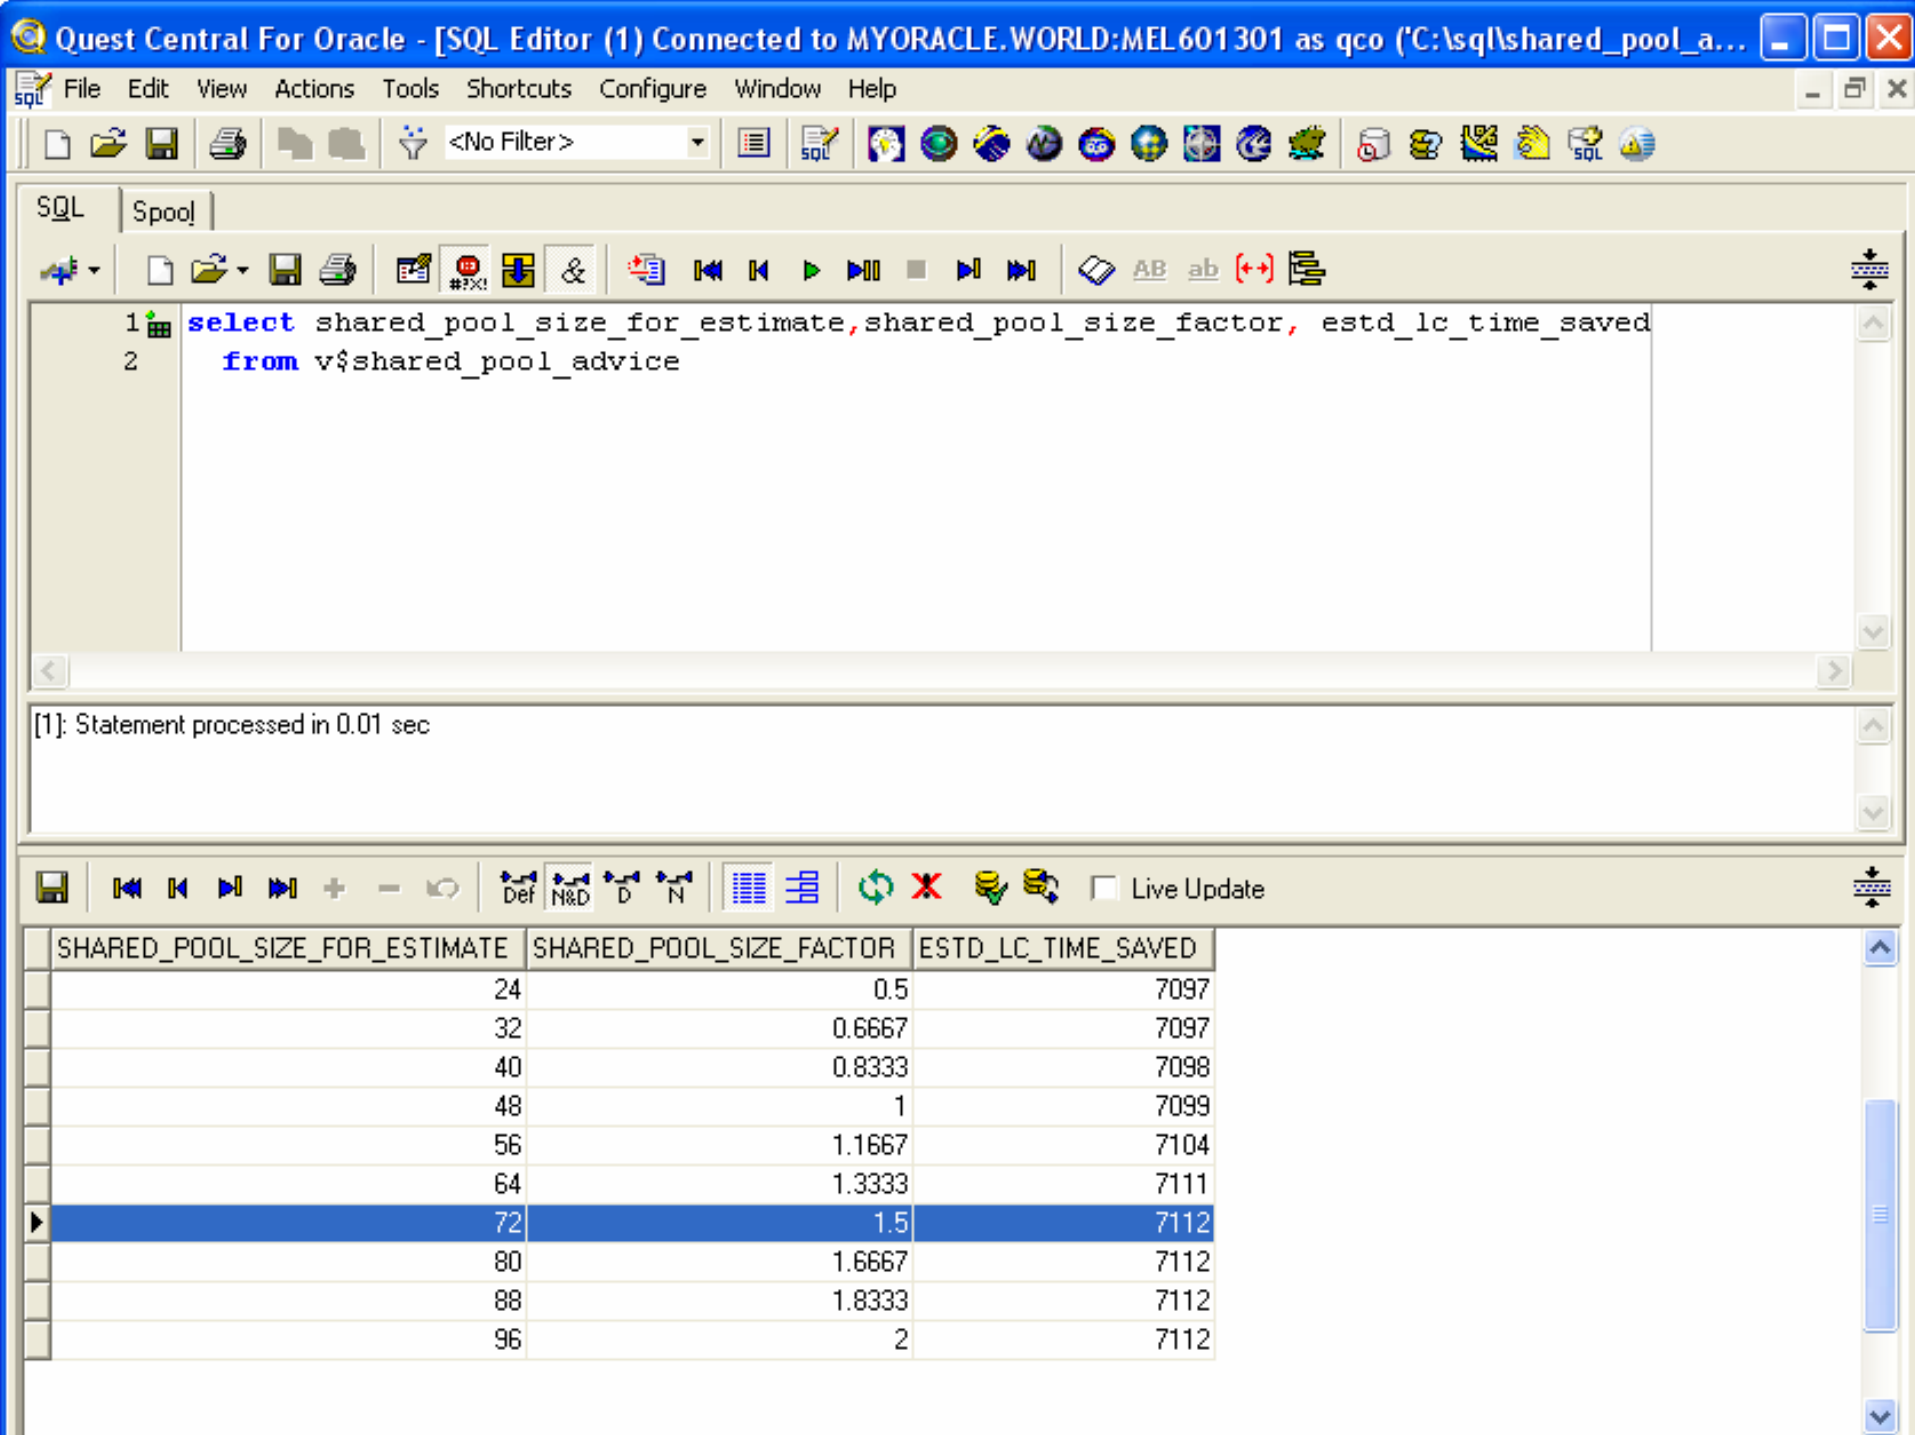
Quest Central For Oracle - [SQL Editor (1) Connected to MYORACLE.WORLD:MEL601301 as qco ('C:\sql\shared_pool_a...
File Edit View Actions Tools Shortcuts Configure Window Help
<No Filter>
SQL Spool
1 select shared_pool_size_for_estimate,shared_pool_size_factor, estd_lc_time_saved
2 from v$shared_pool_advice
[1]: Statement processed in 0.01 sec
Def N&D D N
Live Update
SHARED_POOL_SIZE_FOR_ESTIMATE SHARED_POOL_SIZE_FACTOR ESTD_LC_TIME_SAVED
24 0.5 7097
32 0.6667 7097
40 0.8333 7098
48 1 7099
56 1.1667 7104
64 1.3333 7111
72 1.5 7112
80 1.6667 7112
88 1.8333 7112
96 2 7112

PGA

- Besides SGA, the Oracle Database also assigns each server process a private memory region called the “Program Global Area” (PGA).
- A PGA is created for each server process when it is started i.e. when a new session is initiated while using the dedicate server configuration or, when a new shared server process is created.
- It contains data and control information for a server process and, unlike the SGA, each server process has exclusive access to its PGA.

Automatic SQL Memory MGMT

- The Automatic SQL Memory Management feature can be enabled by setting the parameter PGA_AGGREGATE_TARGET to a non-zero value.
- The default value of this parameter is zero meaning that the automatic mode is disabled by default.
- This is primarily to ensure backward compatibility.

Monitoring SQL Memory Use

- **V$PGASTAT**
  - This view should be the primary reference for monitoring PGA memory usage.
  - The information contained in this view includes PGA memory currently used, maximum PGA memory allocated since instance start up and the "**PGA cache hit percentage**".
  - The PGA cache hit percentage is a new concept in Oracle9i defined as the percentage of the total amount of data processed by memory intensive SQL operators which was accommodated in the available PGA memory.

Monitoring SQL Memory Use

- **V$SYSSTAT**
  - The following new statistics have been added to V$SYSSTAT and V$SESSTAT
    - Work area memory allocated (KB)
      - Total amount of PGA memory dedicated to work areas allocated on behalf of a given session (V$SESSTAT) or on the system (V$SYSSTAT).
      - This includes work areas allocated under both MANUAL as well as AUTO mode.
      - For DSS workload, this should represent most of the PGA memory.

Monitoring SQL Memory Use

- **V$SYSSTAT** – Continued
  - Work area executions - optimal size
    - The cumulative count of work areas that were executed in cache mode.
    - A sort area can be said to have an optimal size if it did not need to spill to disk. This applies to hash-join as well

Monitoring SQL Memory Use

- **V$SYSSTAT** – Continued
  - Work area executions - one pass size
    - The cumulative count of work areas using the one pass size.
    - One pass is generally used for big work areas where spilling to disk cannot be avoided.

Monitoring SQL Memory Use

- **V$SYSSTAT** – Continued
  - Work area executions - multipasses size
    - The cumulative count of work areas running in more than one pass.
    - This should be avoided and may be a symptom of a poorly tuned system.

Monitoring SQL Memory Use

- **V$PROCESS**
- The following four new columns have been added to the V$PROCESS view to report the PGA memory allocated and used by an Oracle process.
  - PGA_USED_MEM PGA
    - Memory currently used by the process.
  - PGA_ALLOC_MEM
    - It is the PGA memory currently allocated by the process. It includes free PGA memory not yet released to the OS by the server process.
  - PGA_FREEABLE_MEM
    - Part of the allocated PGA memory that can be freed.
  - PGA_MAX_MEM
    - The maximum PGA memory ever allocated by the process.

Monitoring SQL Memory Use

- **V$SQL_WORKAREA**
  - This view displays information about work areas used by SQL cursors. Each SQL statement stored in the shared pool has one or more child cursors which are listed in the V$SQL dynamic view.
- **V$SQL_WORKAREA_ACTIVE**
  - This view is structurally same as V$SQL_WORKAREA. However, while the latter contains work area for all cached cursors, V$SQL_WORKAREA_ACTIVE view only list those which are currently active.

PGA advisory

- In Oracle9i, SORT_AREA_SIZE, HASH_AREA_SIZE and other parameters that affect the size of the PGA have been superceded by PGA_AGGREGATE_TARGET.
- The PGA_AGGREGATE_TARGET is the total amount of system memory that can be allocated to the PGA of all sessions.
- The PGA advisory shows you how much IO to the temporary tablespace would be saved for different settings of PGA_AGGREGATE_TARGET.

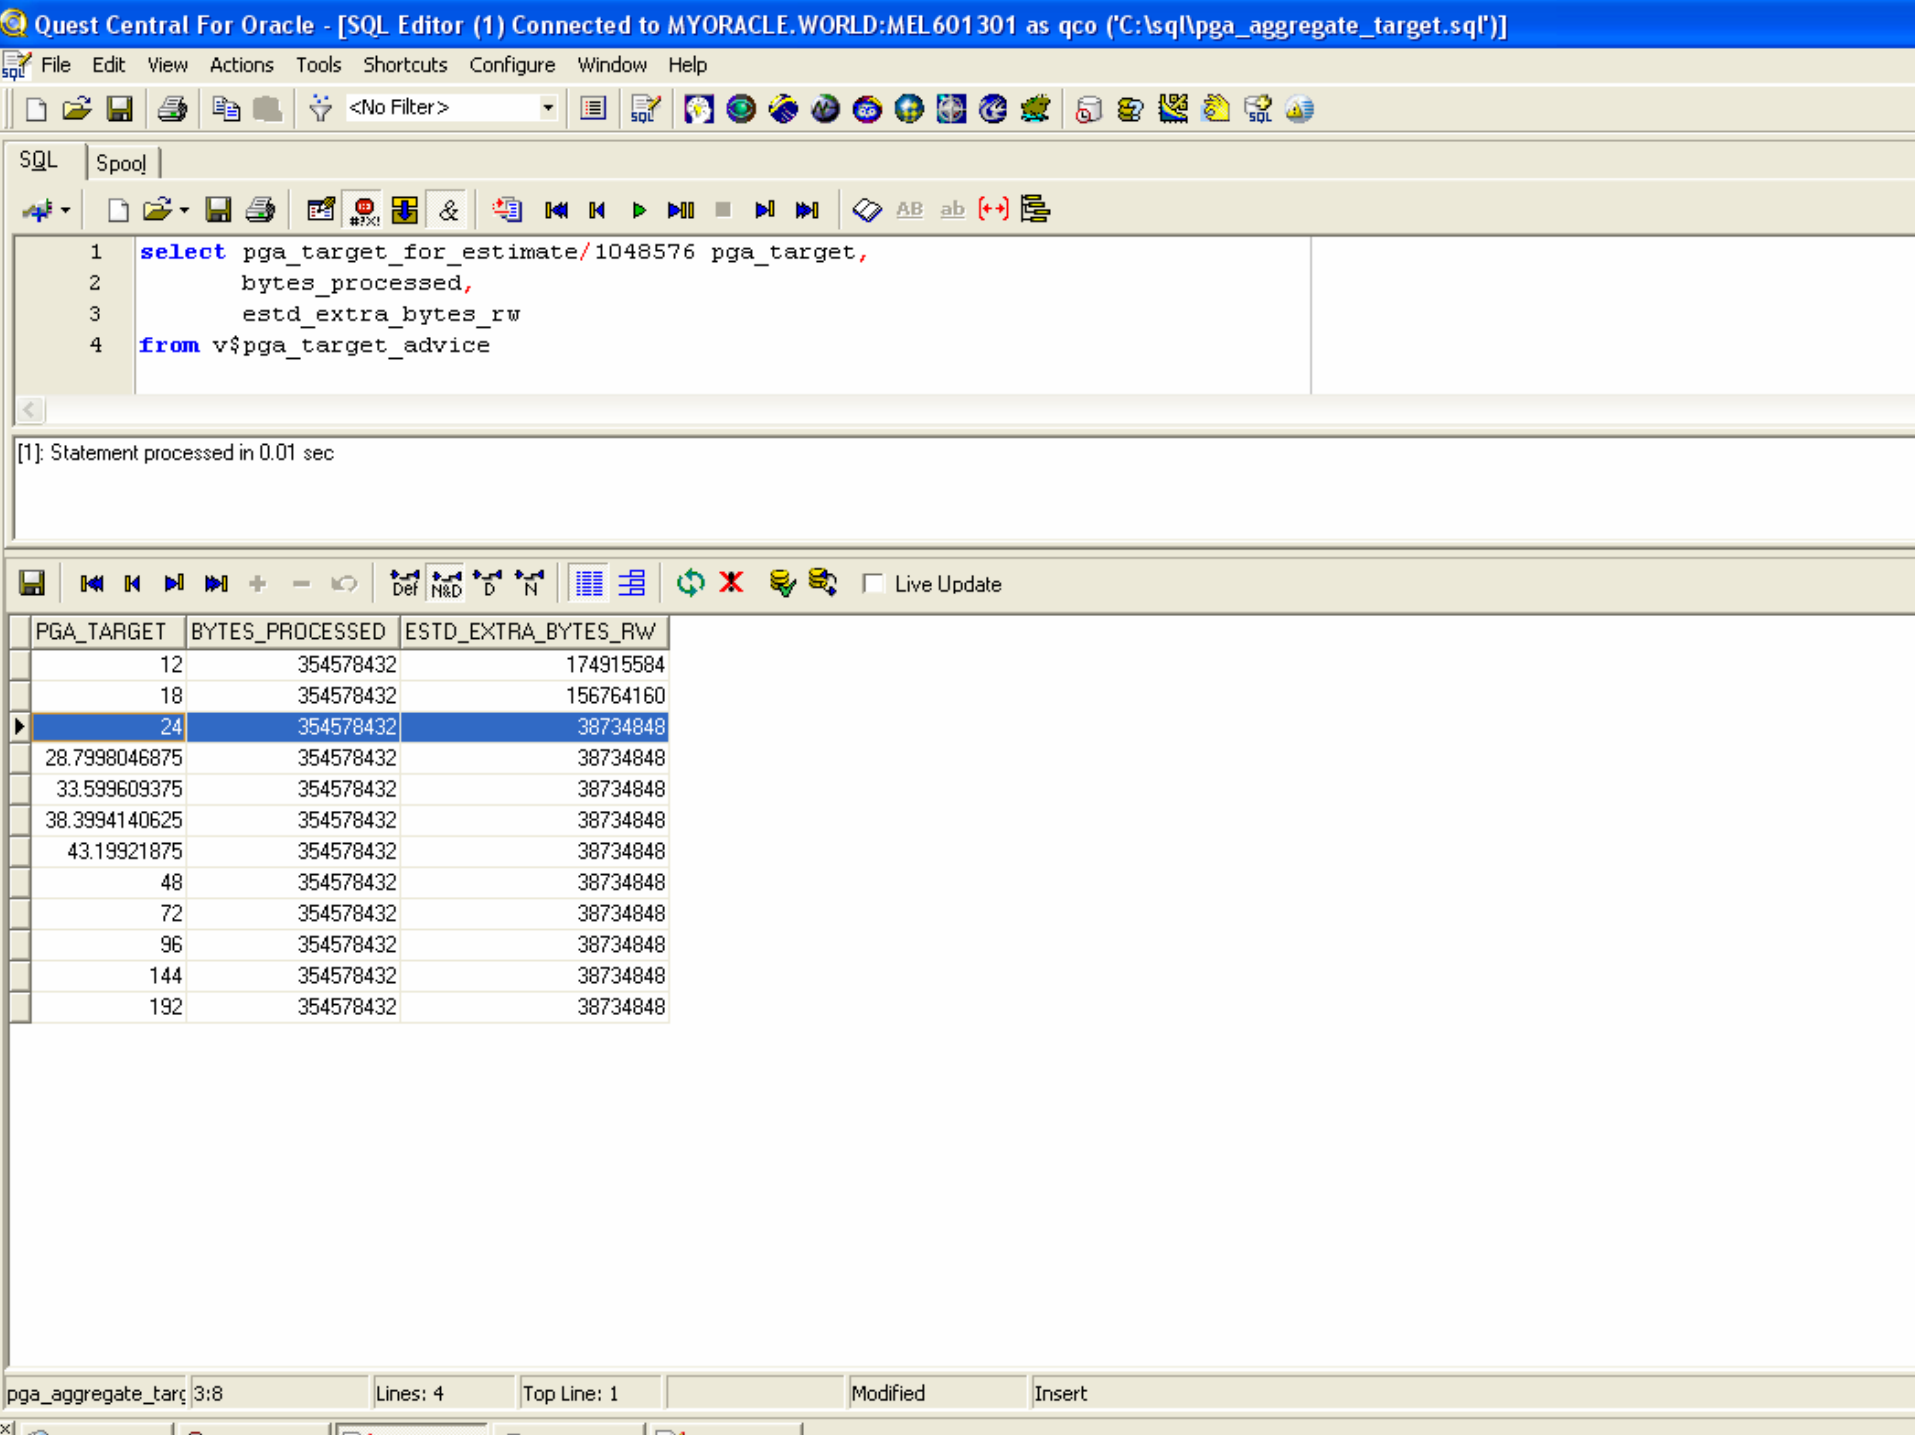
Quest Central For Oracle - [SQL Editor (1) Connected to MYORACLE.WORLD:MEL601301 as qco ('C:\sql\pga_aggregate_target.sql')]
File Edit View Actions Tools Shortcuts Configure Window Help
<No Filter>
SQL Spool
1 select pga_target_for_estimate/1048576 pga_target,
2 bytes_processed,
3 estd_extra_bytes_rw
4 from v$pga_target_advice
[1]: Statement processed in 0.01 sec
Live Update
PGA_TARGET BYTES_PROCESSED ESTD_EXTRA_BYTES_RW
12 354578432 174915584
18 354578432 156764160
24 354578432 38734848
28.7998046875 354578432 38734848
33.599609375 354578432 38734848
38.3994140625 354578432 38734848
43.19921875 354578432 38734848
48 354578432 38734848
72 354578432 38734848
96 354578432 38734848
144 354578432 38734848
192 354578432 38734848
pga_aggregate_targ 3:8 Lines: 4 Top Line: 1 Modified Insert

Using the advisories in combination

- What we really want to know is how to allocate all our available memory to best effect.
- Unfortunately, it is hard to work this out because each advisory reports in different units
  - Shared pool in seconds saved
  - Buffer pool in IO savings
  - PGA in bytes of IO saved

Shared Pool Size Advice - M0920_64

Shared Pool Size Advice

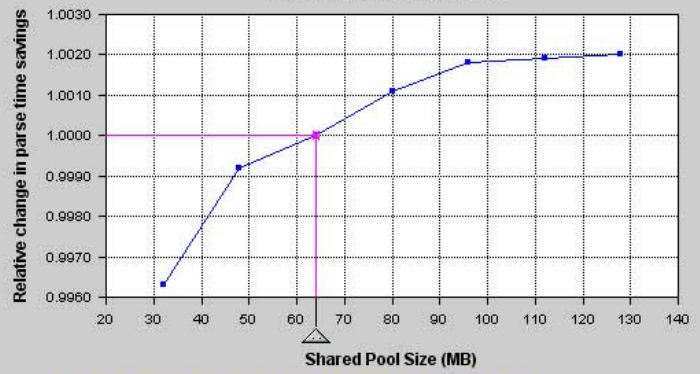

- Change in elapsed parse time savings for various sizes of Shared Pool
- Current Shared Pool size

Shared Pool Size: 64 MB

Estimated Time Saved

Size for estimate (MB)	Size Factor	Time Saved (Seconds)	Time Saved Factor
32	.5	29594	.9963
48	.75	29679	.9992
64	1	29704	1
80	1.25	29738	1.0011
96	1.5	29756	1.0018
112	1.75	29761	1.0019

Refresh OK Cancel Help

PGA Target Advice - M0920_64

PGA Aggregate Target Advice

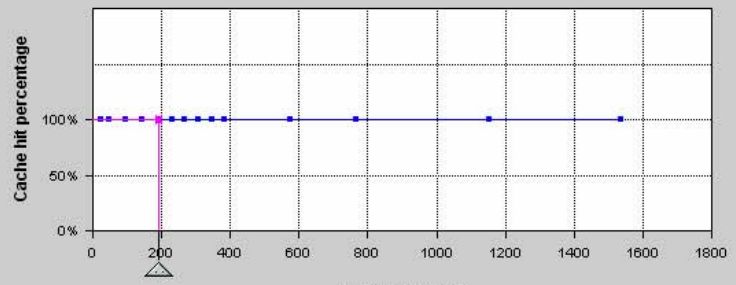

- Variation of cache hit r...
- Current PGA target
- Overflow range

PGA Target: 192 MB

Memory usage details...

Estimated cache hit perce...

PGA target for estimate (MB)
24
48
96
144
192
220

Buffer Cache Size Advice - M0920_64

Buffer Cache Size Advice

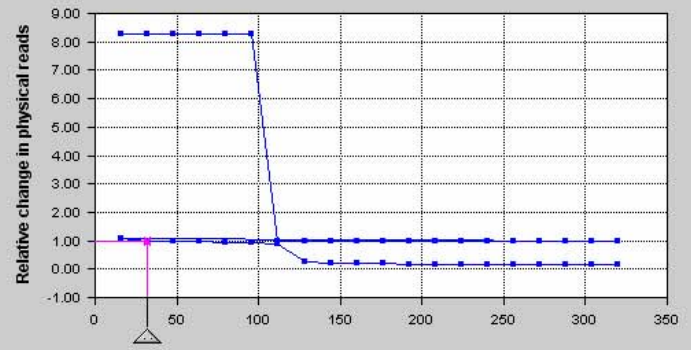

- Change in physical reads for various cache sizes
- Current cache size

Buffer Cache Size: 32 MB

Estimated Physical Reads

Pool	Block Size	Size For Esti...	Buffers For E...	Estd Physical ...	Estd Physical ...
KEEP	8192	16	1985	8.2539	66997
KEEP	8192	32	3970	8.2539	66997
KEEP	8192	48	5955	8.2539	66997
KEEP	8192	64	7940	8.2539	66997
KEEP	8192	80	9925	8.2539	66997
KEEP	8192	96	11910	8.2539	66997

Refresh OK Cancel Help

Configuration
Sessions
Locks
Resource Consum...
Resource Plans
Resource Plan Sch...

Type a question for help

Click to add notes

Draw AutoShapes

Slide 1 of 28 Default Design English (U.S.)

Start | Microsoft A... | Microsof... | Nelson Emai... | Command P... | 7 sqlplus | latch_test4 ... | 3 Internet... | My Documents | 7 Microsof... | guyh2 | 3 Microsof... | Perlide | Quest Cent... | Adobe Acro... | Windows M... | Microsoft P... | Oracle Ent... | 9:21 AM Wednesday

Using the advisories together

- The trick is to convert the output of the advisories into a common unit and the obvious common unit is time.

- Shared pool is already measured in seconds; we need to convert the db cache and PGA advisories to seconds as well.

Converting the PGA advisory to seconds

- PGA advisory shows number of bytes saved if the PGA aggregate target is increased.
- These bytes are read and written using ‘db file direct read’ and ‘db file direct write’ operations.
- We can work out how many seconds each operation takes using v$system_event.
- We can work out how many bytes are read and written by examining trace file information

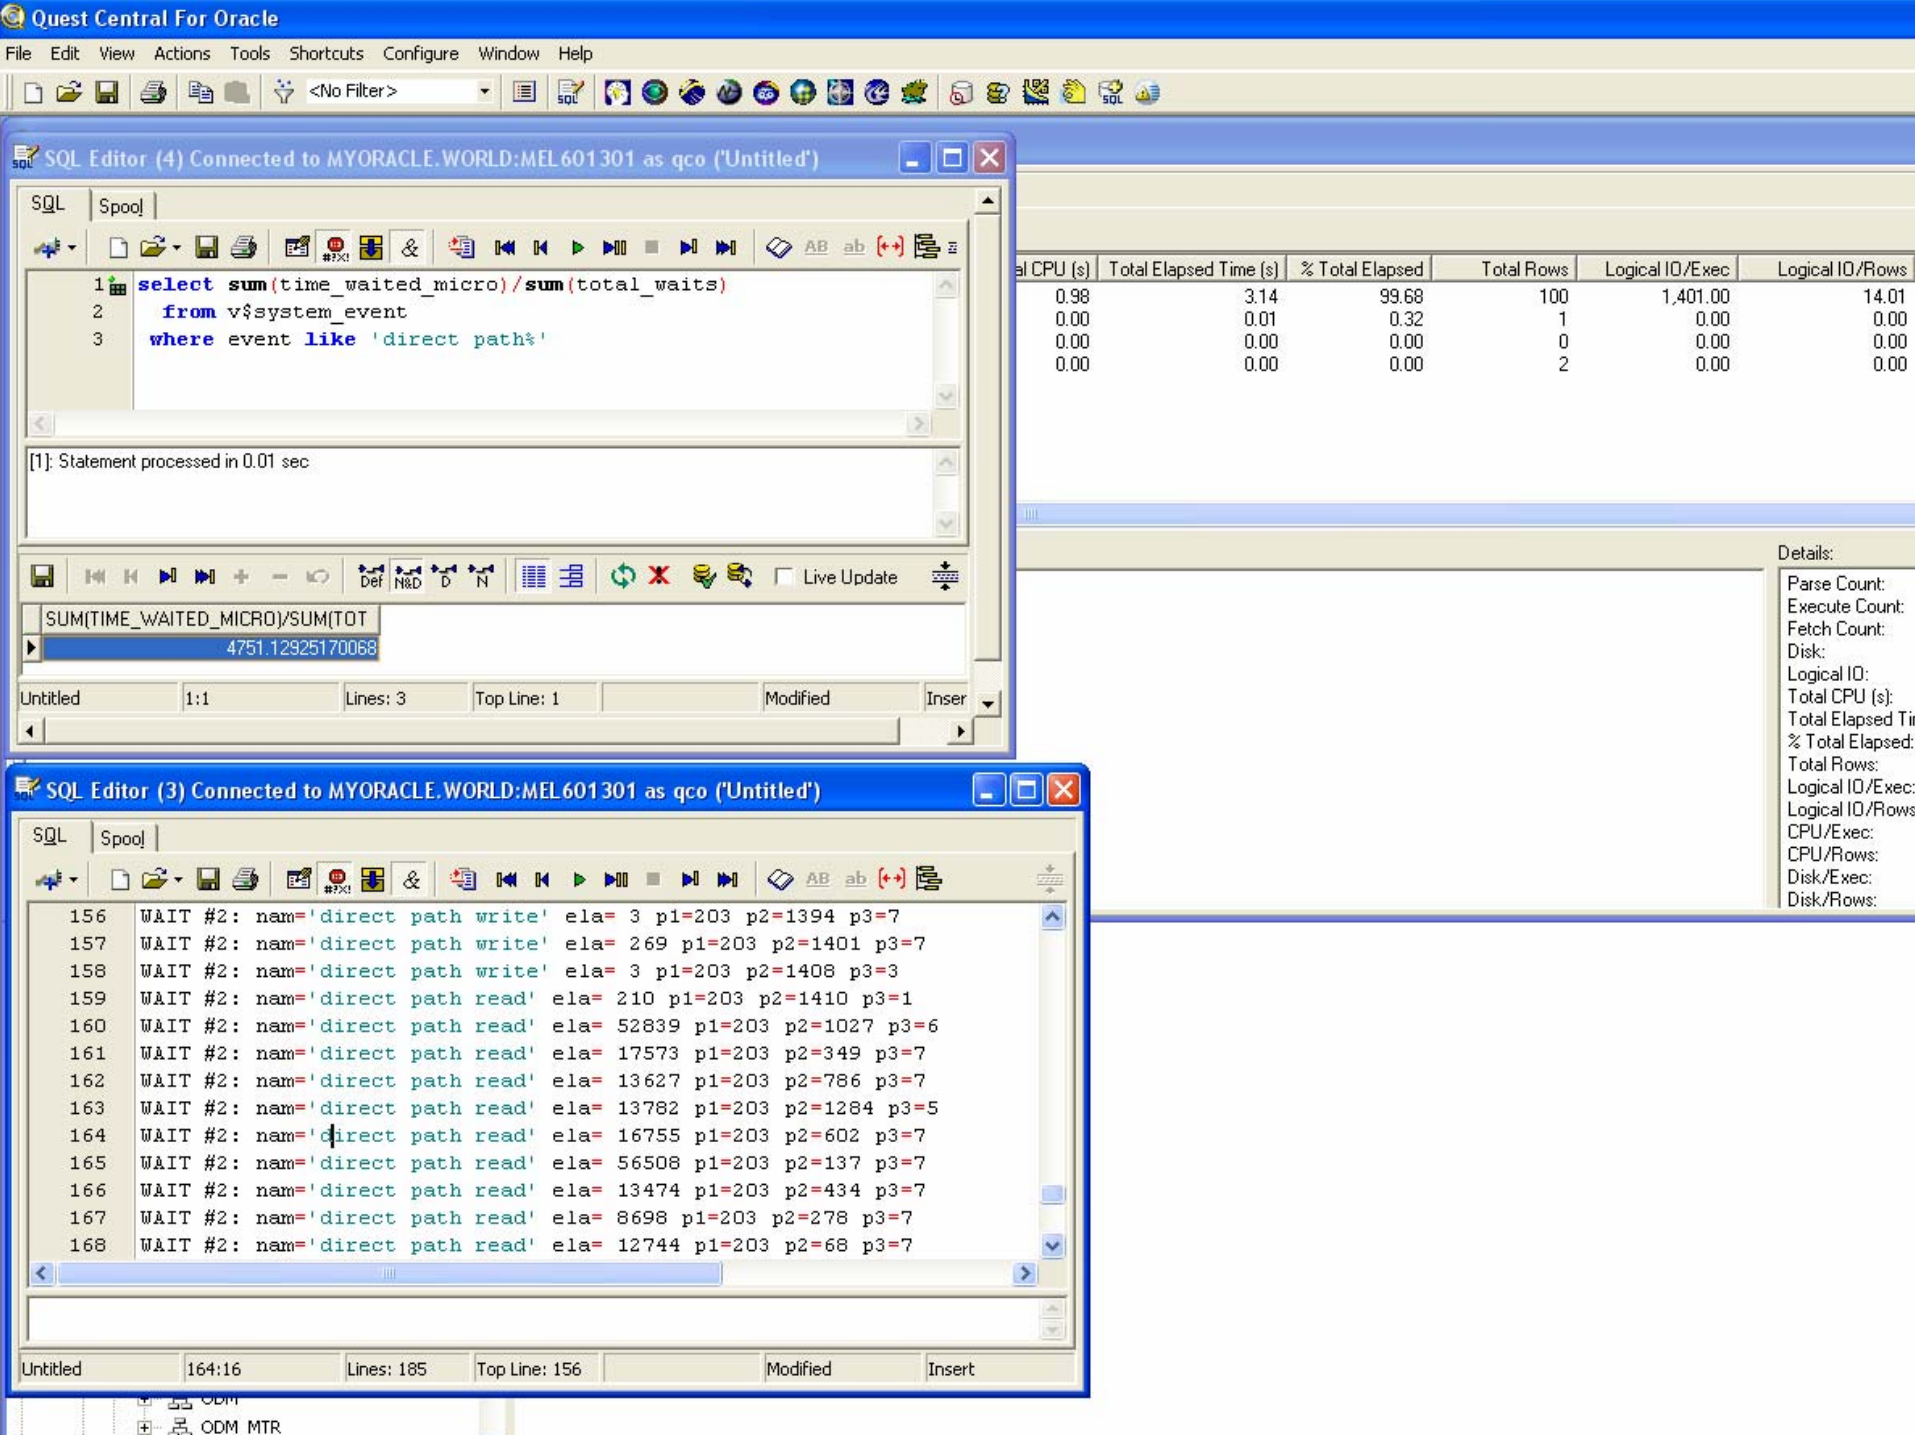
Quest Central For Oracle
File Edit View Actions Tools Shortcuts Configure Window Help
<No Filter>
SQL Editor (4) Connected to MYORACLE.WORLD:MEL601301 as qco ('Untitled')
SQL Spool
1 select sum(time_waited_micro)/sum(total_waits)
2 from v$system_event
3 where event like 'direct path%'
[1]: Statement processed in 0.01 sec
Live Update
SUM(TIME_WAITED_MICRO)/SUM(TOT
4751.12925170068
Untitled 1:1 Lines: 3 Top Line: 1 Modified Inser
SQL Editor (3) Connected to MYORACLE.WORLD:MEL601301 as qco ('Untitled')
SQL Spool
156 WAIT #2: nam='direct path write' ela= 3 p1=203 p2=1394 p3=7
157 WAIT #2: nam='direct path write' ela= 269 p1=203 p2=1401 p3=7
158 WAIT #2: nam='direct path write' ela= 3 p1=203 p2=1408 p3=3
159 WAIT #2: nam='direct path read' ela= 210 p1=203 p2=1410 p3=1
160 WAIT #2: nam='direct path read' ela= 52839 p1=203 p2=1027 p3=6
161 WAIT #2: nam='direct path read' ela= 17573 p1=203 p2=349 p3=7
162 WAIT #2: nam='direct path read' ela= 13627 p1=203 p2=786 p3=7
163 WAIT #2: nam='direct path read' ela= 13782 p1=203 p2=1284 p3=5
164 WAIT #2: nam='direct path read' ela= 16755 p1=203 p2=602 p3=7
165 WAIT #2: nam='direct path read' ela= 56508 p1=203 p2=137 p3=7
166 WAIT #2: nam='direct path read' ela= 13474 p1=203 p2=434 p3=7
167 WAIT #2: nam='direct path read' ela= 8698 p1=203 p2=278 p3=7
168 WAIT #2: nam='direct path read' ela= 12744 p1=203 p2=68 p3=7
Untitled 164:16 Lines: 185 Top Line: 156 Modified Insert
l CPU (s) Total Elapsed Time (s) % Total Elapsed Total Rows Logical IO/Exec Logical IO/Rows
0.98 3.14 99.68 100 1,401.00 14.01
0.00 0.01 0.32 1 0.00 0.00
0.00 0.00 0.00 0 0.00 0.00
0.00 0.00 0.00 2 0.00 0.00
Details:
Parse Count:
Execute Count:
Fetch Count:
Disk:
Logical IO:
Total CPU (s):
Total Elapsed Ti
% Total Elapsed:
Total Rows:
Logical IO/Exec:
Logical IO/Rows
CPU/Exec:
CPU/Rows:
Disk/Exec:
Disk/Rows:
ODM
ODM MTR

Converting the PGA advisory to seconds

- An average direct path IO operation took 4.7 ms on my laptop and the average blocks in each IO was 6.
- My block size is 8K, so the amount of time taken to read each bytes is (4.7)/(6*8*1024)=.09ms
- We can now multiply that figure into the pga advisory to calculate time saved.

Converting the buffer pool advisory to seconds

- This is easier, v$system_event can easily give us the average time per IO:

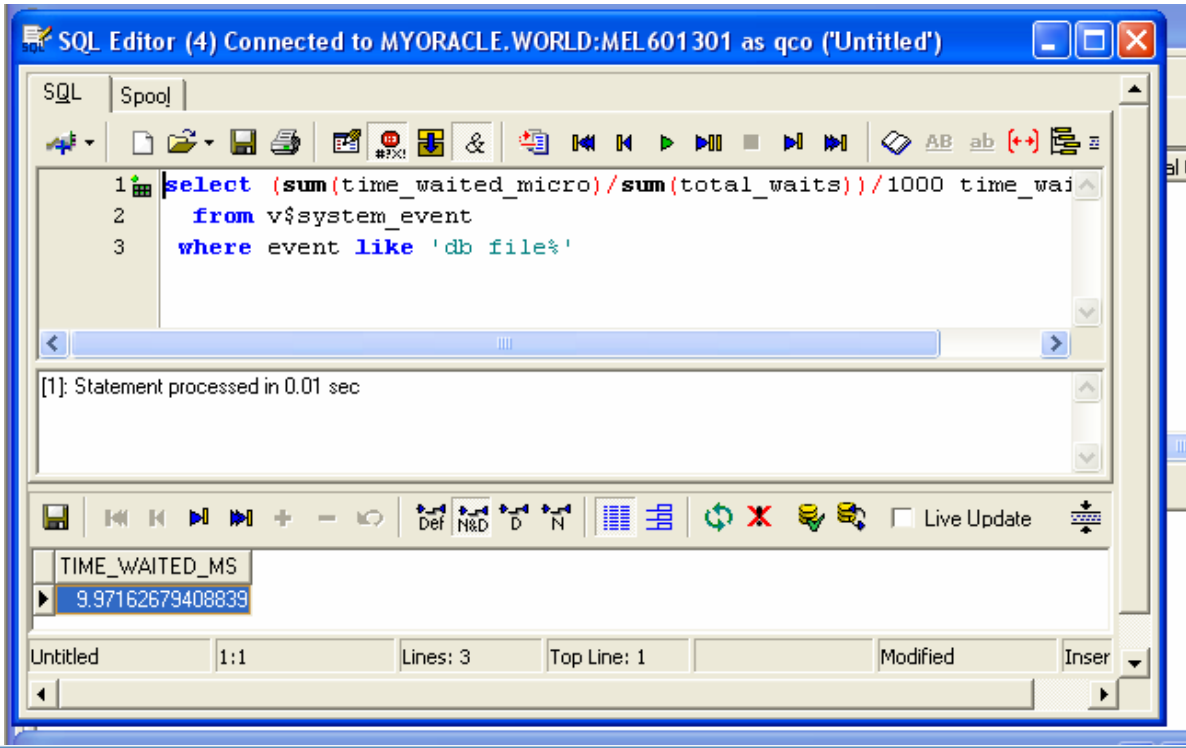

Comparing the advisories

- Finally, convert the advisory data into seconds for the pga and buffer pool advice.
- Find the combination of memory settings that minimize elapsed time but still fit into your available memory
- However, be wary of reducing the size of the shared pool

Advisory support in Quest Central

- Quest Central supports the algorithm outlined above:
  - Coverts advisory data into estimated elapsed time savings
  - Works out the combination of memory settings that minimized elapsed time.
  - Allows you to browse the advisories and work out the best combination.
  - Applies the new settings and lets you see the results

Spotlight on Oracle - mo920_64

File View Favorites Tools Help

mo920_64

Tuning

Performance | 9i Memory Management | Latch

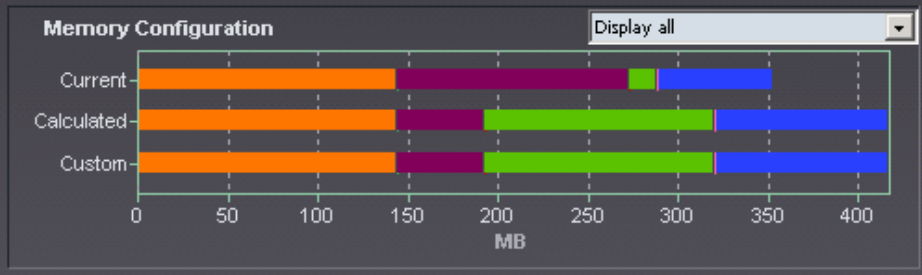

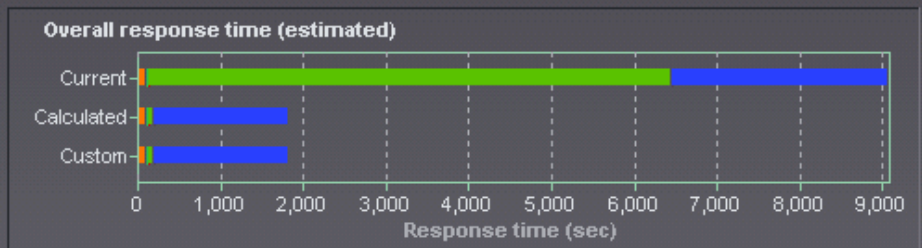

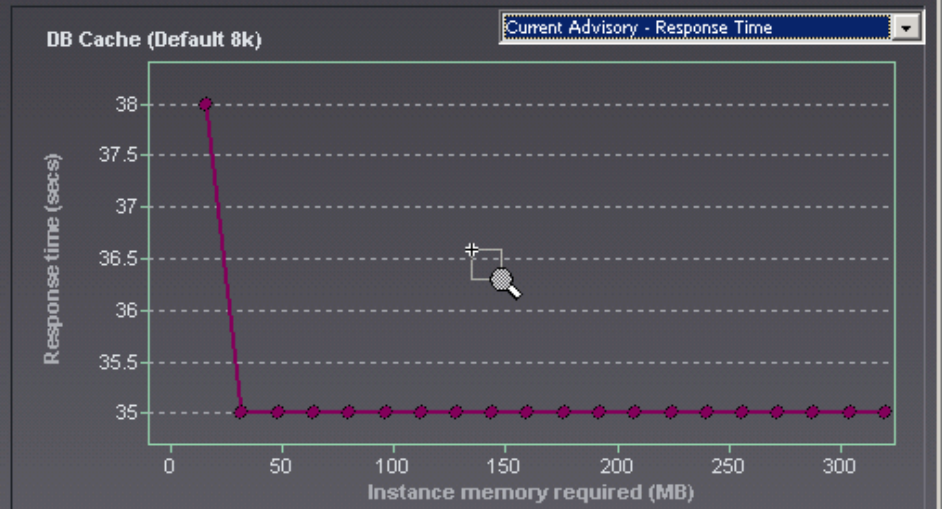

Component	Current	Min.	Max.	Calculated	Custom	Locked	Status
Fixed	0.70						
Variable	144.00			144	144		
Database Buffers	144.00			176	176		
DB 2k Cache	0.00	0	0	0	0		
DB 4k Cache	0.00	0	0	0	0		
DB Cache (Default 8k)	128.00	24	25,600	48	48		
DB 16k Cache	0.00	0	0	0	0		
DB 32k Cache	0.00	0	0	0	0		
DB Keep Cache	16.00	800	256	128	128		
DB Recycle Cache	0.00	0	0	0	0		
Redo Buffers	0.77						
PGA Aggregate Target	64.00	48	192	96	96		
Total	**353.47**			**417.47**	**417.47**		

Max. SGA size :
513.47 MB
Unused SGA memory :
224 MB
Free Physical Memory :
0 MB
Granule size :
16 MB
Statistics level :
TYPICAL
Maximum Instance Memory :
420 MB
Time Recommended :
N/A

Start | Windows Task ... | Rptsched | C:\WINNT\Sys... | Spotlight | Quest Central ... | 2:39 PM

Spotlight on Oracle - mo920_64

File View Favorites Tools Help

mo920_64

Tuning

Performance | 9i Memory Management | Latch

Memory Configuration

Display all

Overall response time (estimated)

DB Keep Cache

Current Advisory - Response Time

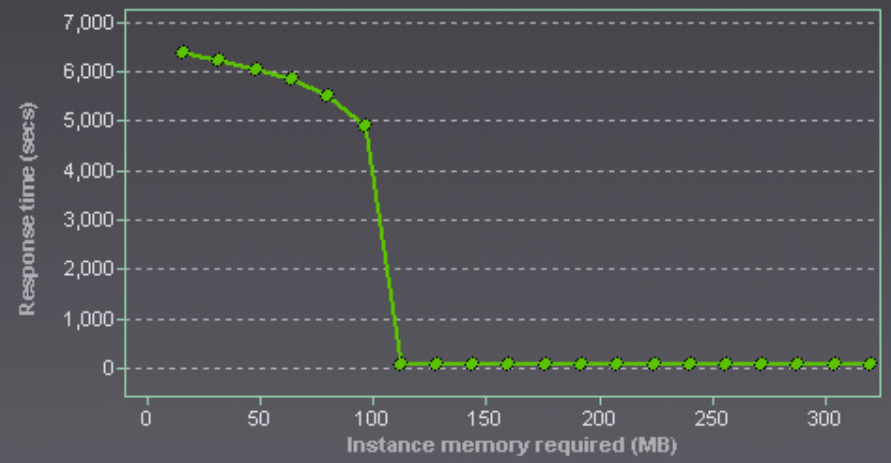

Component	Current	Min.	Max.	Calculated	Custom	Locked	Status
Fixed	0.70						
⊞ Variable	144.00			144	144		
⊟ Database Buffers	144.00			176	176		
DB 2k Cache	0.00	0	0	0	0		
DB 4k Cache	0.00	0	0	0	0		
DB Cache (Default 8k)	128.00	24	25,600	48	48		
DB 16k Cache	0.00	0	0	0	0		
DB 32k Cache	0.00	0	0	0	0		
DB Keep Cache	16.00	800	256	128	128		
DB Recycle Cache	0.00	0	0	0	0		
Redo Buffers	0.77						
PGA Aggregate Target	64.00	48	192	96	96		
Total	**353.47**			**417.47**	**417.47**		

Max. SGA size :
513.47 MB
Unused SGA memory :
224 MB
Free Physical Memory :
0 MB
Granule size :
16 MB
Statistics level :
TYPICAL
Maximum Instance Memory :
420 MB
Time Recommended :
N/A

Start | Windows Task ... | Rptsched | C:\WINNT\Sys... | Spotlight | Quest Central ... | 2:40 PM

Spotlight on Oracle - mo920_64

File View Favorites Tools Help

mo920_64

Tuning

Performance | 9i Memory Management | Latch

Memory Configuration — Display all

Overall response time (estimated)

PGA Aggregate Target — Current Advisory - Response Time

Component	Current	Min.	Max.	Calculated	Custom	Locked	Status
Fixed	0.70						
⊞ Variable	144.00			144	144		
⊟ Database Buffers	144.00			176	176		
DB 2k Cache	0.00	0	0	0	0		
DB 4k Cache	0.00	0	0	0	0		
DB Cache (Default 8k)	128.00	24	25,600	48	48		
DB 16k Cache	0.00	0	0	0	0		
DB 32k Cache	0.00	0	0	0	0		
DB Keep Cache	16.00	800	256	128	128		
DB Recycle Cache	0.00	0	0	0	0		
Redo Buffers	0.77						
PGA Aggregate Target	64.00	48	192	96	96		
Total	**353.47**			**417.47**	**417.47**		

Max. SGA size :
513.47 MB
Unused SGA memory :
224 MB
Free Physical Memory :
0 MB
Granule size :
16 MB
Statistics level :
TYPICAL
Maximum Instance Memory :
420 MB
Time Recommended :
N/A

Start | Windows Task ... | Rptsched | C:\WINNT\Sys... | Spotlight | Quest Central ... | 2:40 PM

Spotlight on Oracle - mo920_64

File View Favorites Tools Help

mo920_64

Tuning

Performance | 9i Memory Management | Latch

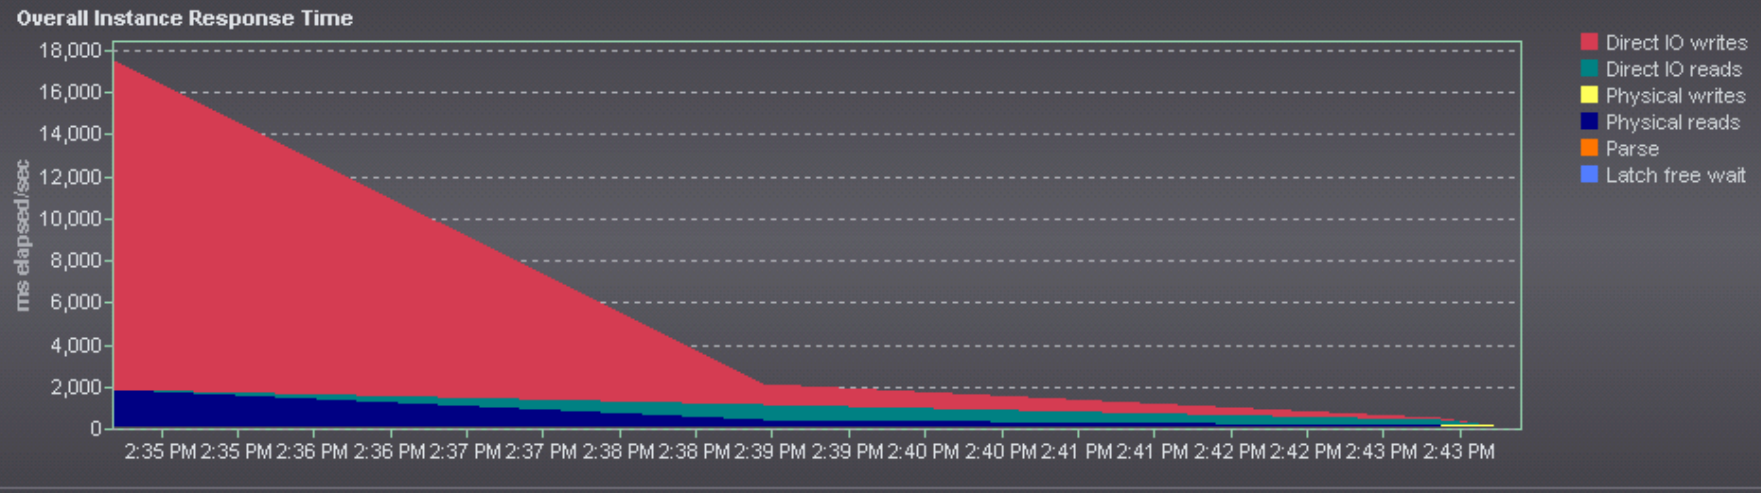

Implementation

Memory Updates

Component	Action	Initial size	Target size	Final size	Start time	End time	Status
DB Cache (Default 8k)	SHRINK	128	32	32	21/04/2003 2:43:46 PM	21/04/2003 2:43:53 PM	NORMAL
DB Keep Cache	GROW	16	128	128	21/04/2003 2:43:54 PM	21/04/2003 2:43:54 PM	NORMAL
PGA Aggregate Target		64	192				

Restart Stop Abort

Start | Windows Task ... | Rptsched | C:\WINNT\Sys... | Spotlight | Quest Central ... | 2:43 PM

THANK YOU
FOR LISTENING

Any Questions?

QUEST SOFTWARE